paloalto NETWORKS

TECH**DOCS**

IoT Security Administrator's Guide

February 2024

Contact Information

Corporate Headquarters:
Palo Alto Networks
3000 Tannery Way
Santa Clara, CA 95054
www.paloaltonetworks.com/company/contact-support

About the Documentation

- For the most recent version of this guide or for access to related documentation, visit the Technical Documentation portal docs.paloaltonetworks.com.

- To search for a specific topic, go to our search page docs.paloaltonetworks.com/search.html.

- Have feedback or questions for us? Leave a comment on any page in the portal, or write to us at documentation@paloaltonetworks.com.

Copyright

Palo Alto Networks, Inc.
www.paloaltonetworks.com

Last Revised

February 14, 2024

Table of Contents

IoT Security Solution

The architectural components that constitute the IoT Security solution are introduced here. Learn about the various components, how they work together, and how to set them up. Also learn about all the educational resources available for IoT Security.

- IoT Security Solution Structure
- IoT Security Solution Setup
- IoT Security Documentation

IoT Security Solution Structure

Using AI and machine learning, IoT Security automatically discovers and identifies all network-connected devices and constructs a data-rich, dynamically updating inventory. In addition to identifying IoT devices and IT devices (laptops and servers for example), IoT Security provides deep visibility into network behaviors, establishing what's normal and discerning what's suspicious. When it detects a device vulnerability or anomalous behavior posing a threat, IoT Security notifies administrators, who can then take action to investigate and remediate the issue.

To accomplish all this, the cloud-based IoT Security app works with Palo Alto Networks next-generation firewalls, logging service, and update server, and optionally with Panorama and integrated third-party products. These elements of the IoT Security solution collaborate to carry out the following tasks:

- Firewalls with IoT Security subscriptions collect information about network traffic and forward their logs to the logging service, which streams metadata to IoT Security for analysis.

- The update server provides firewalls and Panorama with a regularly updated device dictionary file of device attributes (profile, vendor, category, and so on) that Security policy rules use for device identification, or *Device-ID*.

- IoT Security recommends Security policy rules based on Device-ID to firewalls. When Panorama provides centralized firewall management, IoT Security works through it to recommend Security policy rules to managed firewalls. When Panorama is not in use, IoT Security interacts directly with firewalls.

- IoT Security maps IP addresses to devices and notifies firewalls of their corresponding device attributes so they can enforce Device-ID-based Security policy rules that reference attributes in IP address-to-device mappings.

With a third-party integrations add-on license for your IoT Security account, you are able to expand IoT Security capabilities to include product-specific features and those of the integrated products to include IoT.

Learn about the major components that constitute the IoT Security solution:

1 - Device Data Collection

2 - Data Analysis

1 - Device Data Collection

For IoT Security to identify IoT devices and establish a baseline of their acceptable network behaviors, it needs to analyze their network activity. That's where next-generation firewalls come in. They log network traffic to which they apply Security policy rules and then forward logs to the logging service where IoT Security accesses them. Depending on whether your IoT Security subscription includes data storage, the logging service either streams metadata to your IoT Security account and Strata Logging Service instance or just to your IoT Security account.

Various IoT devices generate different types of traffic on the network: ARP, DHCP, HTTPS, and other protocols.

If the firewall is in a position to receive this traffic, it generates logs when enforcing Security policy rules and then forwards the logs to the logging service.

The logging service streams network traffic metadata to IoT Security. Depending on the IoT Security subscription type, it can also store data in Cortex Data Lake.

Detailed Instructions

Onboard IoT Security

Prepare Your Firewall for IoT Security

2 - Data Analysis

IoT Security uses AI and machine-learning algorithms to analyze numerous aspects of the network behavior of a device and classify it within three levels or tiers. At the broadest tier, IoT Security identifies behavioral similarities that enable its algorithms to assign a device to a device category, such as security camera, even if it doesn't yet know the exact vendor and model. At the next tier, IoT Security gathers more granular behavioral attributes shared by certain vendors and models of security cameras to assign it a device profile. At the third tier, the algorithms create a model of unique behaviors for this individual security camera, such as its usage pattern.

In addition to device identification, IoT Security applies proprietary and supplemental machine-learning technologies to threat detection. It automatically detects device vulnerabilities and notifies IoT Security administrators. It also detects anomalous network behavior indicative of attack or reconnaissance and generates security alerts.

IoT Security uses AI and machine learning to identify devices, learn their normal behaviors, and detect anomalies.

Detailed Instructions

Introduction to IoT Security

Discover IoT Devices and Take Inventory

Detect IoT Device Vulnerabilities

Respond to IoT Security Alerts

3 - IoT Device Protection

IoT Security coordinates with next-generation firewalls to recommend Security policy rules for IoT device traffic. After identifying devices and establishing a baseline of acceptable network behavior, IoT Security automatically generates recommended Security policy rules for device profiles based on the network behavior it observes. Panorama or firewall administrators then import the recommendations to Panorama or directly to firewalls where they decide which ones to add to their policy set.

Firewalls and Panorama must have a list of device profiles or other device attributes for Device-ID-based Security policy rules. This list is provided as a device dictionary file from the update server, which firewalls and Panorama check regularly for updates to download.

So that firewalls apply imported Device-ID-based rules appropriately, IoT Security continually sends the firewall IP address-to-device mappings, which include the profile and other attributes of all devices monitored and protected by IoT Security.

IoT Security also integrates with Prisma Access to identify and secure devices.

Detailed Instructions

4 - Third-party Integrations

In addition to protecting IoT devices by coordinating with next-generation firewalls, IoT Security also integrates with third-party products to do the following:

- Increase device inventory and enrich device context—sometimes for IoT Security and sometimes for the integrated third-party product

- Broaden the coverage of specific features in integrated products to include IoT

- Expand the capabilities of IoT Security; for example, through integrations that allow you to do vulnerability scanning, quarantine devices with critical vulnerabilities or security alerts, and apply access control lists (ACLs) to IoT devices

IoT Security integrates with other products through a third-party integrations add-on, which is based on a Cortex XSOAR module.

Detailed Instructions

IoT Security Integration Guide

5 - Using Prisma Access instead of Next-generation Firewalls

When using IoT Security with Prisma Access, the process for collecting device data is similar to the previous description of data collection except that you substitute Prisma Access for firewalls. In addition, IoT Security can coordinate with Prisma SD-WAN ION devices to collect data at branch sites. When Prisma Access and SD-WAN forward data logs to the logging service, Strata Logging Service must be used.

Various IoT devices generate different types of traffic on the network: ARP, DHCP, HTTPS, and other protocols.

If Prisma Access is in a position to receive this traffic, it logs it when enforcing Security policy rules and then forwards its logs to Cortex Data Lake.

Cortex Data Lake stores the network traffic data it receives and streams metadata to IoT Security for analysis.

IoT Security sends Security policy rule recommendations through Panorama to Prisma Access. It sends IP address-to-device mappings to Prisma Access directly. Likewise, the update server sends device dictionary updates directly to Prisma Access as well as to Panorama.

Detailed Instructions

Prisma Access

IoT Security Integration with Prisma Access

IoT Security Integration Status with Prisma Access

Strata Logging Service

Prisma SD-WAN

IoT Security Solution Setup

The following is an overview of the main steps involved in setting up the IoT Security solution with particular focus on the following three components:

- Palo Alto Networks Next-Generation Firewalls with or without Panorama management
- Logging service with or without a Strata Logging Service instance
- IoT Security application

The solution also makes use of the update server for device dictionary file updates and the Customer Support Portal and hub for IoT Security user management. Optionally, IoT Security integrates with Prisma Access and SD-WAN and, through XSOAR, with third-party products.

Learn about the main steps involved in the IoT Security solution setup:

1 - Check Firewall Support and Prerequisites

2 - Onboard IoT Security

3 - Prepare Firewalls

4 - Install Certificates and Licenses

5 - Configure Logging

1 - Check Firewall Support and Prerequisites

Most current Palo Alto Networks firewall models support IoT Security with a few exceptions and with different degrees of functionality depending on the PAN-OS release:

- PAN-OS 8.1, PAN-OS 9.0, and PAN-OS 9.1: Device visibility and manually configured Security policy enforcement
- PAN-OS 10.0 or later: Device visibility and automated Security policy enforcement through Device-ID

Although IoT Security is a cloud application and is always running its latest software version, make sure the firewall models and PAN-OS versions on them support the level of functionality you want.

In addition, there are several prerequisites. For example, each firewall that integrates with IoT Security must have an IoT Security subscription. Not all firewalls on your network must subscribe to IoT Security; only those that collect network traffic and forward logs to it and those after PAN-OS 10.0 that receive policy rule recommendations and IP address-to-device mappings from it.

Customer Account

Cortex Data Lake instance
+
Premium Support license
or better

- - -

Premium Support license
or better

Firewall

Subscriptions

Check IoT Security support on your firewall models
- ⊘ PA-200
- ⊘ PA-3050
- ⊘ PA-5020
  ...

NGFW

IoT Security Subscription (with Cortex Data Lake)

or

IoT Security Subscription - Doesn't Require Data Lake

Recommendations
- Threat Prevention license (Highly recommended)
- DNS Security license
- URL Filtering license
- Enable SSL decryption

Check the PAN-OS release
- 8.1 - 9.0.2 with Panorama
- 9.0.3 or later with or without Panorama
- 8.1, 9.0, 9.1 - IoT visibility + manual policy enforcement
- 10.0 or later - Automatic policy enforcement through Device-ID

Medical IoT
Use application content version 8367-6513 or later

Prisma Access integration with IoT Security
Run Prisma Access 2.0-Innovation or later with the IoT Security add-on

Detailed Instructions

Firewall and PAN-OS Support of IoT Security

IoT Security Prerequisites

Medical IoT

IoT Security Integration with Prisma Access

2 - Onboard IoT Security

IoT Security onboarding is a six-step process that starts from an **Activate** link in an email from Palo Alto Networks. (If you have an Enterprise License Agreement, it starts either in the Customer Support Portal or in the hub). During the IoT Security onboarding process, do the following depending on what you're activating:

- Create an IoT Security tenant

- (IoT Security Subscription) Activate a new Strata Logging Service instance or associate an existing one with your IoT Security tenant

  or

  (IoT Security Subscription - Doesn't Require Data Lake) Specify the data ingestion region

- Subscribe firewalls to IoT Security services

- Optionally activate a third-party integrations add-on

Detailed Instructions

Onboard IoT Security

3 - Prepare Firewalls

For IoT Security to discover network-connected devices and assess their network behavior patterns, it needs quality network metadata from next-generation firewalls. Therefore, it's essential that firewalls are placed on the network and configured to collect metadata from traffic and forward it for IoT Security to access. In particular, DHCP traffic is important because it links dynamically assigned IP addresses to device MAC addresses, making them trackable over time.

Firewalls must also provide IoT Security with metadata for other types of traffic that devices generate. They do this by enforcing policy on network traffic, creating logs, and then forwarding them to the logging service, which then streams the metadata to IoT Security.

Detailed Instructions

Firewall Deployment for Device Visibility

DHCP Data Collection by Traffic Type

Firewall Deployment Options for IoT Security

4 - Install Certificates and Licenses

Logging service and device licenses permit next-generation firewalls to connect to the logging service and IoT Security. Logging service and device certificates authenticate these connections. Firewalls need these licenses and certificates to integrate with IoT Security.

Firewalls running PAN-OS 8.1–10.0 use logging service certificates to secure communications with the logging service so they can forward various logs to it. From PAN-OS 10.0, when Device-ID was introduced, firewalls use device certificates to secure communications with IoT Security to get IP address-to-device mappings and recommended policy rules. (Note: Panorama-managed firewalls can get recommended policy rules either directly from IoT Security or indirectly from IoT Security through Panorama.) From PAN-OS 10.1, firewalls use just one device certificate to secure connections to both the logging service and IoT Security. Panorama also uses a device certificate to secure communications with IoT Security.

Detailed Instructions

Onboard IoT Security

Prepare Your Firewall for IoT Security

Install a Device Certificate

Install the Panorama Device Certificate

5 - Configure Logging

Configure Security policy rules on firewalls to log traffic and forward logs to the logging service where IoT Security accesses it. The more network traffic metadata IoT Security has for analysis, the more quickly and confidently it identifies devices and establishes a baseline of their normal network behaviors. This results in a broader application of Security policy rules based on Device-ID (IoT Security sends firewalls IP address-to-device mappings only when it has a high confidence in their identities and the devices have sent or received traffic within the past hour) and broader and deeper insight into device risk and real and potential security threats.

Detailed Instructions

Prepare Your Firewall for IoT Security

Configure Policies for Log Forwarding

IoT Security Integration Status with Firewalls

IoT Security Documentation

Learn about the IoT Security technical documentation and training that's available.

IoT Security Documentation Set

The following set of technical documents in the Palo Alto Networks Tech Docs portal constitute the primary documentation for IoT Security.

IoT Security Best Practices	This reference recommends IoT Security best practices for the following main phases of deployment:
	Plan Your IoT Security Deployment Using Best Practices
	Deploy IoT Security Using Best Practices
	Monitor Your IoT Security Deployment Using Best Practices
IoT Security Administrator's Guide	The administrator's guide describes IoT Security features and explains how to configure and use them. Some chapters are about administering the IoT Security application:
	IoT Security Solution
	Get Started with IoT Security
	IoT Security Overview
	Manage IoT Security Users
	Other chapters explain how to work with device- and security-related data:
	Discover IoT Devices and Take Inventory
	Discover IoT Device Applications
	Detect IoT Device Vulnerabilities
	Respond to IoT Security Alerts
	Recommend Security Policies

	Medical IoT
IoT Security Integration Guide	This guide explains how to integrate IoT Security through Cortex XSOAR with third-party products. The guide provides configuration instructions for both sides of each integration. IoT Security supports integration with the following types of systems:
	Asset Management
	Endpoint Protection
	Network Management
	Wireless Network Controllers
	Security Information and Event Management
	Network Access Control
	Vulnerability Scanning
IoT Security API Reference	This reference provides explanations and examples of the IoT Security API and is divided into two parts: commonly used parameters and individual API requests and responses.

Useful Learning Resources

Palo Alto Networks also offers these resources to learn about IoT Security.

Device-ID in the PAN-OS Administrator's Guide	PAN-OS documentation describes how Device-ID works, how to prepare for its deployment, and how to configure and manage it. It also includes useful CLI commands for troubleshooting.
IoT Security Deployment Design Guide	This document captures typical deployment scenarios and recommendations for IoT Security.
Securing IoT Environments: Reference Architecture Guide	This guide is intended for solution architects and engineers and provides architectural guidance for deploying the IoT Security solution.

IoT Security privacy statement	The privacy statement contains information about how the IoT Security solution captures, processes, and stores information.
IoT Security webpage	This page on the Palo Alto Networks website provides an overview of the IoT Security product and various materials such as briefs, datasheets, reports, and case studies.
Knowledge base articles	Several articles answer common questions about IoT Security and how it works.
Monthly release notes	Release notes summarize new features and enhancements, changes in appearance and behavior, and known and addressed issues each month. They're available in the IoT Security portal by clicking the Help icon (?) in the lower right corner and selecting **Product Release**.
IoT Security digital online training	The training course enables you to describe the fundamentals of IoT Security, configure it, and prepare your firewall and logging service to work as part of the IoT Security solution.

20

Get Started with IoT Security

Learn how to onboard the IoT Security app and how to prepare your firewall to work with IoT Security.

- Firewall and PAN-OS Support of IoT Security
- IoT Security Prerequisites
- Onboard IoT Security
- Onboard IoT Security on VM-Series with Software NGFW Credits
- Firewall Deployment for Device Visibility
- DHCP Data Collection by Traffic Type
- Firewall Deployment Options for IoT Security
- Configure a Pre-PAN-OS 10.0 Firewall with a DHCP Server
- Configure a Pre-PAN-OS 10.0 Firewall for a Local DHCP Server
- Use a Tap Interface for DHCP Visibility
- Use a Virtual Wire Interface for DHCP Visibility
- Use SNMP Network Discovery to Learn about Devices from Switches
- Use ERSPAN to Send Mirrored Traffic through GRE Tunnels
- Use DHCP Server Logs to Increase Device Visibility
- Plan for Scaling when Your Firewall Serves DHCP
- Prepare Your Firewall for IoT Security
- Configure Policies for Log Forwarding
- Control Allowed Traffic for Onboarding Devices
- Support Isolated Network Segments
- IoT Security Integration with Prisma Access
- IoT Security Licenses
- Offboard IoT Security Subscriptions

Firewall and PAN-OS Support of IoT Security

For Palo Alto Networks next-generation firewalls running PAN-OS 8.1, PAN-OS 9.0, or PAN-OS 9.1, the IoT Security solution provides visibility of discovered IoT devices based on the logs it receives from the firewall. IoT Security also uses machine learning (ML) to identify vulnerabilities and assess risk in devices based on their network traffic behaviors and dynamically updated threat feeds. Although these PAN-OS versions don't support automated policy enforcement of IoT devices through the Device-ID™ framework, which is available from PAN-OS 10.0, you can still use the policy rule recommendations that IoT Security generates as a reference when manually adding rules to your firewalls. IoT Security always generates Security policy rule recommendations regardless of the PAN-OS version.

Firewalls running PAN-OS 10.0 or later automate policy enforcement through Device-ID. This is a mechanism that identifies devices by attributes such as device type, vendor, model, or operating system and then applies device-based policy rules to those with matching attributes.

All Palo Alto Networks next-generation firewalls running PAN-OS 10.0 or later fully support IoT Security with the following exceptions.

IoT device visibility and the manual application of policy recommendations but not Device-ID

- PA-200 with PAN-OS 8.1
- PA-500 with PAN-OS 8.1
- PA-3020 with PAN-OS 8.1, PAN-OS 9.0, or PAN-OS 9.1
- PA-3050 with PAN-OS 8.1, PAN-OS 9.0, or PAN-OS 9.1
- PA-3060 with PAN-OS 8.1, PAN-OS 9.0, or PAN-OS 9.1
- PA-5020 with PAN-OS 8.1
- PA-5050 with PAN-OS 8.1
- PA-5060 with PAN-OS 8.1

No IoT Security support

- CN-Series firewalls before PAN-OS 11.1
- VM-50
- VM-200

When choosing firewalls to subscribe to IoT Security services, consider the type of IoT Security functionality they support. Another factor to consider is when various firewall models will reach the end of sales and service support and when you plan to update them to newer models. However, even if you subscribe a firewall to IoT Security and then decide to retire it while its IoT Security license still has time remaining, you can transfer the license from that firewall to another one where IoT Security will continue to operate for the remainder of its subscription period.

IoT Security Prerequisites

Ensure that your environment meets all prerequisites for deploying IoT Security with Palo Alto Networks next-generation firewalls:

- One or more firewalls running PAN-OS 8.1 to PAN-OS 9.0.2 with Panorama management, or PAN-OS 9.0.3 or later with or without Panorama management.

  Firewalls running PAN-OS 8.1, PAN-OS 9.0, and PAN-OS 9.1 support IoT Security for device visibility and manual policy enforcement. Firewalls running PAN-OS 10.0 or later support IoT Security for both device visibility and automatic policy enforcement through Device-ID.

- One IoT Security license per firewall.

  The license controls whether IoT Security ingests log data that a firewall forwards to the Palo Alto Networks cloud-based logging service to identify IoT devices and assess risk. The license also controls whether a firewall can pull IP address-to-device mappings and policy rule recommendations from IoT Security and the device dictionary from the update server for use in its security policy rules.

  (A note about IP address-to-device mappings: IoT Security uses patented multi-tier machine-learning algorithms to profile device behaviors and identify the device type, make, model, OS, and OS version. It bundles this set of attributes into a logical object, maps it to the IP address of a device, and sends it to the firewall. This object is called an IP address-to-device mapping.)

  When you buy an IoT Security subscription, you have a 90-day grace period to activate the license on a firewall. If you activate it within the first 90 days, the subscription starts on the activation date. Otherwise, it starts 90 days after the purchase date.

  A Panorama management server does not require an IoT Security license.

- When using IoT Security Subscription, which stores data in Strata Logging Service, you need one Strata Logging Service license per account. (When using IoT Security, Doesn't Require Data Lake Subscription, you do not need a Strata Logging Service license.)

  Your Strata Logging Service subscription can either be new or an existing one, and the data lake can be in the Americas, European Union, or Asia-Pacific region. Regardless of the use of the data lake, firewalls stream logging data automatically and continuously to the IoT Security infrastructure where it is retained for varying periods of time based on data type. For details about data retention, see IoT/OT Security Privacy.

  For a new Strata Logging Service instance, figure out the amount of storage you'll need with the Cortex sizing calculator. When making your calculations, enter the number of firewalls with an IoT Security license and select IoT Security.

- Using the logging service requires a Premium Support license or better. This is required when using the logging service with either of the two IoT Security subscription types: IoT Security Subscription and IoT Security Subscription - Doesn't Require Data Lake. (A Premium Support license is automatically included with the purchase of a Strata Logging Service instance.)

- A Threat Prevention license is required for IoT Security to get all the traffic and threat logs necessary to fully assess risk and detect vulnerabilities.

- The following licenses and firewall capability provide additional value to IoT Security:

  - A DNS Security license helps IoT Security detect DNS-related threats and risks.

  - A Wildfire license enhances the detection of malware and file-related vulnerabilities.

  - A URL Filtering license controls the online content devices can access and how they can interact with it.

  - Enabling SSL decryption on the firewall improves the coverage and accuracy of device identification. It also helps IoT Security with risk assessment and threat detections.

- When using IoT Security on networks with medical equipment, make sure the application content version on your firewalls is 8367-6513 or later; that is, the major version, which is identified by the first four digits, is 8367 or above (8368, 8369, 8370, and so on), starting from 8367-6513. These versions include healthcare-specific applications that allow IoT Security to discover medical equipment and provide utilization data. They also allow firewall Security policy rules to include healthcare-specific applications.

- When integrating IoT Security with Prisma Access, Prisma Access must be running the Prisma Access 2.0-Innovation release or later with an IoT Security add-on. To learn about other requirements, see IoT Security Integration with Prisma Access.

- When Panorama manages firewalls running PAN-OS 10.2, it requires the 3.1 cloud services plugin.

Onboard IoT Security

Follow the onboarding workflow to create a URL for your IoT Security portal and activate IoT Security subscriptions for your firewalls. Through the onboarding process, you can optionally activate a Strata Logging Service instance to store data and a third-party integration add-on for IoT Security to expand its capabilities.

It is important to keep the IoT Security activation email you received from Palo Alto Networks. It not only contains confidential activation-related data but if you still have unused IoT Security licenses after completing the onboarding process, you can click the **Activate** button in the email again to repeat the process and activate more firewalls later.

(Enterprise License Agreement) When you have an Enterprise License Agreement (ELA), begin the activation process by entering the authorization code that Palo Alto Networks sends you in your Customer Support Portal account. For complete step-by-step instructions, see Activate an Add-on Enterprise License Agreement through Common Services.

When you have IoT Security subscriptions, the onboarding process consists of the following main steps.

STEP 1 | Click **Activate** in the IoT Security activation email from Palo Alto Networks.

STEP 2 | Log in to the Palo Alto Networks hub.

STEP 3 | Activate IoT Security.

STEP 4 | Add devices (firewalls) to the tenant service group (TSG) and associate IoT Security, and possibly other applications as well, with the firewalls.

STEP 5 | (Optional) Manage identity and access to IoT Security.

STEP 6 | Set up IoT Security and firewalls to work together.

For instructions for these first six steps, see Common Services: Subscription & Tenant Management. Then return here to continue the setup.

STEP 7 | Log in to the IoT Security portal.

Click the **IoT Security** link on either the Tenant Management or Device Associations page.

A welcome page appears displaying the status of the logging service and several links to useful learning resources.

Resource Center

Search devices, vulnerabiliti... 🔍

Search devices, alerts, vulnerabilities by queries   🔍 Search

Congratulations on activating IoT Security!

We have started monitoring your network and will discover new devices soon. This usually takes couple of hours depending on your network and IoT devices.

Here are the firewall logs status. View details.

EAL Logs

We are receiving logs from **5 of your 705** firewalls. Please wait 30 minutes and check again. If we still aren't receiving logs from one or more firewalls, check that they're properly configured to forward logs to the logging service.

DHCP Logs

We are receiving logs for DHCP traffic.

Traffic Logs

We are receiving logs from **4 of your 705** firewalls. Please wait 30 minutes and check again. If we still aren't receiving logs from one or more firewalls, check that they're properly configured to forward logs to the logging service.

Setup Checklist

Set up your firewalls to get the full benefits of IoT Security:

1.   Generate an OTP or PSK to onboard firewalls with IoT Security.   **Start**

2.   Deploy firewalls with visibility into DHCP and network traffic.   **Start**

3.   Configure the firewalls to work with IoT Security.   **Start**

Recommended Resources

Here are some selected tutorials and articles to help you start protecting the devices in your network!

▶   **IoT Security Overview**

Provide an overview of the challenges with securing the IoT devices, how the IoT Security solution addresses these challenges, key values it provides.

Useful Links

Knowledge Base
Customer Support
Hub

STEP 8 | To access the rest of the web interface, use the navigation menu on the left.

If you are a user with owner privileges and the portal doesn't have a predetermined vertical theme, IoT Security will prompt you to select a theme when you attempt to navigate away from the welcome page: Enterprise IoT Security Plus, Industrial OT Security, or Medical IoT Security. If you don't select a theme, you will use the Enterprise IoT Security Plus theme by

default. IoT Security will continue to prompt you to select a theme every time you log in until you make a selection, or another user with owner privileges does.

29

If you are a user without owner privileges and an owner hasn't yet selected a vertical theme, you will see the Enterprise IoT Security Plus theme by default. Otherwise, if the portal theme was already determined by the IoT Security product purchased or if an owner already set a theme, then that is the one you see.

There might not be any data in the portal when you first log in. Firewalls create network traffic data logs and forward them to the logging service, which streams them to the IoT Security Cloud. On average, devices begin showing up in the IoT Security portal within the first 30 minutes. Depending on the size of the network and the amount of activity of the devices on it, it can take several days for all the data to show up.

> *Click **Administration** > **Sites and Firewalls** > **Firewalls** in the IoT Security portal to see the status of logs that the logging service is streaming to the IoT Security app. For more information, see* IoT Security Integration Status with Firewalls

After the IoT Security portal has had time to use its machine-learning algorithms to analyze the network behavior of your IoT devices (1-2 days), consider following the typical workflow of an IoT Security user:

- Device visibility – Learn about the IoT devices on the network
- Application visibility – Learn about the applications and protocols these devices use
- Device vulnerabilities – Learn about IoT device vulnerabilities and take steps to mitigate them, first on the most critical devices and then on others
- Security alerts – Respond to security alerts as they occur, prioritizing your response on the urgency of the alert and the importance of the targeted device or network segment
- Security policy rule recommendations – Based on observed network behavior, the IoT Security app can generate recommended security policy rules that you can then sync with those on your next-generation firewall.

Depending on the PAN-OS versions running on your firewalls, you must generate an OTP or PSK and install certificates on firewalls so they will connect securely with the logging service and with IoT Security. There are also firewall configurations necessary to enable logging and log forwarding to IoT Security. For Enterprise IoT Security Plus, Industrial OT Security, and Medical IoT Security, you must also configure IoT Security and PAN-OS to apply Device-ID to enforce Security policy rules. To continue, see Prepare Your Firewall for IoT Security.

Onboard IoT Security on VM-Series with Software NGFW Credits

A Palo Alto Networks VM-Series is a virtualized form factor of a Palo Alto Networks next-generation firewall and is intended for use in a virtualized or cloud environment. When you use Software NGFW credits to fund VM-Series with either fixed or flexible virtual CPUs (vCPUs), you can include IoT Security in the deployment profile during the firewall registration process.
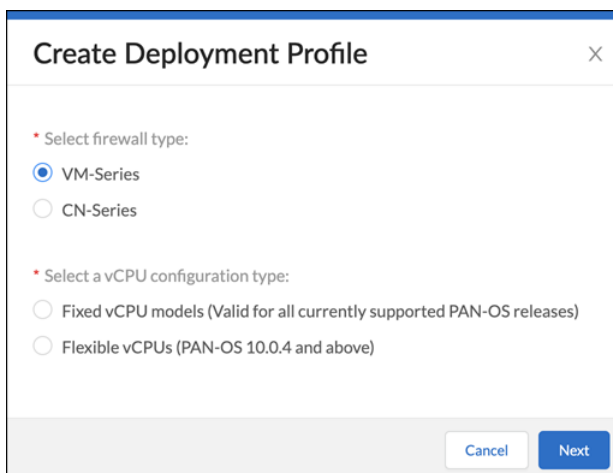
> *You can also use Software NGFW credits to fund CN-Series with an IoT Security subscription as long as the firewalls are under Panorama management. For onboarding instructions of a CN-Series with IoT Security, see* IoT Security Support for CN-Series.

The following onboarding procedure is for VM-Series with an IoT Security subscription. It assumes that you have already purchased Software NGFW credits and activated them. At this point, you can use the Software NGFW credits to purchase VM-Series.

STEP 1 | Create one or more deployment profiles for VM-Series.

Create a deployment profile for each type of VM-Series model you want to deploy.

1. Log in to the Customer Support Portal (CSP), and—if you have multiple accounts—choose the account you want to use.

2. Select **Products** > **Software NGFW Credits** to view the Software NGFW Credits Dashboard.

3. Locate your purchased NGFW Credits pool on the dashboard and **Create Deployment Profile**.

4. Select **VM Series** and either **Fixed vCPU models (Valid for all currently supported PAN-OS releases)** or **Flexible vCPUs (PAN-OS 10.0.4 and above)** and then click **Next**.

5. Assuming you selected **Fixed vCPU models (Valid for all currently supported PAN-OS releases)**, configure the following and then **Create Deployment Profile**:

Profile Name: Enter a name for the deployment profile.

Number of Firewalls: Enter the maximum number of firewalls that can be associated with this deployment profile.

Fixed vCPU model: Choose a VM-Series model from the list.

Security Use Case: Choose **Custom**.

Customize Subscriptions: Clear all preselected items and select **IOT**.

IOT Subscription: Choose the type of IoT Security subscription to activate on the VM-Series. The different types are based on vertical themes with or without traffic log retention in Strata Logging Service.

Use Credits to Enable VM Panorama: (clear all)

After creating the deployment profile, it appears in the Current Deployment Profiles table on the **Assets** > **Software NGFW Credits** page.

6. (Optional) After you click **Create Deployment Profile**, you can return to the configuration and click **Calculate Estimated Cost** to see an estimation of how many Flex credits will be deducted from your account and your remaining balance. If you hover your cursor over the question mark next to the estimate, you can see the credit breakdown for each component.

7. If you have other types of firewall models to deploy, create additional deployment profiles, one for each type.

STEP 2 | Activate IoT Security subscriptions based on the deployment profile in Common Services.

1. Log in to the hub with your Palo Alto Networks Customer Support credentials.

    The hub fetches available deployment profiles for this account from the CSP.

2. Select **Common Services** > **Subscriptions & Add-ons**.

    The deployment profile you created appears in the Ready for Activation section at the top of the page.

3. Click **Activate Now**.

   The Activate Subscriptions based on Deployment Profile(s) page appears.

4. Configure the following IoT Security subscription activation settings:

   Customer Support Account: Choose your CSP account with the deployment profile.

   Recipient: Use an existing tenant or create a new one.

   > *To create a new tenant, hover your cursor over **All Tenants** at the top of the Select Tenant drop-down list and then click the **Add** icon (**+**) that appears on the right. Enter a unique name for the tenant service group (TSG) and choose a business vertical.*

   Select Region: When activating an IoT Security subscription that doesn't require a Strata Logging Service, select the region where the logging service will ingest network traffic logs that the VM-Series send it for IoT Security to access and analyze.

   When activating an IoT Security subscription that does require a Strata Logging Service, you must first already have an activated Strata Logging Service instance in the same tenant service group (TSG). IoT Security will then use this instance by default. The TSG might already have another product with an activated Strata Logging Service (PA+CDL or AIOps +CDL for example), or you might have migrated an activated standalone Strata Logging Service instance to the TSG before activating the IoT Security subscription. In either case, the region will be automatically populated based on the region of the existing data lake in the TSG.

   Select Deployment Profile(s): Select the deployment profile you previously created.

   There are two sections for deployment profiles: **Available** and **Unavailable**. Deployment profiles appear in the Unavailable section if a required component is missing. For example, if the IoT Security subscription in the deployment profile requires a Strata Logging Service but the tenant service group (TSG) doesn't have one, the deployment profile will be in the

Unavailable section. You will need to activate the required Strata Logging Service before attempting to activate IoT Security in such scenarios.

📋 *When you create multiple deployment profiles, it's possible that they have different IoT Security subscriptions. When using them in the same IoT tenant, the IoT Security subscription type in the first deployment profile takes precedence over others added afterward.*

Configure Subscription URL(s): Enter a unique subdomain to complete the <subdomain>.iot.paloaltonetworks.com URL for your IoT Security application. This will be the URL where you log in to the IoT Security portal.

5. **Agree to the Terms and Conditions** and then **Activate**.

The hub displays the Tenant Management page where you can see the IoT Security initialization status for the TSG. The initialization generally takes a few minutes to complete.

STEP 3 | Associate firewalls through the deployment profile with the IoT Security subscription in the TSG.

1. Register a VM-Series using one of the two methods described in Register the VM-Series (Software NGFW Credits) and then **Submit** the registration.

    📋 *When registering a VM-Series that cannot access the CSP, you must enter a UUID, a CPU ID, the number of vCPUs on the firewall, and the amount memory allocated to the firewall. This information is in the General Information section on the **Dashboard** page of the web interface on your firewall. You can copy it from there and paste it in the Register Firewall form. You can also download this information from the firewall web interface to a text file by selecting **Device > Licenses > Activate Feature using Auth Code > Download Authorization File**. Then on the Register Firewall page in the CSP, **Upload a File for UUID & CPUID**.*

    After you submit the firewall registration, the CSP associates this firewall through the deployment profile with the TSG. It typically takes a few minutes for the registration and association to complete. When completed, you can see the firewall on the **Common Services** > **Device Associations** tab in the hub.

    During the firewall registration, the number of Software NGFW credits needed to fund the virtual firewall are automatically deducted from your pool of credits.

2. Associate more firewalls to the TSG through the same deployment profile or, if they are different types of firewall models, through other deployment profiles you have created for them.

    📋 *It's not currently possible to extend, renew, or offboard IoT Security licenses that have been activated on VM-Series funded by Software NGFW credits. In addition, Enterprise License Agreements (ELA) and IoT Security FedRAMP Moderate licenses are not supported.*
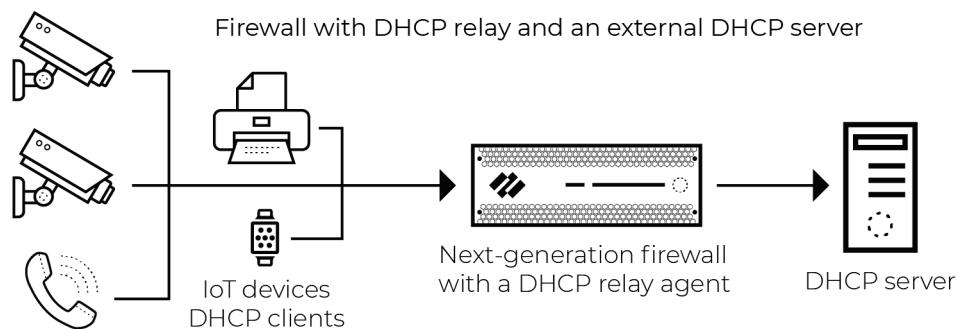
STEP 4 | Configure the VM-Series to provide network traffic logs with IoT Security.

Now that you've onboarded IoT Security onto your VM-Series, follow the steps in Prepare Your Firewall for IoT Security to configure it to log network traffic and forward the traffic logs to the logging service, which then streams network traffic metadata to IoT Security for analysis.
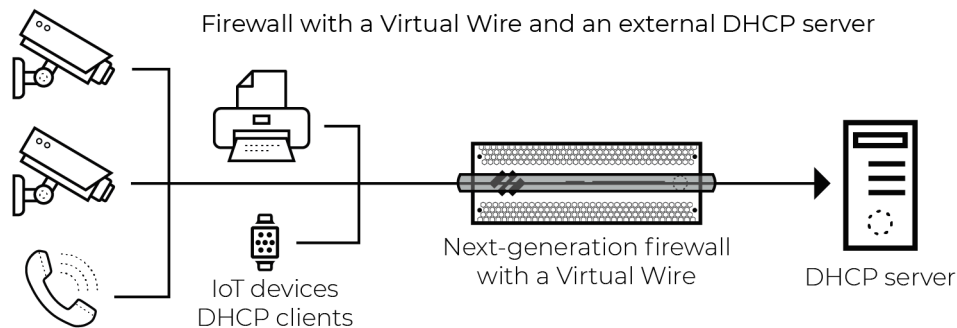
Firewall Deployment for Device Visibility

The Palo Alto Networks IoT Security app uses machine learning to classify IoT devices based on the network traffic for which these devices are either a source or destination. To accomplish this, it relies on Enhanced Application logs (EALs) generated by the Palo Alto Networks next-generation firewall.

DHCP traffic is of particular importance to the IoT security solution. DHCP provides a way to create an IP address-to-device mapping (that is, an IP address-to-MAC address mapping) that is required for classification to take place. However, a firewall typically only generates an EAL entry when it receives a unicast DHCP message; for example, when there is centralized Internet Protocol address management (IPAM) and either the firewall or another local device acts as a DHCP relay agent. Below is an example architecture that illustrates a common case where the firewall generates EALs for unicast DHCP traffic.

The firewall generates an EAL entry for broadcast DHCP traffic when the packet is seen on a virtual wire (vWire) interface with multicast firewalling enabled, as shown below.

DHCP Data Collection by Traffic Type

The tables below show Enhanced Application log (EAL) coverage when the firewall interface receiving unicast and broadcast DHCP traffic is in different modes.

Unicast DHCP Traffic

Firewall Interface Deployment Mode	DHCP EAL Generated
Virtual Wire	Yes

Firewall Interface Deployment Mode	DHCP EAL Generated
Tap	Yes
Layer 2	Yes
Layer 3	Yes

Broadcast DHCP Traffic

Firewall Interface Deployment Mode	DHCP EAL Generated
Virtual Wire	Yes
Tap	No
Layer 2	No
Layer 3	No
DHCP server on the firewall (L3, L2 with VLAN interface)	Yes*
DHCP relay agent on the firewall (L3, L2 with VLAN interface)	Yes

* The method for generating EALs when the firewall is the DHCP server is dependent on its PAN-OS version:

- A firewall running a PAN-OS 10.0 release or later natively generates EALs when a DHCP server is configured on an interface, **DHCP Broadcast Session** is enabled, and there's a Security policy rule that allows DHCP traffic to reach the server and has EAL forwarding enabled. For more information, see Prepare Your Firewall for IoT Security and Configure Policies for Log Forwarding.

- A firewall running a PAN-OS 8.1 - 9.1 release requires a configuration-only workaround to generate DHCP EALs when a DHCP server is configured on one of the firewall interfaces. For more information, see Configure a Pre-PAN-OS 10.0 Firewall with a DHCP Server.
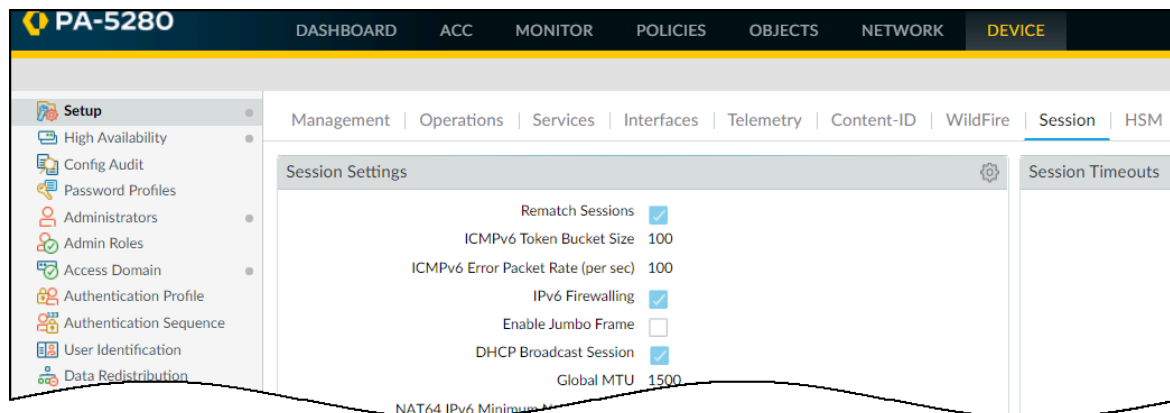
Firewall Deployment Options for IoT Security

When assessing deployment options for IoT device visibility, there are two fundamental considerations:

- The firewall must see traffic for the IoT application to use network traffic data for classification and analysis and for the enforcement of policy rules on the firewall itself. This includes regular operational traffic in addition to DHCP traffic.

- With the exceptions outlined below, the firewall must see unicast DHCP traffic to generate the data that allows IoT Security to create the required IP address-to-device mappings.

Exceptions to the Unicast Rule

- Virtual Wire: When the firewall has Virtual Wire interfaces with multicast firewalling enabled, it generates Enhanced Application logs (EALs) for broadcast DHCP sessions.

- A DHCP server is configured on the firewall:

  - PAN-OS 8.1 - 9.1: A configuration-only workaround is required for the firewall to generate EALs when a DHCP server on one of its interfaces receives broadcast DHCP traffic.

  - PAN-OS 10.0 and later: No workaround is required for the firewall to generate EALs when a DHCP server on one of its interfaces receives broadcast DHCP traffic. Just enable **DHCP Broadcast Session** at **Device** > **Setup** > **Session**.

  When the firewall receives DHCP broadcast traffic and applies a policy rule with an Enhanced Application log forwarding profile, it logs the DHCP traffic and forwards it to the logging service. From there, IoT Security accesses the data for analysis.

- A DHCP relay agent is configured on the firewall:

  - The firewall generates EALs for broadcast DHCP traffic when a DHCP relay agent is configured on one of its interfaces.

Tap Interfaces

Considerations – If you use a Tap interface to gain visibility into DHCP traffic that the firewall doesn't ordinarily see, consider the following:

- Place the tap "north" of any routed boundary where DHCP is configured. This will ensure that the captured traffic is unicast rather than broadcast. (If the firewall with the Tap interface is in the same broadcast domain as the switch that's mirroring traffic to it, enable **DHCP Broadcast Session** at **Device** > **Setup** > **Session**.)

- If adding a tap interface to an existing firewall, consider the available capacity on the firewall before implementing. While tap interfaces don't forward traffic, traffic seen on the tap port still consumes resources for processes such as the session table and packet buffers. For guidance on mitigating performance impact, see Use a Tap Interface for DHCP Visibility.

Use Cases for Tap interfaces

- Evaluations

- Networks where DHCP is configured on a device "south" of the firewall

- Monitor networks that don't naturally traverse the firewall

Virtual Wire Interfaces

Considerations – You might have to use a Virtual Wire (vWire) interface on the firewall to gain visibility into DHCP traffic that the firewall wouldn't normally see. Consider the following when using a tap interface in this manner:

• Ensure the Virtual Wire has multicast firewalling enabled.

• Ensure the Virtual Wire is in the path for DHCP traffic. This traffic can be either broadcast or unicast.

• Ensure that a security policy rule allowing DHCP exists and that a proper log-forwarding profile is applied to the rule.

• Ensure the firewall has the available capacity to process the additional traffic. For guidance on mitigating performance impact, see Use a Virtual Wire Interface for DHCP Visibility.

Use Case for Virtual Wire interfaces – When the DHCP server and the firewall interface are on the same network segment, the firewall sees only broadcast DHCP traffic. Placing the DHCP server behind a Virtual Wire interface enables the firewall to create EALs for this broadcast traffic.

Layer 2 and Layer 3 Interfaces

Considerations – Layer 2 (L2) and Layer 3 (L3) deployments both require unicast DHCP traffic to generate EALs. When using a VLAN interface in an L2 deployment, the considerations are the same as a deployment using Layer 3 interfaces:
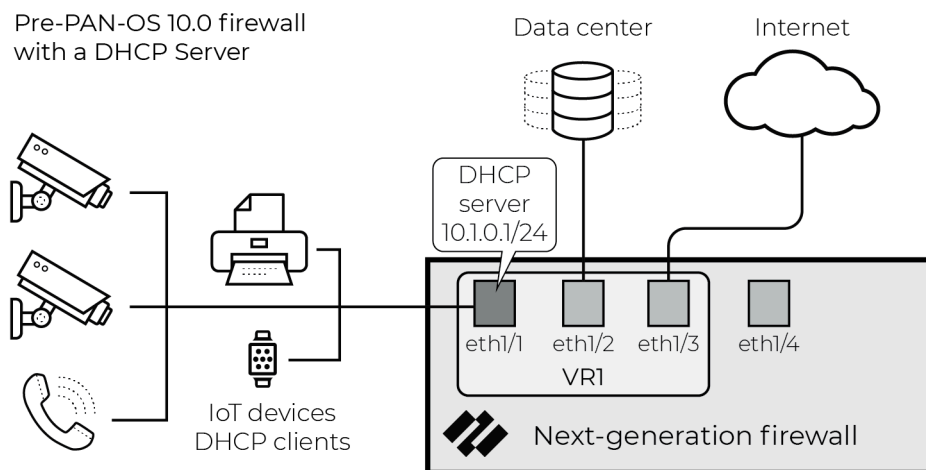
• Unicast DHCP packets traversing the firewall generate an EAL.

• When an L3 or VLAN interface is configured as a DHCP relay agent, the firewall generates an EAL.

• When an L3 or VLAN interface is configured as a DHCP server the firewall might generate an EAL. For more information, see DHCP Data Collection by Traffic Type in Firewall Deployment for Device Visibility.

Configure a Pre-PAN-OS 10.0 Firewall with a DHCP Server

The primary challenge is that PAN-OS versions before 10.0 do not generate Enhanced Application logs (EALs) when the firewall is the DHCP server, which is common in branch office and retail use cases. When the firewall is also the DHCP server, some reconfiguration of the firewall is required to generate EALs for DHCP traffic. You can do this by introducing a DHCP relay agent into its configuration.
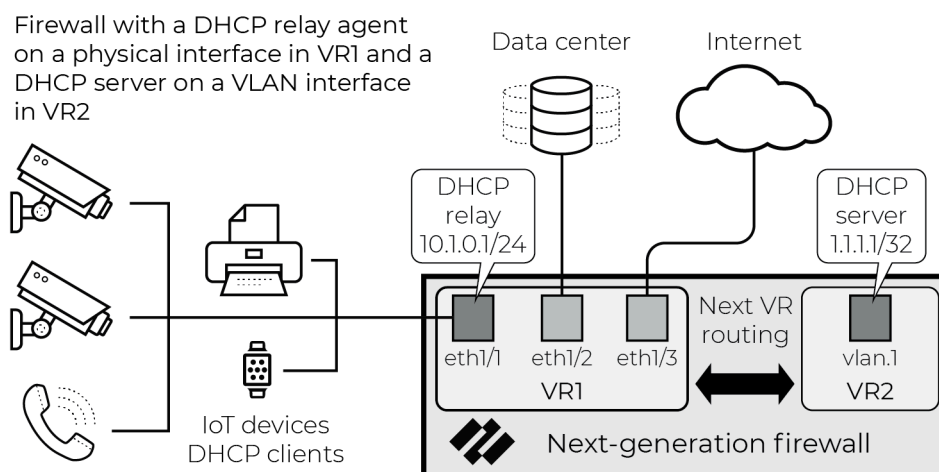
> *For the rest of this section on DHCP visibility, the firewall is assumed to be running a version of PAN-OS 9.1 or earlier.*

Solution: Configure a DHCP Relay Agent on a Physical Interface and a DHCP Serer on a VLAN Interface

Add a DHCP relay agent on the firewall so that unicast DHCP messages go through content scanning and the firewall generates EAL entries for them. Create a VLAN interface on the firewall to host a DHCP server and configure the physical interface of the firewall as a DHCP relay agent.

Analysis

When clients in the diagram above broadcast DHCPDISCOVER messages, the DHCP relay agent configured on ethernet1/1 receives them. You configure the relay agent to unicast the DHCPDISCOVER messages to the IP address of the vlan.1 interface which hosts a DHCP server. Note the following points:

- The vlan.1 interface can have an IP address with a 32-bit netmask to use address space efficiently when scaling this solution beyond one physical interface.

- The vlan.1 interface is in a separate virtual router. This forces the unicast DHCP messages to go through the data plane, which triggers the firewall to generate EAL entries.

- The DHCP server is configured with IP pools consistent with the subnet configured on ethernet1/1.

- You use Next-vr host routes to route unicast DHCP messages between ethernet1/1 and vlan.1.

Because this solution uses a virtual interface for the DHCP server, it can be implemented through configuration only without the need to physically reconfigure the network. Additionally, it can be implemented even when all the physical interfaces are in use.
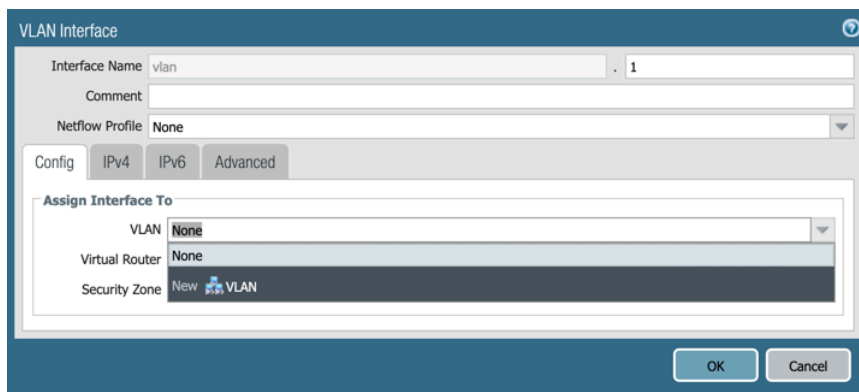
Configuration

**STEP 1 |** Save a snapshot of the current configuration.

STEP 2 | Configure a new virtual router.

STEP 3 | Configure a VLAN interface. In the VLAN drop-down list, click **New** to create a new VLAN.

STEP 4 | Enter a name for the new VLAN and then click **OK**.

The VLAN Interface configuration window appears.

STEP 5 | In the Assign Interface To section on the Config tab, select the virtual router you just created and the same security zone that the existing DHCP server is configured on.

If you choose a different zone, or create a new one, you must configure a security policy rule that allows DHCP between the two zones (see Configure an Interzone Policy in Configure Policies for Log Forwarding).

STEP 6 | Enable log forwarding.

Log forwarding enables the firewall to send enhanced application logs to the logging service. IoT Security then ingests metadata from there for analysis.

STEP 7 | If you use the same security zone, remember to enable logging and log forwarding for the intrazone policy rule.

For more information, see Configure an Intrazone Policy in Configure Policies for Log Forwarding.

STEP 8 | On the IPv4 tab, configure a host IP address—that is, an address with a 32-bit netmask—and then click **OK**.

> 📋 *For testing and troubleshooting purposes, assign an interface management profile that allows the VLAN interface to respond to pings. If the VLAN interface and physical interface are in different zones, see details in Configure an Interzone Policy in* Configure Policies for Log Forwarding*.*

STEP 9 | Open the existing virtual router and configure a host route to the IP address assigned to the VLAN interface configured above.


> 📋 *When there are multiple DHCP servers, replace the host route with a network route to simplify the configuration. For details, see* Plan for Scaling when Your Firewall Serves DHCP*.*

STEP 10 | Leave Interface set as **None** and select **Next VR** as the next hop. In the drop-down list below Next Hop, select the new virtual router you created.

STEP 11 | Click **OK** in the Static Route dialog box and then click **OK** in the Virtual Router dialog box.

STEP 12 | Open the new virtual router and configure a route to the network that the DHCP server serves.

The configuration is similar to that shown below where the Next Hop settings are Next VR and the name of the existing virtual router.

> 📋 *Creating a network route rather than a host route to the DHCP relay agent enables the probe feature of the DHCP server to function.*

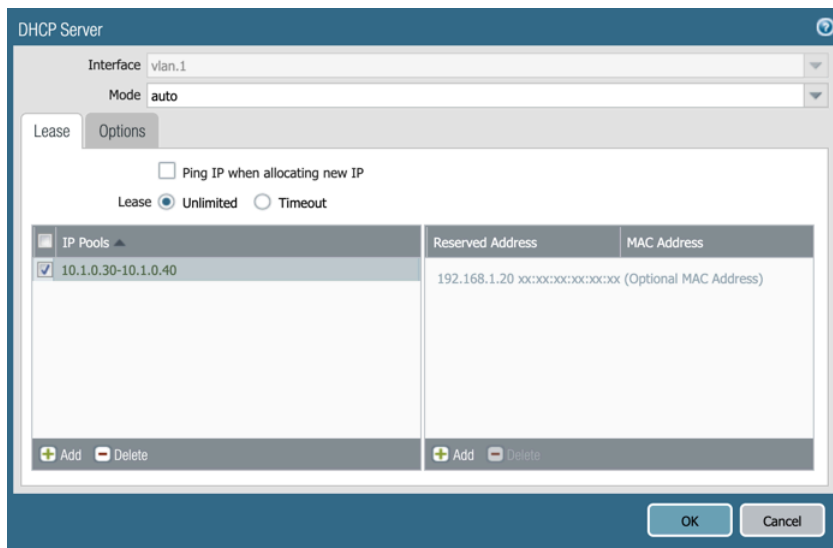**STEP 13 |** Commit these changes.

STEP 14 | Test your configuration.

If you assigned an interface management profile allowing ping to the VLAN interface, test your configuration by logging into the CLI and pinging from the physical interface to the VLAN interface:

```
ping source <phy_interf_ip-addr> host <vlan_interf_ip-addr>
```

STEP 15 | Configure a DHCP server on the VLAN interface.

Include the appropriate IP pools and options such as gateway and DNS servers and then click
OK.

STEP 16 | Configure a DHCP relay agent on the physical interface that connects to the local network
and then click **OK**.


STEP 17 | Commit the configuration.

STEP 18 | Test DHCP release and renew functionality by connecting a client to the local network
segment.

Configure a Pre-PAN-OS 10.0 Firewall for a Local DHCP Server

When the firewall is not receiving unicast DHCP packets—either as a DHCP server or relay agent
—you must arrange for it to generate or receive them. Instructions for doing this to provide DHCP
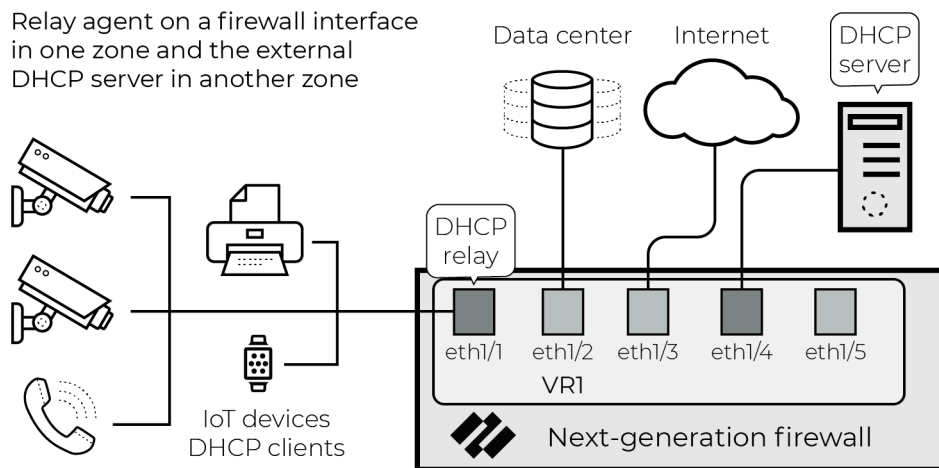traffic visibility in PAN-OS 8.1, PAN-OS 9.0, and PAN-OS 9.1 are provided in this section.

49

In the example below, there is a DHCP server on the local network segment. The firewall receives the DHCPDISCOVER messages that DHCP clients broadcast, but it is not configured as a DHCP server.

External DHCP server and pre-PAN-OS 10.0 firewall

Solution 1: Move the DHCP Server to a Different Zone

This solution involves moving the DHCP server to a different zone on the firewall and configuring a DHCP relay agent on the firewall interface that connects to the clients. This forces the generation of unicast DHCP traffic, which the firewall can then use to generate Enhanced Application logs (EALs).

Solution 2: Place the DHCP Server behind a Virtual Wire

Placing the DHCP server behind a Virtual Wire interface enables the firewall to generate EALs for all packets in the exchange. After proper configuration and physical network changes, the network looks similar to the illustration below:

DHCP exchanges between clients and server cross a Virtual Wire on the firewall

Use a Tap Interface for DHCP Visibility

To gain complete visibility of DHCP traffic, deploy a Tap interface on the firewall. This guide assumes familiarity with PAN-OS configuration, including Tap configuration. For details on configuring Tap interfaces, see the PAN-OS Networking Administrator's Guide.
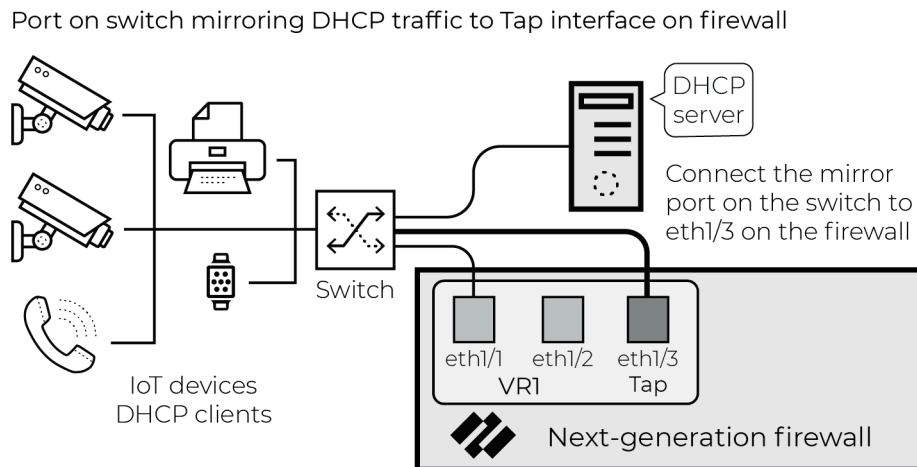
Considerations

Sending additional traffic to a Tap interface on the firewall results in additional session load. There are two causes for this:

- Any flow from the DHCP server to the internet, data center, or some other destination that would normally cross the firewall is inspected twice.

- Flows that normally would not be inspected are inspected when the Tap interface receives them; for example, flows bound for other hosts on the local network segment.

The following configuration section includes options for minimizing performance impact.

Network Architecture

The figure below illustrates the general idea of this solution. The actual topology can vary depending on the location of the DHCP server and the use of technologies such as RSPAN (Remote Switched Port Analyzer).

Port on switch mirroring DHCP traffic to Tap interface on firewall

The purpose of this configuration is to gain visibility into DHCP traffic that the firewall wouldn't normally see based on its current configuration and network topology.

Configuration

STEP 1 | Configure a Tap interface and zone.

Interface	Interface Type	Link State	IP Address	Virtual Router	Tag	VLAN / Virtual-Wire	Security Zone	Comment
ethernet1/1	Layer3		10.1.0.1/24	default	Untagged	none	Users	Local_Net
ethernet1/2	Layer3		172.16.1.2/30	default	Untagged	none	IT_Infra	Corp_Connection
ethernet1/3	Tap		none	none		none	DHCP_Tap	DHCP_Tap

STEP 2 | Configure policy rules for Tap traffic.

	Name	Type	Source Zone	Source Address	Destination Zone	Destination Address	Application	Service	Action	Profile	Options
1	Allow_DHCP_Tap	universal	DHCP_Tap	any	any	any	dhcp	application-d...	Allow	none	
2	Drop_Tap	universal	DHCP_Tap	any	any	any	any	application-d...	Drop	none	

- The first policy rule matches DHCP traffic and uses the same log forwarding profile that the rest of the rule base uses.

- The second rule drops all other traffic, minimizing additional session load on the firewall. Log forwarding profile is not enabled.

- Neither of the rules use security profiles.

**STEP 3 |** Connect the Tap interface to the port mirror on the switch.

Use a Virtual Wire Interface for DHCP Visibility

To gain complete visibility of DHCP traffic, deploy a Virtual Wire (vWire) in front of the DHCP server. This guide assumes familiarity with PAN-OS configuration, including Virtual Wire configuration. For details on configuring Virtual Wire interfaces, see the PAN-OS Networking Administrator's Guide.
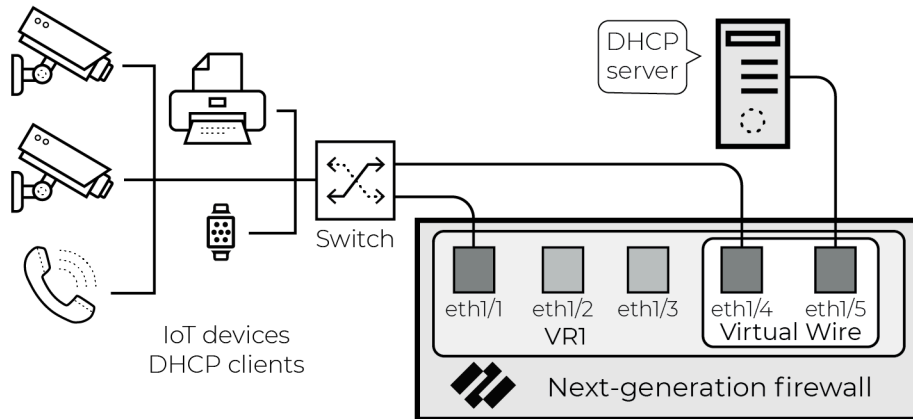
Network Architecture

This solution is for networks where a DHCP server is on the same network segment as the firewall interface, as shown in the figure below.

External DHCP server and pre-PAN-OS 10.0 firewall

For full visibility of all four DHCP messages, place the DHCP server behind a Virtual Wire interface. Doing so enables the firewall to generate Enhanced Application logs (EALs) for all packets in the exchange. After proper configuration and physical network changes, the network looks similar to the following illustration:

External DHCP server and firewall with a Virtual Wire on the network segment

Configuration

STEP 1 | Configure a Virtual Wire interface, complete with zones.

The configuration of the Virtual Wire object must include multicast firewalling:

STEP 2 | Configure a policy rule to allow traffic between the two Virtual Wire interface zones.

Configure this policy rule to allow all the existing traffic that the server currently sees use the same log forwarding object as the rest of the rule base. The Policy Optimization section below covers optimizing this policy rule set and preventing double logging.

			Source		Destination						
	Name	Type	Zone	Address	Zone	Address	Application	Service	Action	Profile	Options
1	DHCP_Host_Allow	universal	DHCP_Network_Side DHCP_Server_Side	any	DHCP_Network_Side DHCP_Server_Side	any	any	application-d...	Allow	none	

STEP 3 | Connect the external DHCP server to one side of the Virtual Wire and connect the network switch to the other side.

Instead of connecting the DHCP server host directly to the firewall, you can use an isolated VLAN to minimize cabling in the switching infrastructure.

Policy Optimization

The goal of this solution is to gain visibility into DHCP payloads while minimizing performance impact on the firewall. To that end, configure the following policy rule set for the Virtual Wire zones:

			Source		Destination						
	Name	Type	Zone	Address	Zone	Address	Application	Service	Action	Profile	Options
1	DHCP_Traffic	universal	DHCP_Network_Side DHCP_Server_Side	any	DHCP_Network_Side DHCP_Server_Side	any	dhcp	application-d...	Allow	none	
2	DHCP Ping	universal	DHCP_Server_Side	any	DHCP_Network_Side	any	ping	application-d...	Allow	none	
3	DHCP_Host_Allow	universal	DHCP_Network_Side DHCP_Server_Side	any	DHCP_Network_Side DHCP_Server_Side	any	any	application-d...	Allow	none	

- The "DHCP_Traffic" policy rule allows DHCP to and from the DHCP server. This rule uses the standard log forwarding profile with EALs enabled.
- The "DHCP Ping" policy rule allows pings from the DHCP server to the rest of the subnet. This enables DHCP servers to check if an IP address is active before assigning it as a lease to a new request. This rule does not forward logs.
- The "DHCP_Host_Allow" policy rule allows everything else to and from the DHCP server and does not forward logs for traffic matches.

To minimize the performance impact of the additional sessions that the firewall sees as a result of this Virtual Wire configuration, security profiles are not assigned to the above policy rules. If you want to microsegment the DHCP server, replace the "DHCP_Host_Allow" rule with a more granular policy rule set that allows applications in accordance with best practices. You can use security profiles in that policy rule set.

Use SNMP Network Discovery to Learn about Devices from Switches

To identify devices, assess risk, and help next-generation firewalls enforce security policy rules based on Device-ID, IoT Security requires network traffic metadata for analysis. Next-generation firewalls extract and log this metadata when they apply security policy rules that have logging enabled. When the rules also have log forwarding enabled, the firewalls send the logs to the logging service, which then streams the metadata to IoT Security.

However, depending on where the firewalls are placed, they might not have visibility into all network traffic, resulting in device discovery gaps and lower efficacy in device identification, behavior monitoring, and Device-ID rule enforcement. To extend visibility further into the network, IoT Security supports several options:

- Mirror traffic on network switches and use Encapsulated Remote Switched Port Analyzer (ERSPAN) to send mirrored traffic through GRE tunnels to a firewall. The firewall inspects the traffic, logs it, and then forwards the logs to the logging service for IoT Security to access.

- Configure a DHCP server to send its server logs as syslog messages to a firewall. The firewall then forwards the messages as Enhanced Application Logs (EALs) with a subtype of dhcp-syslog through the logging service to IoT Security.

- Integrate IoT Security with third-party products that provide services such as asset management and network management. IoT Security connects to these systems through Cortex XSOAR and retrieves additional device data from them to enhance the metadata learned from next-generation firewalls and optionally from network switches and DHCP servers.

In environments using DHCP to assign devices with network settings, IP addresses are leased dynamically for limited periods of time. An essential part of monitoring network behaviors to identify devices, assess risk, and enforce Device-ID security policy rules is the ability to link the dynamically assigned IP address of each device to its unique, unchanging MAC address. Next-generation firewalls can do this when they receive traffic containing both IP and MAC addresses. When firewalls don't receive traffic from all devices or when they do but it contains only IP addresses—possibly because the traffic crossed Layer 2 domains and the device MAC address was changed to that of the forwarding device—they can still gather IP address-to-MAC address bindings by using SNMP to query switches throughout the network.

When using SNMP to query network switches and other forwarding devices, firewalls first develop a network topography by requesting the Link Layer Discovery Protocol (LLDP) neighbors and Cisco Discovery Protocol (CDP) neighbors of one switch (the entry point switch) and then repeating the request with neighboring switches and child switches one by one throughout the network. After obtaining a list of switches and forwarding devices throughout the network, or within a limited area of the network, the firewall next queries each one for its ARP table as well as other information. The ARP table contains the IP address-to-MAC address binding information for the devices connected through the switch to the network. Other device details for which firewalls query include the physical interfaces or ports on the switch to which devices connect, their VLANs and subnets, and DHCP and DNS server IP addresses. After the firewall receives this information, it creates logs and sends them through the logging service to IoT Security.

The following are sample object identifiers (OIDs) that SNMP queries on UDP port 161 for information about LLDP neighbors and CDP neighbors, device IP address-to-MAC address bindings, and interface or port information:

- OID: 1.0.8802.1.1.2.1.4 lldpRemoteSystemsData (LLDP neighbors)

- OID: 1.3.6.1.4.1.9.9.23 ciscoCdpMIB (CDP neighbors)

- OID: 1.3.6.1.2.1.4.22.1.2 ipNetToMediaPhysAddress (IP-to-MAC address bindings from ARP)

- OID: 1.3.6.1.2.1.4.22.1.1 ipNetToMediaIfIndex (Interface or port information)

IoT Security provides SNMP Network Discovery as part of the IoT Security Third-party Integrations Add-on license, which must be purchased. From PAN-OS 11.1, SNMP network discovery is available to next-generation firewalls as a free plugin and does not require the add-

on license. While the version using the add-on license supports multiple sets of jobs for different networks and network segments per IoT Security tenant, the version with the free plugin supports just one set for one network or network segment per firewall.

> *The SNMP network discovery process cannot traverse switches that don't support CDP or LLDP.*

STEP 1 | Log in to the web interface of your firewall or Panorama and install the SNMP network discovery plugin.

This plugin allows a firewall to send SNMP queries to switches and routers on the network and then process the responses it receives.

Next-generation firewall

Select **Device** > **Plugins**, search for `network_discovery`, click **Download** in the Actions column, and then **Install** the plugin on the firewall.

Panorama

1. Select **Panorama** > **Plugins** , search for `network_discovery`, click **Download** in the Actions column, and then **Install** the plugin on Panorama.

2. Select **Panorama** > **Device Deployment** > **Plugins**, click **Install** in the Actions column, select the firewalls on which to install the plugin, and then click **OK**.

STEP 2 | Configure SNMP network discovery parameters.

The following instructions are for the SNMP network discovery configuration using the PAN-OS web interface on an individual next-generation firewall. To configure SNMP network discovery on Panorama, use templates and template stacks, and template stack variables for the IP addresses of the entry switch, discovery scope, and interfaces as needed.

1. Select **Device** > **IoT Security** > **Network Discovery** and then click **Edit** (gear icon).

    The SNMP Network Discovery Settings dialog box appears with the Schedule Settings tab active.

2. In the Network Discovery Job section, schedule how often the firewall runs a job to learn all the switches and other network forwarding devices that run LLDP and CDP on the network or within a defined scope of the network. The default is once a day, which usually is often enough.

3. In the Network Data Refreshment Job section, schedule how often the firewall runs a job to query switches and other forwarding devices for information about the network and devices connected to them. Consider how often DHCP lease times renew and schedule the job to run at half the lease time, which is when DHCP clients start requesting lease renewals

and could receive different IP addresses. In environments without DHCP, consider running the network data refreshment job once every hour, which is the default setting.

4. Click the **Discovery Scope Settings** tab, and enter the following:

Entry Point switch: Enter the IP address of the entry point switch with which to begin the SNMP discovery process.

A good choice for the entry point switch is a core switch because it would commonly have the broadest access to various distribution-layer and access-layer switches throughout the network.

Device IP Address Scope: Enter the prefix for the IP CIDR block to define the scope of the switches and endpoint devices to learn. Optionally, don't set a scope by entering **None** and SNMP will collect network topology for the entire network.

Service Route: If your firewall uses a data interface rather than the management interface to do SNMP network discovery, set a service route specifying that interface and the network segment to query.

*Service routes configured on **Device** > **Setup** > **Services** > **Service Route Configuration** are not applied. SNMP network discovery only uses service routes configured here.*

5. Click the **SNMP Settings** tab and set the SNMP version and configure the required settings for the version and options you use.

SNMP Version: Choose the SNMP version that your switches support, either **V2** (SNMPv2c) or **V3**. If you choose **V2**, configure the **Community String**. If you choose **V3**,

configure the **Username**, **Security Level**, **Authentication Protocol** and **Password**, and **Privacy Protocol** and **Password** settings.

Community String (for SNMP V2): Enter the SNMP community string configured on the switches to permit read-only access.

Username (for SNMP V3): Enter a username for an SNMP user account with read-only access. This is the account the firewall uses when accessing an SNMP server running on a switch.

Security Level (for SNMP V3): Choose the security level for accessing an SNMP server on a switch.

- **noAuthNoPriv**: Choose this to not authenticate and encrypt communications between the SNMP agent on the firewall and an SNMP server on a switch.

- **authNoPriv**: Choose this to require authentication based on either MD5 or SHA hashes but not encrypt communications between the firewall and the switches.

- **authPriv**: Choose this to require both authentication and encryption.

Authentication Protocol (for SNMP V3): Choose the algorithm for authenticating communications between the firewall and the switches: **MD5** (Message Digest Algorithm 5) or **SHA** for SHA-1 (Secure Hash Algorithm 1).

Authentication Password (for SNMP V3): Enter the password used during the authentication process.

Privacy Protocol (for SNMP V3): Choose the algorithm for encrypting communications between the firewall and the switches: **DES** (Data Encryption Standard) or **AES** (Advanced Encryption Standard).

Privacy Password (for SNMP V3): Enter the password used during the encryption process.

6. Select **Enable SNMP Network Discovery Settings** and then click **OK**.

   After enabling this feature, the settings are sent to the plugin, which checks the source interface IP address that will send and receive SNMP traffic and schedules the following tasks:

   - Send SNMP queries for Network Discovery using CDP and LLDP OIDs.

   - Send SNMP queries for Network Data Refresh using various OIDs for VLANs, subnets, switch interface or port information, device IP-to-MAC address bindings, and other attributes on a per-device level.

   After the SNMP jobs are run, the resulting SNMP data is stored in files and converted to Enhanced Application logs. The firewall then sends the logs to the logging service. The logging service then streams the data to IoT Security, which updates its database and displays the SNMP discovery network topology data in the IoT Security portal.

Use ERSPAN to Send Mirrored Traffic through GRE Tunnels

Unless device traffic is visible to a firewall, the firewall cannot include it in the logs it forwards to IoT Security. When you need to collect data for devices whose traffic doesn't pass through a firewall, mirror their traffic on network switches and use Encapsulated Remote Switched Port Analyzer (ERSPAN) to send it to the firewall through a Generic Routing Encapsulation (GRE) tunnel. After the firewall decapsulates the traffic, it inspects it similar to traffic received on a TAP

port. The firewall then creates enhanced application logs (EALs) and traffic, threat, WildFire, URL, data, GTP (when GTP is enabled), SCTP (when SCTP is enabled), tunnel, auth, and decryption logs. It forwards them to the logging service where IoT Security can access and analyze the IoT device data.

You can use this feature for any deployments where traffic from remote switches needs to be inspected. IoT Security is just one use case.

The network switch mirrors (copies) IoT device traffic on one or more source ports or VLANs, uses ERSPAN to encapsulate it in a GRE tunnel, and sends it to a destination port on the firewall.

The firewall terminates the tunnel and decapsulates the mirrored traffic. It logs the traffic and then forwards the logs to the logging service, which streams the IoT device data to IoT Security for analysis.

This feature requires switches that support ERSPAN such as Catalyst 6500, 7600, Nexus, and ASR 1000 platforms.

STEP 1 | Configure a switch that supports ERSPAN to mirror traffic on one or more source ports or VLANs, and forward it through a GRE tunnel to a destination port on a next-generation firewall.

For configuration instructions, see the Cisco documentation for your switch.

STEP 2 | Enable ERSPAN support on the firewall.

By default, ERSPAN support is disabled.

1. Log in to the firewall and select **Device** > **Session**.
2. Click the **Edit** icon for Session Settings, select **Enable ERSPAN Support**, and then click **OK**.

The ERSPAN Support check box in the Session Settings section is now selected.

STEP 3 | **Commit** your change.

STEP 4 | Create a Layer 3 security zone specifically to terminate the GRE tunnel and receive mirrored IoT device traffic from the source port on the network switch.

1. Select **Network** > **Zones** and then **Add** a zone.

2. Enter the following and leave the other settings at their default values:

Name: Enter a meaningful name for the zone such as **ERSPAN-IoT-data**.

Log Setting: Select **IoT Security Default Profile** or another log forwarding profile that sends the required types of logs to the logging service for IoT Security.

> *You must already have logging services enabled on the firewall.*

Type: **Layer3**

3. Click **OK**.

STEP 5 | Create a Layer 3 interface and bind it to the zone you just created.

1. Select **Network** > **Interfaces** > **Ethernet** and then click the Ethernet interface on which you want to terminate the GRE tunnel from the switch. Optionally, use a subinterface.

2. Enter the following and leave the other settings at their default values:

   Comment: Enter a meaningful note about the interface for later reference.

   Interface Type: Layer3

   Virtual Router: Choose the virtual router you want to route to the interface. Consider using a separate virtual router exclusively for ERSPAN traffic.

   Security Zone: Choose the zone you just created.

3. Click **IPv4**, select **Static** for the address type, and **Add** an IP address for the interface.

   The switch uses this in its GRE tunnel configuration as the IP address of its peer.

4. Click **Advanced** and either add a **New Management Profile** or select a previously defined profile that allows the Ethernet interface to accept different types of administrative traffic.

5. Click **OK** to save the new interface management profile and then click **OK** again to save the Ethernet interface configuration.

STEP 6 | Create a tunnel interface with an IP address in the same subnet as that of the corresponding tunnel interface on the switch and bind it to the zone you just created.

1. Select **Network** > **Interfaces** > **Tunnel** and then **Add** the logical tunnel interface for the GRE tunnel from the switch.

2. Enter the following and leave the other settings at their default values:

   Interface Name: The field on the left is read-only and contains the text "tunnel". Enter a number in the field on the right to complete the name. For example, enter **8** to make the name `tunnel.8`.

   Virtual Router: Choose the same router you used for the Layer 3 interface.

   Security Zone: Choose the same zone to which you bound the Layer 3 interface.

3. Click **IPv4** and **Add** an IP address that's in the same subnet as the IP address of the logical tunnel interface on the switch.

4. Click **Advanced** and either add a **New Management Profile**, or select a previously defined profile, to allow the tunnel interface to accept different types of administrative traffic.

5. Click **OK**.

STEP 7 | Configure static routes for the virtual router (VR) for ERSPAN.

1. Select **Network** > **Virtual Routers** and then click the virtual router for ERSPAN.

2. Click **Static Routes** and then click **+ Add**.

3. Enter the following and leave the other settings at their default values:

Name: Enter a name for the static route.

Destination: **0.0.0.0/0**

> *If you know the subnets beyond the switch, create individual static routes for each of them. Otherwise, use a separate VR for ERSPAN and set a default route.*

Interface: **ethernet1/3** (the interface you previously configured)

Next Hop: **None**

4. Click **OK**.

STEP 8 | Configure a GRE tunnel with ERSPAN enabled.

1. Select **Network** > **GRE Tunnels** and then click **+ Add**.
2. Enter the following and leave the other settings at their default values:

   Name: Enter a name for the GRE tunnel; for example, **GRE-ESPAN-for-IoT-data**

   Interface: Choose the Layer 3 interface you configured for GRE tunnel termination.

   Local Address: Choose **IP** and the IP address of the Layer 3 interface where the GRE tunnel terminates.

   Peer Address: Enter the IP address of the switch egress interface from which it initiates the GRE tunnel.

   Tunnel Interface: Choose the logical tunnel interface you configured for the GRE tunnel.

   ERSPAN: (select)

3. Click **OK**.

   The IP addresses of the Ethernet and tunnel interfaces in relation to each other and the rest of the network look like this.

STEP 9 | **Commit** your changes.

Use DHCP Server Logs to Increase Device Visibility

IoT Security relies on IP address-to-MAC address bindings to ascribe observed network behaviors to IoT devices and uniquely track them. IoT Security typically uses DHCP traffic to learn IP address-to-MAC address bindings and track IP address changes. However, in designs where the next-generation firewall is not in the DHCP data path, you can use this method to ingest DHCP server logs and expand DHCP traffic visibility.

In areas of the network where it's difficult to route DHCP traffic to or through a firewall, configure DHCP servers to send their server logs as syslog messages to the firewall. The firewall then forwards the messages as Enhanced Application Logs (EALs) with a subtype of dhcp-syslog through the logging service to IoT Security. IoT Security parses them to learn the IP address-to-MAC address bindings and add newly learned devices to its inventory. IoT Security also learns device hostnames from the server logs, with the exception of logs from Cisco DHCP servers.

DHCP server	Next-generation firewall	Logging service	IoT Security
A DHCP server sends syslog messages with IP address-to-MAC address bindings to a next-generation firewall.	The firewall forwards the syslog messages in Enhanced Application Logs (EALs) to the logging service.	The logging service streams the EALs to IoT Security to parse and ingest.	IoT Security parses the EALs, analyzes the metadata, and adds newly learned devices to its inventory.

Prerequisites

- A DHCP server with syslog capabilities configured to send messages to a syslog server running on a next-generation firewall

- A next-generation firewall running PAN-OS 11.0 or later with an active IoT Security subscription

> *DHCP server log ingestion is not available on CN-, M-, and WF-series next-generation firewalls.*

Set up the Next-generation Firewall

Set up your next-generation firewall to receive syslog messages from one or more DHCP servers. The firewall will automatically forward the syslog messages it receives as EALs to the logging service, which streams them to IoT Security to parse and analyze.

STEP 1 | Add a DHCP server to the next-generation firewall.

1. Log in to your next-generation firewall, select **Device** > **IoT** > **+ Add**, configure the following, and then click **OK**:

Name: Enter a name for the DHCP server. It can be up to 32 characters, including spaces.

Description: Enter a note about the DHCP server for future reference. It can be up to 256 characters, including spaces.

Enabled: Select to enable the firewall to listen for connections from the DHCP server and process them when they come.

IP Address: Enter the IP address from which the DHCP server will connect to the firewall. The address can be in IPv4 or IPv6 format. An FQDN is not allowed.

Protocol: Select **TCP**, **UDP**, or **SSL**. When making your choice, consider what's important for the connection between the DHCP server and firewall. TCP provides transmission reliability but not security. UDP provides low processing overhead and faster speeds but lacks reliability and security. SSL provides reliability and security but incurs more overhead.

> *The firewall listens for DHCP server connections using TCP and UDP on port 10514 and connections using SSL on port 16514.*

STEP 2 | Repeat the previous step to add more DHCP servers.

Add more DHCP servers and expand visibility of DHCP traffic throughout your network as needed. All next-generation firewalls support a maximum of 100 DHCP servers per firewall.

Set up DHCP Servers for Syslog

Configure your DHCP servers to send syslog messages of their server logs to the management interface on the next-generation firewall. Make sure to configure the DHCP server to use the same protocol configured for it on the firewall: TCP, UDP, or SSL. You can use DHCP servers such as Windows, Linux, Cisco, or Infoblox for example. See the documentation for your DHCP servers for configuration instructions.

Check DHCP Server Connection Status

To see all the configured DHCP servers, select **Device** > **IoT**

A green circle next to a DHCP server name means it was configured in Panorama and is read-only when viewed in the web interface of the local next-generation firewall.

When a DHCP server using TCP or SSL is currently connected to the firewall, "Connected" appears in the Status column. "Connected" also appears in this column if a DHCP server using UDP has been connected within the past two hours. At all other times, the Status column is empty, indicating that the server isn't currently connected to the firewall.

The following CLI commands are also useful for checking DHCP server settings, the status of their connections, and the data they're providing for IoT Security.

`show iot dhcp-server status { all	server <server-name> }`	Entering `all` shows a table with all DHCP servers configured and enabled on the firewall, the port numbers on which they connect, and their current connection status. Entering `server <server-name>` shows detailed information about a specific DHCP server and its recent activity.
`show iot eal dhcp-syslog-eal`	This command shows information related to EALs carrying DHCP server syslog messages.

Plan for Scaling when Your Firewall Serves DHCP

This section discusses scaling the solution for when the firewall provides DHCP services as described in Configure a Pre-PAN-OS 10.0 Firewall with a DHCP Server.

Align Numbers of VLAN Subinterfaces with Physical Interfaces

For consistency, align the VLAN subinterface numbers with the physical interface numbers they serve. For example, interface vlan.1 serves DHCP for the network attached to ethernet1/1. This allows you to associate them with each other faster and troubleshoot issues more easily later.

Interface	Interface Type	Management Profile	Link State	IP Address
ethernet1/1	Layer3	Inside_Manage...		10.1.0.1/24
ethernet1/2	Layer3	Inside_Manage...		10.2.0.1/24
ethernet1/3	Layer3	Inside_Manage...		10.3.0.1/24

Interface	Management Profile	IP Address	Virtual Router
vlan		none	none
vlan.1	Inside_Manage...	1.1.1.1/32	DHCP_VR
vlan.2	Inside_Manage...	1.1.1.2/32	DHCP_VR
vlan.3	Inside_Manage...	1.1.1.3/32	DHCP_VR

Conserve IP Addresses for VLAN Subinterfaces

When production IP address space is used for the VLAN interfaces, giving them IP addresses with 32-bit netmasks will conserve address space. You can use addresses from a single network (for example, 1.1.1.0/24) for all the VLAN interfaces. Because these interfaces exist solely to serve DHCP to a local network, the addresses assigned to the VLAN interfaces don't need to be routable in the rest of the enterprise. Operationally, this means that the same network space and addresses can be used for VLAN interfaces on all firewalls in the enterprise.

Configure a Network Route to all VLAN Interfaces

When configuring this solution for multiple interfaces, the routing configuration changes slightly. On the default (production) virtual router, you can configure a network route to the VLAN interfaces instead of a collection of host routes. In the figure below all of the VLAN interfaces have addresses that can be summarized using a 1.1.1.0/24 route.

On the DHCP virtual router, add network routes for each network for which a VLAN interface serves DHCP and set the default (production) virtual router as the next hop. Adding network rather than host routes for the DHCP relay agents allows the probe feature on the DHCP servers to function.

Prepare Your Firewall for IoT Security

The following steps describe how to enable logging service on a next-generation firewall and configure it to obtain and log network traffic metadata. It then explains how to forward the collected metadata to the cloud-based logging service where IoT Security uses it to identify various IoT devices on the network.

The steps below assume you already completed the IoT Security onboarding process but still need to do the following:

- Install a device license and a logging service license on your firewalls.
- Install certificates on your firewalls (if they aren't installed already).
- Configure your firewalls to collect network traffic metadata.
- Configure your firewalls to forward the collected metadata in logs to the logging service.
- Enable Device-ID on zones with devices that you want to monitor and protect with Security policy rules.
- (Optional) Create service routes and Security policy rules to permit firewalls to communicate with the logging service, IoT Security, and update server through a data interface.

> *For additional details about configuring a firewall for IoT Security, see* Device-ID.

STEP 1 | Install licenses required for IoT Security to function.

After onboarding IoT Security, take one of the following actions to install the licenses your firewalls need to use IoT Security:

Next-generation firewalls: Log in to each of your firewalls, select **Device** > **Licenses**, and then select **Retrieve license keys from license server** in the License Management section.

or

Panorama: Log in to Panorama, select **Panorama** > **Device Deployment** > **Licenses**, and then **Refresh**. Select the devices onboarded with IoT Security and **Refresh**.

This installs the licenses for IoT Security and the logging service on the firewall.

> *When the time comes to renew IoT Security licenses, use this retrieval function on your firewalls so that they extend their license expiration dates.*

STEP 2 | If necessary, generate a one-time password (OTP) and pre-shared key (PSK) to get device and logging service certificates.

> 📋 *This step only applies to firewalls with an IoT Security, Doesn't Require Data Lake Subscription. If your firewalls have an IoT Security Subscription, which requires a Strata Logging Service, see the* Strata Logging Service Getting Started *for details about generating certificates and installing them on your firewalls.*

- **Firewalls with PAN-OS 10.1 or later**

  > 📋 *Skip this step if your firewalls run PAN-OS 10.1 or later and already have a* device certificate *installed. Any firewalls on which you've previously installed a device certificate for another Palo Alto Networks product already have this certificate and don't require a new one. You can check if your firewall has a valid certificate in the General Information section on the Dashboard page in the PAN-OS web user interface.*

  Firewalls running PAN-OS 10.1 or later require a device certificate but not a logging service certificate.

  The following next-generation firewall models automatically install a device certificate when they first connect to the Customer Support Portal (CSP); therefore, you don't have to install one manually on any of these firewalls running these PAN-OS versions:

  - **PAN-OS 10.1**: PA-410, PA-440, PA-450, PA-460, and PA-5450 firewalls

  - **PAN-OS 10.2**: PA-410, PA-440, PA-450, and PA-460 firewalls; PA-1400 Series and PA-3400 Series firewalls; and PA-5410, PA-5420, PA-5430, and PA-5450 firewalls

  - **PAN-OS 11.0**: PA-400 Series, PA-1400 Series, PA-3400 Series, PA-5400 Series, and PA-5450 firewalls

  Also any firewalls on which you've previously installed a device certificate for another Palo Alto Networks product already have a device certificate and don't require a new one.

  Check the following questions and answers to determine when to generate and install a device certificate on a firewall.

Do firewalls already have a device certificate?	Do firewalls already have a logging service certificate?	Are firewalls managed by Panorama?	What to do?
Yes	N/A	N/A	Skip this step.
No	N/A	Yes	Enter the Panorama serial number, generate an OTP in the Customer Support Portal, and enter it in Panorama to generate a device certificate.

Do firewalls already have a device certificate?	Do firewalls already have a logging service certificate?	Are firewalls managed by Panorama?	What to do?
No	N/A	No	Generate an OTP in the Customer Support Portal and install a device certificate on the firewall.

- **Firewalls with PAN-OS 10.0**

  📋 *Skip this step if your firewalls run PAN-OS 10.0 and already have* device and logging service certificates *installed. Any firewalls on which you've previously installed a device certificate and logging service certificate for another Palo Alto Networks product already have these certificates and don't require new ones. You can check if your firewall has valid certificates in the General Information section on the Dashboard page in the PAN-OS web user interface.*

  Check the following questions and answers to determine when to generate and install a device and logging service certificate on a firewall.

Do firewalls already have a device certificate?	Do firewalls already have a logging service certificate?	Are firewalls managed by Panorama?	What to do?
Yes	Yes	N/A	Skip this step.
Yes	No	Yes	Enter the Panorama serial number, copy the OTP, and enter it when installing the Cloud Services plugin on Panorama.
Yes	No	No	Copy the preshared key and paste it in a PAN-OS firewall to generate a logging service certificate.
No	Yes	Yes	Enter the Panorama serial number and use Panorama to generate and install a device certificate on one or more firewalls.

Do firewalls already have a device certificate?	Do firewalls already have a logging service certificate?	Are firewalls managed by Panorama?	What to do?
No	Yes	No	Generate an OTP in the Customer Support Portal and install a device certificate on the firewall.
No	No	Yes	Copy the OTP, and enter it when installing the Cloud Services plugin on Panorama. When Panorama pushes a configuration requiring logging services and IoT Security to a firewall that doesn't have a logging service and device certificate, the firewall responds to Panorama by requesting the certificates.
No	No	No	Generate an OTP in the Customer Support Portal and install a device certificate on the firewall. Copy the preshared key and paste it in a PAN-OS firewall to

Do firewalls already have a device certificate?	Do firewalls already have a logging service certificate?	Are firewalls managed by Panorama?	What to do?
			generate a logging service certificate.

- **Panorama-managed Firewalls Running PAN-OS 8.1 – 9.1**

  Skip this step if your firewalls are managed by Panorama, run PAN-OS 8.1-9.1, and already have a logging service certificate *installed. Any firewalls on which you've previously installed a logging service certificate for another Palo Alto Networks product don't require a new one. You can check if your firewall has a valid certificate in the General Information section on the Dashboard page in the PAN-OS web user interface.*

  Check the following questions and answers to determine when to generate and install a logging service certificate on a firewall.

Do firewalls already have a device certificate?	Do firewalls already have a logging service certificate?	Are firewalls managed by Panorama?	What to do?
N/A	Yes	Yes	Skip this step.
N/A	Yes	No	Skip this step if the firewalls are running PAN-OS 9.0.3-9.1 with or without Panorama management. Panorama is required for firewalls running PAN-OS 8.1–9.0.2 to get a logging service certificate. If you aren't using Panorama to manage firewalls with these PAN-OS versions, then your firewalls cannot send logs to the logging service to support IoT Security.
N/A	No	Yes	Copy the OTP, and enter it when installing the Cloud services plugin on

Do firewalls already have a device certificate?	Do firewalls already have a logging service certificate?	Are firewalls managed by Panorama?	What to do?
			Panorama. When Panorama pushes a configuration requiring logging services to a firewall that doesn't have a logging service certificate, the firewall responds to Panorama by requesting it.
N/A	No	No	Firewalls running PAN-OS 8.1–9.0.2 require Panorama management to get a logging service certificate; they cannot support IoT Security without Panorama. For firewalls running PAN-OS 9.0.3-9.1 without Panorama management, copy the preshared key and paste it in a PAN-OS firewall to generate a logging service certificate.

📋 *For information about the sites that next-generation firewalls contact to authenticate certificates when communicating with IoT Security, see* IoT Security Integration with Next-generation Firewalls*.*

1. Log in to the IoT Security portal as a user with owner privileges. To be able to generate OTPs and PSKs, your user account must have been created in the Customer Support Portal (CSP) and assigned a superuser role in the relevant tenant service group (TSG) in Identity & Access. A superuser role in the hub provides owner privileges in IoT Security.

2. Select **Administration** > **Firewalls** > **Certificate Generation**.

3. If you manage your firewalls with Panorama, choose **Yes** and enter its serial number. This will link your Panorama management server with the applications in this TSG. You can find the Panorama serial number in your Customer Service Portal account in **Assets** > **Devices**. After you choose **Yes** and enter your Panorama serial number, IoT Security displays the

materials you need to get the certificate or certificates that firewalls need to secure their connections with IoT Security and the logging service.

To get a device certificate, click the link to the Customer Support Portal, log in to your account, and then follow the instructions below. To generate a logging service certificate, copy the OTP or PSK and follow the instructions below.

If you don't use Panorama, choose **No**. Because an OTP for a logging service certificate applies only to Panorama, it isn't shown.

Consider the following points when deciding which certificates you need and how to generate them:

Device Certificate: From PAN-OS 10.0, firewalls require a device certificate to authenticate with IoT Security and, from PAN-OS 10.1, to also authenticate with the logging service. To generate and install a device certificate on firewalls directly and through Panorama:

- Generate and install a device certificate on each firewall.
- Use Panorama to generate and install a device certificate on one or more firewalls.

> *When a device certificate is installed on a firewall so it can authenticate itself to the logging service and IoT Security, the firewall cannot decrypt encrypted traffic to inspect it and enforce policy rules on it. Therefore, don't try to use decryption policy rules on firewalls that have a device certificate installed on them.*

Logging Service Certificate – One-Time Password: An OTP is necessary for Panorama to verify itself with its logging service instance and obtain logging service certificates

for Panorama-managed firewalls running PAN-OS 8.1-10.0. A logging service certificate authenticates firewalls with the logging service.

1. Regenerate the OTP if necessary and copy it.

2. Log in to the Panorama web interface as an admin user. and select **Panorama** > **Setup** > **Management** > **Device Certificate** and **Get certificate**.

3. Paste the OTP and then click **OK**.

Logging Service Certificate – Pre-Shared Key: A PSK is necessary to generate a logging service certificate on firewalls without Panorama management running PAN-OS

9.0.3-10.0.x. A logging service certificate authenticates firewalls with the logging service. To generate a logging service certificate:

1. Regenerate the PSK if necessary and copy it.

2. Log in to your PAN-OS 9.0.3-10.0.x firewall and select **Device** > **Setup** > **Management**.

3. In the Strata Logging Service section, click **Connect** next to Onboard without Panorama.

   This opens the Onboard without Panorama dialog box.

4. Paste the PSK and **Connect**.

   The firewall first connects to the Customer Support Portal, submits the PSK, and downloads a logging service certificate. It then uses the certificate to authenticate itself and connect securely to the logging service.

5. Click the **Edit** icon (gear) for Strata Logging Service. Select **Enable Strata Logging Service** and **Enable Enhanced Application Logging**.

   or

   If you have an IoT Security–Doesn't Require Data Lake license, select **Enable Duplicate Logging (Cloud and On-Premises)** and **Enable Enhanced Application Logging**.

6. Choose the region where the logging service will ingest logs from your firewalls.

   For PA-7000 and PA-5200 models, enter the number of connections for sending logs from the firewall to the logging service. The range is 1-20 and the default is 5.

7. When done, click **OK**.

> *The term "Strata Logging Service" is a bit of a misnomer. The firewall forwards logs to the logging service, which only streams them to Strata Logging Service if you're using it for data retention. An IoT Security, Doesn't Require Data Lake subscription doesn't use Strata Logging Service at all, but it still requires that this setting be enabled.*

STEP 3 | Make sure your firewall is set up to apply policy to DHCP traffic between DHCP clients and their DHCP server and to log their traffic.

For detailed instructions about setting up firewalls to capture and log DHCP traffic, see Firewall Deployment for Device Visibility.

If the firewall is running a PAN-OS 10.0 release or later with a DHCP server on one of its interfaces, enable **DHCP Broadcast Session** on **Device** > **Setup** > **Session**. This setting is supported on all firewalls running PAN-OS 10.1.10 or later, PAN-OS 10.2.4 or later, and PAN-OS 11.0.1 or later. (For more information, see Firewall Deployment Options for IoT Security.)

> *In addition to detecting devices with dynamically assigned IP addresses, IoT Security also discovers and identifies devices with static IP addresses. To learn about the multiple methods IoT Security uses to do this and how you can assist, see Devices with Static IP Addresses.*

STEP 4 | To forward logs to the logging service, click **Objects** > **Log Forwarding** and then click **Add**.

Configure a log forwarding profile on the firewall to send enhanced application logs to the logging service so the IoT Security app can ingest network traffic data. Optionally, instead of adding a new profile, you can edit an existing one.

STEP 5 | In the Log Forwarding Profile, enter a name such as Log-Forwarding, click **Enable enhanced application logging to Strata Logging Service (including traffic and url logs)**, and then click **OK**.

> *Enhanced application logging was introduced in PAN-OS 8.1.*

A list of enhanced application Logs automatically populates the page and forwards all logs per type to the logging service. Selecting **Enable enhanced application logging to Strata Logging Service (including traffic and url logs)** enables the firewall to capture packet payload data (EALs) in addition to session metadata (regular logs) for these different log types. When this log forwarding profile is attached to a Security policy rule to control traffic, the firewall forwards both types of data to the logging service. You cannot delete any of these logs from the profile nor modify any of the filters in the Filter column, which are the default "All Logs" filter.

The following describes each log type, explains if IoT Security uses it, and what its purpose is:

- **traffic** – Traffic logs contain entries for the end of each network session and, optionally, the start of a network session. IoT Security uses traffic logs to identify devices, generate policy rule recommendations, risk assessment, device behavior anomaly detection, correlate sessions, and raise security alerts.

- **threat** – Threat logs contain entries for when network traffic matches one of the security profiles attached to a next-generation firewall Security policy rule. IoT Security uses threat logs to assess risks, detect vulnerabilities, raise security alerts, and generate policy rule recommendations.

- **wildfire** – WildFire® logs contain entries for when WildFire security profiles are attached to a Security policy rule and files are traversing the network. IoT security uses WildFire

logs to detect IoT-specific file-based attacks, raise security alerts, and generate policy rule recommendations.

- **url** – URL logs are written whenever network traffic matches a URL filtering profile attached to a Security policy rule. IoT Security does not currently use URL filtering logs.

- **data** – Data logs can represent either a successful file data transfer or an attempted file transfer that was blocked by the firewall. IoT Security does not currently use data logs.

- **gtp** (When GTP is enabled) – GTP logs are written whenever a firewall is processing traffic from 3G, 4G, and 5G cellular devices. IoT Security uses the metadata from this traffic to identify cellular devices and their network behaviors. If such traffic isn't on the network, firewalls don't generate GTP logs, and you can safely ignore the red icon that appears in the Status column for it on **Administration** > **Firewalls** in the IoT Security portal.

- **sctp** (When SCTP is enabled) – SCTP logs are written whenever a firewall is processing Stream Control Transmission Protocol traffic. IoT Security does not currently use SCTP logs.

- **tunnel** – Tunnel logs are written whenever a firewall is processing Generic Routing Encapsulation (GRE) or null encryption IPsec traffic. They contain metadata about the traffic inside these types of tunnels. IoT Security does not currently use tunnel logs.

- **auth** – Auth logs contain information about authentication events seen by the firewall. These occur when users access network resources which are controlled by authentication policy rules. IoT Security does not currently use auth logs.

- **decryption** – Although IoT Security uses decrypted SSL data to improve device identification, risk assessment, and threat detections, it doesn't use decryption logs, which are helpful when troubleshooting issues with decryption.

If you name the log forwarding profile "default" (all lowercase), the firewall will automatically apply it to new Security policy rules when they're created—or when they're imported from IoT Security. Doing this will save you time and effort when importing Security policy rule recommendations from IoT Security. Because imported rule recommendations don't include a log forwarding profile, you have to add one manually to each rule after you import it. However, by naming the profile "default", you can avoid this step. (Note that the "default" log forwarding profile will be applied when adding new Security policy rules, but it won't be retroactively applied to existing rules.)

STEP 6 | Enable log forwarding on Security policy rules.

On Security policy rules that apply to traffic whose data you want to collect, enable log forwarding and choose the log forwarding profile you just created to send enhanced application logs for this traffic to the logging service. For information, see Configure Policies for Log Forwarding.

STEP 7 | Enable Device-ID in each zone where you want to use it to detect devices and enforce your Security policy rules.

For detailed configuration instructions, see Configure Device-ID in the PAN-OS Administrator's Guide.

STEP 8 | (Optional) Create service routes.

By default, firewall uses its Management interface to send data logs to the logging service, get recommended policy rule sets and IP address-to-device mappings from IoT Security, and download device dictionary files from the update server. When a firewall uses its Management interface for all this, a service route and a Security policy rule are not needed.

However, when a firewall accesses the logging service, IoT Security, and update server through a data interface, then you must add a service route identifying the source data interface, source interface IP address, and service type. In addition, you must add an interzone Security

policy rule permitting Data Services from 127.168.0.0/16 to the destination zone where the logging service, IoT Security, and update server are.

> *When a firewall generates traffic that it sends through a data interface, it uses an IP address in the 127.168.0.0/16 subnet as its internal source and then translates it to the IP address of the source interface. Because Security policy rules are applied to the original source IP address before NAT, the source IP address must be 127.168.0.0/16 instead of the IP address of the source interface.*

1. If necessary, configure the data interface you want to use as the source interface for required IoT Security communications.
2. Select **Device** > **Setup** > **Services** > **Service Route Configuration** and then select **Customize**.
3. On the IPv4 tab, select **Data Services** and then choose the data interface you want to use as the Source Interface.

   Its IP address autofills the Source Address field. This service route is for forwarding enhanced application logs (EALs) to the logging service.

   > *Device-ID and IoT Security do not support IPv6.*

4. Click **OK**.
5. Click **IoT**, choose the same data interface as the Source Interface, and then click **OK**.

   This service route is for pulling IP address-to-device mappings and policy recommendations from IoT Security.
6. Click **Palo Alto Networks Services**, choose the same data interface, and then click **OK**.

   This service route is for forwarding other logs besides EALs to the logging service and for pulling device dictionary files from the update server.
7. Click **OK** to save your configuration changes.

87

STEP 9 | (Optional) If you created service routes in the previous step, add Security policy rules permitting services required for the firewall to use IoT Security.

1. Select **Policies** > **Security** > **+ Add**.
2. On the General tab, enter a name for the Security policy rule and choose **interzone** as the Rule Type.
3. On the Source tab, select **Any** as the source zone and then **Add 127.168.0.0/16** as the source address.
4. On the Destination tab, **Add** the destination zone with IoT Security, and **Add** the edge services FQDN for your region as the destination address.
5. On the Application tab, **Add paloalto-iot-security**.

   The firewall uses this application to pull IP address-to-device mappings and policy recommendations from IoT Security.

6. On the Actions tab, choose **Allow** and then click **OK**.
7. If you have an intranet policy rule that allows all intranet traffic in the zone where the logging service and update server are, you can use that rule to allow the firewall to forward logs to the logging service and pull dictionary files from the update server.

   Otherwise, create an intranet policy rule that allows the firewall to send these three applications to the logging service and update server from the IP address of the firewall interface in the same zone:

   paloalto-shared-services to forward EALs and session logs to the logging service

   paloalto-logging-service to forward other logs besides EALs to the logging service

   paloalto-updates to pull device dictionary files from the update server

STEP 10 | Commit your configuration changes.

After the configuration is committed, the firewall begins generating logs and forwarding them to the logging service. You can use the Explore app in the hub to see the progress of log forwarding between the firewall and the logging service.

Configure Policies for Log Forwarding

Enable log forwarding so that the firewall sends Enhanced Application logs (EALs) to the Palo Alto Networks cloud-based logging service. IoT Security then fetches metadata from there for analysis.

Configure an Interzone Policy

If the VLAN interfaces are set in different L3 security zones from the Ethernet interfaces with which they're paired, Security policy rules must be configured for the solution to work. The figure below shows example rules when multiple VLAN interfaces have been configured to support multiple Ethernet interfaces.

	Name	Type	Source Zone	Source Address	Source User	Source HIP Profile	Destination Zone	Destination Address	Application	Service	Action
1	Allow_DHCP_Relay	universal	Office Prod_Line Wireless	10.1.0.1 10.2.0.1 10.3.0.1	any	any	DHCP Zone	1.1.1.0/24	dhcp	application-d...	Allow
2	Allow_DHCP_Probe	universal	DHCP Zone	1.1.1.0/24	any	any	Office	10.1.0.0/24 10.2.0.0/24 10.3.0.0/24	ping	application-d...	Allow
3	Allow_Interface_Testing	universal	Office Prod_Line Wireless	10.1.0.1 10.2.0.1 10.3.0.1	any	any	DHCP Zone	1.1.1.0/24	ping	application-d...	Allow
4	intrazone-default	intrazone	any	any	any	any	(intrazone)	any	any	any	Allow
5	interzone-default	interzone	any	any	any	any	any	any	any	any	Deny

Policy rule 1: This policy rule allows relayed unicast DHCP messages from the zones assigned to interfaces ethernet1/1 - ethernet1/3 to the DHCP zone. In addition, enable log forwarding and choose the log-forwarding profile you previously created to send EALs for this traffic to the logging service.

💡 *If you name the log forwarding profile "default" (all lowercase), the firewall will automatically apply it to new Security policy rules when they're created—or when they're* imported from IoT Security. *Doing this will save you time and effort when importing Security policy rule recommendations from IoT Security. Because imported rule recommendations don't include a log forwarding profile, you have to add one manually to each rule after you import it. However, by naming the profile "default", you can avoid this step. (Note that the "default" log forwarding profile will be applied when adding new Security policy rules, but it won't be retroactively applied to existing rules.)*

Policy rule 2: This rule allows ping (ICMP echo requests) from the VLAN interfaces in the DHCP zone to networks configured on ethernet1/1 - ethernet1/3.

Policy rule 3: This rule allows ping from the IP addresses assigned to ethernet1/1 - ethernet1/3 to VLAN interfaces configured in the DHCP zone.

Configure an Intrazone Policy

You must override the logging and log forwarding settings in the default intrazone policy rule so that the firewall will forward logs to the logging service.

If the interface hosting the DHCP server is in the same zone as the interface your clients are on, the default intrazone policy rule applies to this traffic, which, by default, allows all traffic within this zone but does not have logging and log forwarding enabled. Therefore, you must override this by enabling log forwarding on your default intrazone policy rule.

> *Even for cases where the DHCP server is in a different zone from the DHCP clients and an interzone policy is applied to their DHCP traffic, we still recommend that you enable log forwarding on the default intrazone policy rule to capture the enhanced application logs for traffic within that zone.*

STEP 1 | Click **Policies** > **Security**, select **intrazone-default**, and then click **Override**.

The Security Policy Rule configuration window appears.

STEP 2 | Click **Actions**, select **Log at Session End**, choose the log forwarding profile you just configured from the Log Forwarding drop-down list, and then click **OK**.

Control Allowed Traffic for Onboarding Devices

When new devices join the network, they must be allowed to function normally so that IoT Security can identify them by analyzing their normal network behavior. However, firewalls are typically configured with Zero Trust security policy rules that allow only the network activities that devices need based on their function. As a result, the rules might inadvertently block traffic for a new device that, if allowed, would have allowed IoT Security to determine its identity.

To overcome this, you can configure one or more onboarding policy rules that use Device-ID to apply the rules only to devices that have been recently detected on the network but have not yet been confidently identified. For the firewall to enforce the rule, a device must be categorized as an Onboarding Device. IoT Security places low-confidence devices in this category during a customizable period of time that starts when IoT Security first detects them on the network. Devices continue to be categorized as "Onboarding Device" until IoT Security confidently identifies them with a confidence score above 90 or until the time period ends. The policy rule does not apply to other, previously identified devices and must be configured to allow new devices enough network access for IoT Security to identify them. Once IoT Security identifies them, it switches them over to an appropriate category for what they are. The firewall can then apply appropriate policy rules based on their identities. If IoT Security cannot confidently identify one or more devices and the time period expires, it still switches them to a category it considers appropriate, but because their confidence scores are below 90, IoT Security doesn't generate any security rule recommendations.

STEP 1 | Configure a security policy rule that allows certain types of traffic from any device whose Device-ID attribute for Category is "Onboarding Device".

1. Log in to the PAN-OS or Panorama web portal and configure a security policy rule that allows the basic types of traffic that devices in certain VLANs or in different IP address subnets would be expected to generate. For example, a rule for a VLAN that contains printers should allow only typical printer-specific traffic, whereas a rule for a VLAN that contains medical scanning equipment should only allow typical types of traffic for scanners.

2. Add a Device-ID component to the rule and specify **Onboarding Device** as the category that a device must match for the firewall to apply the rule. (In short, **Add** a security policy rule on **Policies** > **Security**. Select the **Source** tab, click **Add** in the Source Device section, and then click **Device**. In the Device Object dialog box that appears, choose **Onboarding Device** in the Category list.)

3. Create additional security policy rules that specify **Onboarding Device** as the category in the Device-ID section of the rule configuration.

STEP 2 | Enable the new device onboarding feature based on Device-ID in IoT Security.

1. Log in to the IoT Security portal as a user with owner privileges.

2. Select **Policy Sets** > **Settings** and toggle on **Control newly onboarded low-confidence devices through firewall policy rules**.

3. Optionally change the period of time during which IoT Security categorizes a device as an Onboarding Device if it has an identity confidence score below 90. The default onboarding

period is 7 days. There are no maximum and minimum limits. You can also switch from a limited period of time to an unlimited length of time.

After you enable this feature and set a length of time for the onboarding period, IoT Security displays a daily system alert if there are any devices for which the onboarding period will soon be expiring. The alert appears a few days before the expiration and includes a link to the **Assets** > **Devices** page with a filter applied to show just these devices.

4. To see which devices are in the Onboarding Device category, select **Assets** > **Devices** and, if necessary, show the **Onboarding Device** column in the Devices table.

> *If necessary, also show the* ***First Seen*** *column and then sort by this to organize the display of devices based on the order in which IoT Security first discovered them on the network.*

Support Isolated Network Segments

An isolated network segment is a part of a private network that allows an extremely limited set of connections between devices in the segment and devices in any other local segment or in the public network. Because IoT Security is a cloud-based application that relies on network traffic logs to provide its services, there needs to be a way to get the logs to IoT Security without compromising the security of the isolated segment. To accomplish this, you can configure next-generation firewalls as security telemetry gateways (referred to in the PAN-OS web interface as *proxies*) to forward traffic logs from the isolated segment through the non-isolated part of the network to the Palo Alto Networks logging service, where IoT Security can access it. In addition, the security telemetry gateways can forward requests from isolated firewalls for the data and files they need to onboard IoT Security and support Device-ID: licenses, certificates, IP address-to-device mappings, security policy rule recommendations, and dictionary file downloads.

This data path occurs only through security telemetry gateways, and only requests and network traffic logs that next-generation firewalls generate, not actual data from protected devices, are sent on this path through the security telemetry gateway chain.

Importantly, there are no direct connections between devices in the isolated network segment and the cloud, and the status of the security telemetry gateway-to-cloud connection (up or down) has no impact on protected device operations nor on next-generation firewall functions such as policy enforcement and threat detection and prevention. All protected device and firewall operations will continue to operate even if an upstream security telemetry connection goes down.

You can use a single security telemetry gateway or a chain of two or more security telemetry gateways for additional security layering. In this way, Palo Alto Networks can provide IoT Security services to industries that have isolated OT networks as is common in power utilities and oil and gas companies for example. These networks typically consist of two segments: an IT network and OT network. Leveraging existing next-generation firewalls or deploying new ones, you could configure two firewalls as security telemetry gateways, placing one at the boundary between OT and IT networks and the other at the boundary between the IT and public networks. Firewalls in the OT network would send traffic logs to the OT security telemetry gateway, which forwards them to the IT security telemetry gateway, which forwards them to the Palo Alto Networks logging service. Setting up next-generation firewalls in a security telemetry gateway chain like this increases the depth of the logical network segment boundary because the IT security telemetry gateway blocks inbound connections to the OT security telemetry gateway.

The following next-generation firewalls support the security telemetry gateway feature:

- Physical firewalls: PA-1400 series, PA-3400 series, PA-5400 series (except PA-5450)
- VM-300, VM-500, VM-700

The firewalls must be running PAN-OS version 11.0.1-h2 or later.

When deploying firewalls for a network that contains an isolated OT network segment, set up the security telemetry gateways in order from the IT perimeter (the IT security telemetry gateway) toward the deepest part of the OT network: IT security telemetry gateway, then OT security telemetry gateway, and then OT firewalls. By deploying them in this order, you will have the information you need after completing one deployment to deploy the next one. Also, as each firewall comes online, the firewall or firewalls that the next one needs to reach the public network will already be online and reachable.

93

The following illustration shows the logical relationship of next-generation firewalls in a security telemetry gateway chain and the IP addresses and subnets used as examples in the configuration instructions that follow. As shown here, OT firewalls initiate all outbound connections through the OT and IT security telemetry gateways to the logging service, IoT Security cloud, and update server.

📋 *Although having an IT security telemetry gateway in front of an OT security telemetry gateway lets you block inbound connections to the firewall at the perimeter of the OT network, multiple cascading gateways is not required. If you use a single security telemetry gateway at the perimeter of the OT network, it becomes the proxy between OT firewalls and Palo Alto Networks cloud services in the external network instead of hopping through an IT security telemetry gateway.*

Configure the IT Security Telemetry Gateway

The IT security telemetry gateway is the next-generation firewall that forwards the traffic logs and requests it receives from the OT security telemetry gateway to the logging service, IoT Security, and update server. It would typically be deployed on the network perimeter.

STEP 1 | Configure a next-generation firewall to act as an IT security telemetry gateway.

1. Access the CLI as a superuser or device administrator and enter the following command to enable the firewall to function as a security telemetry gateway (proxy):

   `set system setting paloalto-networks-service-proxy on`

2. Reboot the firewall.

   📋 *When using Panorama to manage firewalls, enter the above command in the Panorama CLI and then reboot Panorama.*

3. Log in to the firewall web interface as a superuser or device administrator and configure two Layer3 interfaces—one on the IT network and the other on the external network. For example, configure ethernet1/1 with IP address 192.168.10.1/24 for the IT network and ethernet1/2 with IP address 1.1.1.1/24 for the external network.

4. Create a loopback interface with an IP address in a different subnet from the other two networks. For example, if the subnet of the IT network is 192.168.10.0/24 and the subnet

of the external network is 1.1.1.0/24, use an IP address that's not in either of these subnets, such as 10.1.2.3, for the loopback interface.

5. Create a virtual router for all three interfaces and add them to it (for example, **vr1**). If the external network interface is a static IP address, add a default route to the gateway in the external network subnet as the next hop.

6. Create a zone for each interface such as **IT**, **external**, and **loop**.

7. Select **Network** > **DNS Proxy** and configure a DNS proxy for the interface in the external zone. For example, create a configuration called **dns-proxy** that does DNS lookups on a DNS server at **8.8.8.8** from **ethernet1/2**.

8. Select **Objects** > **URL Category** and create the following URL group:

   Name: Give the URL list a name; for example, **iot_cloud_traffic**.

   URL List: Add the following URLs (and IP address) to the URL list. These are the only destinations that proxied traffic must be allowed to access.

   • *.paloaltonetworks.com/
   • *.panservicetest.com/
   • ocsp.godaddy.com/
   • certificates.godaddy.com/
   • *.gpcloudservice.com/
   • *.lencr.org/
   • 34.122.191.141

   > 📋 *When using Panorama to manage firewalls, create the URL category as "shared".*

9. Select **Policies** > **Security**, and create a universal policy rule that allows any application from the IT zone to the external zone for destinations in the **iot_cloud_traffic** URL category and position it above other policy rules.

10. Select **Policies** > **NAT**, and create a policy that translates source addresses of devices and interfaces in the IT and loop zones to the IP address of the egress interface in the external zone. In our example, this would be 1.1.1.1, which is the IP address of ethernet1/2.

11. Select **Network** > **Proxy**, click the settings icon for Proxy Enablement, choose **Palo Alto Networks Service Proxy** and then click **OK**.

12. Click the settings icon for Palo Alto Networks Service Proxy Configuration and enter the following:

   Connect Timeout: **5** (default)

   Listening: Enter the name of the IT network interface; for example, **ethernet1/1**.

   Upstream interface: **loopback.1**

   Proxy IP: Enter the IP address of the interface in the IT zone; for example, **192.168.10.1**.

   DNS-Proxy: Enter the name of the DNS proxy you defined previously; for example, **dns-proxy**.

   Allowed URL Category: Enter the name of the allowed URL group you defined previously, for example, **iot_cloud_traffic**.

Next Hop Proxy Server: Leave empty.

Next Hop Proxy Port: Leave empty.

STEP 2 | (Optional) To use IoT Security for device identification, risk assessment, and vulnerability detection in the IT network, subscribe the firewall acting as the IT security telemetry gateway to IoT Security.

> *If you don't want the firewall acting as the IT security telemetry gateway to use IoT Security services in the IT network, it's unnecessary to subscribe it to IoT Security and you can skip this step.*

1. Onboard IoT Security on the IT security telemetry gateway.
2. Install licenses for the logging service and IoT Security on the IT security telemetry gateway and download a device certificate to the IT security telemetry gateway to authenticate its connections with the logging service and IoT Security.
3. Configure the IT security telemetry gateway to support Device-ID and work withIoT Security.

Configure the OT Security Telemetry Gateway

With the IT security telemetry gateway configured and in place, you can next configure the OT security telemetry gateway. The OT security telemetry gateway is the next-generation firewall that forwards the traffic logs it receives from OT firewalls to the IT security telemetry gateway, which in turn forwards them logging service. It also forwards requests from OT firewalls for IP address-to-device mappings, policy rule recommendations, and dictionary files to IoT Security and the update server. It would typically be deployed on the edge of the OT network.

STEP 1 | Configure a next-generation firewall to act as an OT security telemetry gateway.

1. Access the CLI as a superuser or device administrator and enter the following command to enable the firewall to function as a security telemetry gateway (referred to as a *proxy* in PAN-OS):

   set system setting paloalto-networks-service-proxy on

2. Reboot the firewall.

   📋 *When using Panorama to manage firewalls, enter the above command in the Panorama CLI and then reboot Panorama.*

3. Configure two Layer3 interfaces—one on the OT network and the other on the IT network. For example, configure ethernet1/1 with IP address 192.168.100.1 for the OT network and ethernet1/2 with IP address 192.168.10.2 for the IT network.

4. Create a loopback interface with an IP address in a different subnet from the other two networks. Because it's only used for internal routing, you can even use the same IP address as the loopback interface on the IT security telemetry gateway—10.2.3.4, for example.

5. Create a virtual router for all three interfaces and add them to it (for example, **vr1**) and add a default route using ethernet1/2 as the egress interface and 192.168.10.1, the IP address of ethernet1/1 on the IT security telemetry gateway interface, as the next hop.

6. Create a zone for each interface such as **OT**, **IT**, and **loop**.

7. If the next hop security telemetry gateway server is a hostname, select **Network** > **DNS Proxy** and configure a DNS proxy for the interface of the OT security telemetry gateway that's in the IT zone. For example, create a configuration called **dns-proxy** that does DNS

lookups on a local DNS server that the OT security telemetry gateway can reach from ethernet1/2.

> 📋 *If the next hop security telemetry gateway server is an IP address, you don't need to configure a DNS proxy and can skip this step.*

8. Select **Objects** > **URL Category** and create the following URL group:

Name: Give the URL list a name; for example, **iot_cloud_traffic**.

URL List: Add the following URLs (and IP address) to the URL list. These are the only destinations that proxied traffic must be allowed to access.

- *.paloaltonetworks.com/
- *.panservicetest.com/
- ocsp.godaddy.com/
- certificates.godaddy.com/
- *.gpcloudservice.com/
- *.lencr.org/
- 34.122.191.141

> 📋 *When using Panorama to manage firewalls, create the URL category as "shared".*

9. Select **Policies** > **Security**, and create a universal policy rule that allows any application from the OT zone to the IT zone for destinations in the **iot_cloud_traffic** URL category and position it above other policy rules.

> 📋 *Add security policy rules that deny all other outbound connections from the OT network and all inbound connections to the OT network and position them below the rule that allows outbound connections to the destinations in the iot_cloud_traffic URL list.*

10. Select **Network** > **Proxy**, click the settings icon for Proxy Enablement, choose **Palo Alto Networks Service Proxy** and then click **OK**.

11. Click the settings icon for Palo Alto Networks Service Proxy Configuration and enter the following:

Connect Timeout: **5** (default)

Listening: Enter the name of the OT network interface; for example, **ethernet1/1**.

Upstream interface: **loopback.1**

Proxy IP: Enter the IP address of the interface in the OT zone; for example, **192.168.100.1**.

DNS-Proxy: Enter the name of the DNS proxy you defined previously; for example, **dns-proxy**.

Allowed URL Category: Enter the name of the allowed URL group you defined previously, for example, **iot_cloud_traffic**.

Next Hop Proxy Server: Enter the IP address of ethernet1/1 on the IT security telemetry gateway interface; **192.168.10.1** in our example.

Next Hop Proxy Port: 8080

STEP 2 | (Optional) To forward network traffic logs for the OT network from the OT security telemetry gateway as well as from OT firewalls, subscribe the OT security telemetry gateway to IoT Security.

> *If you don't want the firewall acting as the OT security telemetry gateway to use IoT Security services in the OT network, it's unnecessary to subscribe it to IoT Security and you can skip this step.*

1. Onboard IoT Security on the OT security telemetry gateway.
2. Install licenses for the logging service and IoT Security on the OT security telemetry gateway and download a device certificate to the OT security telemetry gateway to authenticate its connections with the logging service and IoT Security.
3. Configure the OT security telemetry gateway to support Device-ID and work withIoT Security.

Configure OT Firewalls

With both the IT and OT security telemetry gateways configured, you can set up the OT firewalls to use the security telemetry gateway chain to access the Palo Alto Networks cloud services necessary to support IoT Security:

- **Logging service** – OT firewalls forward EAL and traffic logs to the logging service, which streams the metadata to IoT Security for analysis to identify devices, assess risk, and detect device vulnerabilities.

- **IoT Security** – OT firewalls retrieve IP address-to-device mappings from IoT Security to enforce Device-ID Security policy rules. OT firewalls also retrieve policy rule recommendations from IoT Security.

- **Update server** – OT firewalls periodically download device dictionary files with a regularly updated list of device attributes used as components in Device-ID Security policy rules.

- **License server** – OT firewalls download activated logging service and IoT Security licenses from the license server.

- **Certificate server** – Firewalls fetch new device certificates from certificate.paloaltonetworks.com and use their existing device certificates—expiring but still valid—to fetch renewed certificates from certificatetrusted.paloaltonetworks.com.

- **Customer Service Portal** and the **hub** – Firewalls connect to the Customer Service Portal to verify admin users and then to the hub to get their role assignments.

STEP 1 | Configure a next-generation firewall to act as an OT firewall.

1. Select **Device** > **Setup** > **Interfaces** > **Management**, configure MGT interface with an IP address on the OT network, and enter the IP address of the OT security telemetry gateway interface in the OT zone as its default gateway; for example:

    IP Type: Static

    IP Address: 192.168.100.2

    Netmask: 255.255.255.0

    Default Gateway: 192.168.100.1

    The OT firewall uses the management interface to onboard IoT Security and fetch certificates and licenses, forward various traffic logs to the logging service, request IP address-to-device mappings and policy rule recommendations from IoT Security, and download dictionary files from the update server.

    > *You can also configure the OT firewall to use one of its Ethernet interfaces when initiating connections through the chain of security telemetry gateways. If you do, you must configure service routes to instruct the firewall to use this interface instead of the management interface. In the service route configuration, select Palo Alto Networks Services, Data Services, and IoT.*

2. Configure interfaces, security zones, and security policy rules as necessary to collect network traffic metadata for IoT Security to analyze. PAN-OS provides various options and you'll need to use whatever methods make sense for your network topology; for example:

    Virtual wire to capture OT traffic – Create a virtual wire zone and a virtual wire object that links two virtual wire interfaces. Add either an intrazone or universal policy rule that allows traffic between devices within the same zone, and enable logging and log forwarding on the rule. Consider placing one or more OT firewalls with this configuration on the OT network at one of the OT Purdue levels (0-3) to capture network traffic at this level and forward traffic logs to the OT security telemetry gateway.

    Tap interface to collect traffic from downstream switches – Create a tap zone with a tap interface to receive traffic from a mirror port on downstream switches. This will capture traffic at other Purdue levels that don't reach OT firewalls, which can then forward it to the logging service.

    Layer 3 interface to collect traffic from a ERSPAN port on downstream switches – Create a Layer 3 zone with a Layer 3 interface on the OT firewall. Configure your switches to use Encapsulated Remote Switched Port Analyzer (ERSPAN) to send mirrored traffic through a Generic Routing Encapsulation (GRE) tunnel to the IP address of the OT network interface on the OT security telemetry gateway. After decapsulating the traffic, the OT security telemetry gateway generates various traffic types of logs and forwards them to the IT

security telemetry gateway, which then forwards them to the logging service where IoT Security can access them for analysis.

3. Select **Device** > **Setup** > **Services**, enter the following settings in the Proxy Server section and leave the other settings with their default values:

Proxy Server

- **Server**: Enter the IP address of the OT security telemetry gateway interface in the OT zone; for example, 192.168.100.1, which is the IP address of the OT security telemetry gateway ethernet1/1 interface.
- **Port**: **8080**
- **Use proxy to send logs to Strata Logging Service**: (select)

4. Select **Policies** > **Security**, and create a universal policy rule that allows the following applications from OT network zones to any zones and position it above other policy rules:

google-base

paloalto-device-telemetry

paloalto-iot-security

paloalto-logging-service

paloalto-shared-services

STEP 2 | Subscribe the OT firewall to IoT Security.

1. Onboard IoT Security on the OT firewall.

2. Install licenses for the logging service and IoT Security on the OT firewall and download a device certificate to the OT firewall to authenticate its connections with the logging service and IoT Security.

3. Configure the OT firewall to support Device-ID and work with IoT Security.

IoT Security Integration with Prisma Access

Prisma Access uses a cloud-based infrastructure that lets you avoid the challenges of sizing firewalls and computing resource allocation while securing remote networks and mobile users. To identify IT and IoT devices at your remote sites, detect IoT device vulnerabilities, and discover threats posed to these devices and the network, Prisma Access can integrate with IoT Security through a purchased add-on. In addition, IoT Security also provides Prisma Access with policy rule recommendations through Panorama to permit only acceptable network behavior and block anomalous behavior from your IoT devices.

For IoT Security to identify IT and IoT devices, and analyze risk levels and detect security alerts on IoT devices, it must be able to access network traffic metadata. The more data it has to work with, the more accurate and faster it can be. Therefore, it's critical to do two things to collect as much traffic metadata as possible. First, design your network strategically so that Prisma Access sees all traffic from your remote sites, including DHCP traffic. Then apply policy rules to as much traffic as you can and enable logging and log forwarding on these rules to send traffic metadata to Strata Logging Service.

DHCP traffic is particularly important to IoT Security. It provides IoT Security with useful data, including a mapping of the IP address to MAC address of each DHCP client, which is a critical element of the IP address-to-device mappings used for device identification. To obtain this data, ensure that a DHCP server is in your data center or in a similar centralized site and a DHCP relay agent is on the customer premises equipment (CPE) where the remote network connection terminates at each site. Each relay agent forwards the DHCP messages it receives from DHCP clients through the Prisma Access service infrastructure to the IP address of the DHCP server. On the policy rule allowing DHCP traffic from the remote sites to the DHCP server, be sure logging and log forwarding are enabled so that Prisma Access sends DHCP traffic logs to Strata Logging Service. In fact, if you have not already done so, enable logging and log forwarding on all policy rules. With log forwarding enabled, Prisma Access sends its logs through Strata Logging Service, which then streams metadata to IoT Security for analysis.

Logical architecture of Prisma Access integration with IoT Security

Prisma Access cannot forward logs to IoT Security for Layer 2 traffic or Layer 3 traffic where both the source and destination are in the same site because such traffic never reaches it. Without ARP and DHCP traffic metadata in particular, identifying devices might take IoT Security longer and its confidence might be lower than it otherwise would be. To counter this, consider deploying SD-WAN ION devices at remote sites where they can log these types of traffic and forward their logs to Strata Logging Service for IoT Security to access. By integrating IoT Security with both Prisma Access and SD-WAN, *IoT Security can gain visibility into traffic that flows between sites and the Internet as well as traffic that stays within a site.*

After IoT Security has sufficient information to identify devices from their network behavior, it provides Prisma Access with IP address-to-device mappings and Panorama with policy recommendations that the Panorama administrator can import and then push to Prisma Access to enforce policy on IoT device traffic. In addition, Prisma Access downloads device dictionary files from the update server. The device dictionary lists various device attributes with which the Panorama administrator can construct Security policy rules. The combination of IP address-to-device mappings, policy recommendations, and device dictionary files comprise the elements of the Device-ID feature introduced in PAN-OS 10.0.

Required Panorama Configuration

Check that you have enabled Enhanced Application Logs on your log forwarding profiles.

1. Log in to Panorama and select **Objects** > **Log Forwarding** under the **Remote_Network_Device_Group** device group or a parent device group.

2. Open your log forwarding profiles and make sure that **Enable enhanced application logging to Strata Logging Service** is selected.

Requirements for using IoT Security with Prisma Access

To use the IoT Security add-on with Prisma Access, check that your deployment meets the following requirements:

1. Prisma Access is running the Prisma Access 2.0-Innovation release or later.

2. You have purchased and activated licenses for Strata Logging Service and the IoT Security add-on for Prisma Access.

   If you are a new Panorama-managed Prisma Access customer as of August 2022, activate new Prisma Access licenses through the Prisma SASE platform.

   If you are an existing Panorama-managed Prisma Access customer from before August 2022, your Prisma Access tenant will be transitioned from the hub to the Prisma SASE platform. After the transition, you will no longer see a Prisma Access app title on the hub. However, there will be a button on the hub to navigate to sase.paloaltonetworks.com where you can activate new Prisma Access licenses through the Prisma SASE platform. Until then, continue to manage your deployment as you've been doing.

3. The deployment of Prisma Access in a particular region requires that the Strata Logging Service instance and IoT Security application it works with to be in a particular location as well. The following table shows the relationship of Prisma Access deployments in different regions to the locations of Strata Logging Service and IoT Security.

	Prisma Access	Strata Logging Service	IoT Security
Americas	Canada	Canada	Canada
	United States	United States	United States
European Union	France	France	Germany
	Germany	Germany	Germany
	Italy	Italy	Germany
	Poland	Poland	Germany
	Spain	Spain	Germany
	Netherlands	Netherlands	Germany
	Switzerland	Switzerland	Switzerland
	United Kingdom	United Kingdom	United Kingdom
Asia-Pacific	Australia	Australia	Australia
	China	China	Singapore

	Prisma Access	Strata Logging Service	IoT Security
	India	India	Singapore
	Indonesia	Indonesia	Singapore
	Japan	Japan	Japan
	Singapore	Singapore	Singapore

4. You're using Panorama 10.0 or later to manage Prisma Access.

> *With a mixed deployment of Prisma Access and on-premises next-generation firewalls, you must use the same Panorama management system to manage them and the same IoT Security tenant for both.*

5. DHCP is being served from a data center or from some other central site.

6. The Prisma Access infrastructure provides routing from remote sites to data center resources, which include the DHCP server.

7. A DHCP relay agent on the VPN terminator at all remote sites points to the IP address of the DHCP server in the data center.

8. Security policy rules in Prisma Access control traffic to the Internet, the data center, and other remote sites. Logging is enabled on these policies and Prisma Access forwards logging data to Strata Logging Service, which streams it to IoT Security.

> *IoT Security uses Enhanced Application logs (EALs), traffic logs (which include DHCP traffic), threat logs, and wildfire logs. Make sure that your policy rules have logging enabled and are forwarding EALs and traffic logs to Strata Logging Service. Although the last two log types are not required for IoT Security to function, we recommend getting licenses for threat prevention and Wildfire and forwarding their logs as well because they help improve risk assessment and malware detection.*

Once these requirements are met, use IoT Security to monitor traffic metadata, identify IoT devices, detect vulnerabilities, discover threats, and prepare policy rule recommendations. Import policy rule recommendations from IoT Security into Panorama or configure Device-ID policy rules directly in Panorama and then push them to Prisma Access for policy enforcement on IoT device traffic.

IoT Security Licenses

You have several options when a license for an IoT Security subscription or a third-party integration add-on expires. If you no longer want a firewall to subscribe to IoT Security services or integrate with third-party systems, you can let the license expire. If you do want to continue using these services or integrations, you can extend trial and eval licenses, renew paid licenses, and even convert licenses from one type to another.

License Extensions

Before buying IoT Security, you might first try it out and evaluate it. The initial term of a trial or eval (evaluation) license is 60 days and can be extended in 30-day increments. To extend the trial or eval term, request a 30-day extension through your Palo Alto Networks sales representative or sales engineer.

License Renewals

As a paid license approaches its expiration date, you can renew it so that there's no break in service, the next license beginning immediately after the current license ends. You can renew the following licenses:

- IoT Security Subscription lab license
- IoT Security Subscription prod (production) license
- IoT Security, Doesn't Require Data Lake (DRDL) Subscription lab license
- IoT Security, DRDL Subscription prod license
- Basic IoT Security Third-party Integrations Add-on license
- Advanced IoT Security Third-party Integrations Add-on license

To renew any of these licenses, contact your Palo Alto Networks sales representative.

License Conversions

A license conversion is the change of one license type to another. The license can be for an IoT Security subscription or a third-party integration add-on.

> You can convert an IoT Security license on a firewall from trial to prod, but not from eval to prod. An eval license is for an eval firewall, which is Palo Alto Networks property and loaned out for temporary use. However, if you create an IoT Security tenant URL for eval licenses on eval firewalls and then replace them with prod licenses on prod firewalls, you can continue using the same IoT Security tenant URL.

Palo Alto Networks supports the following conversions:

IoT Security license conversions

- Trial > Prod
- IoT Security Subscription > IoT Security, Doesn't Require Data Lake (DRDL) Subscription

- IoT Security, DRDL Subscription > IoT Security Subscription

  Activate a Strata Logging Service instance before converting from a subscription that doesn't require a data lake to one that does.

IoT Security Third-party Integrations Add-on license conversions

- Basic > Advanced

- Advanced > Basic

All conversions can be done after the current license expires at the end of its term, but only conversions considered to be upgrades are allowed midterm. Midterm conversions take place immediately, replacing the previous term with the new term. The following conversions are considered to be upgrades:

IoT Security license upgrades

- Trial version of any type of license > Prod version of any type of license

- IoT Security Subscription > IoT Security, DRDL Subscription

IoT Security Third-party Integrations Add-on license upgrades

- Trial version of any type of add-on > Prod version of any type of add-on

- Basic > Advanced

  Converting IoT Security licenses from trial to prod generates a new purchase order with a link to a new onboarding workflow. During the onboarding process, you can select the existing IoT Security tenant you were previously using for trial purposes. The rest of the onboarding workflow follows the same mechanism for activating prod licenses on firewalls as it did for activating trial licenses.

To convert any licenses, contact your Palo Alto Networks sales representative.

Offboard IoT Security Subscriptions

There are three ways to offboard IoT Security services from a firewall:

- Deactivate the IoT Security license on a firewall and optionally transfer it to another firewall
- Transfer a firewall from one customer support portal (CSP) account to another
- Let the subscription expire

Deactivate Firewalls and Transfer Licenses

If you want to remove an IoT Security license from a firewall—and perhaps then use the license on another firewall—you can do so on the Customer Support Portal.

STEP 1 | Log in to your Customer Support Portal account.

STEP 2 | Disassociate IoT Security licenses from one or more firewalls.

1. Select **License Management** > **Activated Licenses**, select the license-to-firewall associations that you want to sever based on firewall serial numbers, and then **Deactivate Licenses**.

2. **Confirm** the deactivation.

> *If you want to apply the deactivated licenses to other firewalls and you have multiple IoT Security license purchase orders, note the number of available licenses in the orders on the Activate Products page before confirming the deactivation. Then when you return to this page after deactivating licenses, you can tell which order they were returned to because the license number will have increased.*

This dissociates the selected IoT Security licenses from the firewall serial numbers and returns them to the pool of available licenses in the original order on the Activate Products page.

Get Started with IoT Security

STEP 3 | Associate licenses with other firewalls or reassociate them with the same firewalls.

1. Select **Activate Products** > **Ready for Activation** and then click **Activate Now** for the order with licenses to activate.

2. Follow the workflow described in Onboard IoT Security.

   When you reach the point in the onboarding workflow when you select firewalls to subscribe to IoT Security, you can see the length of time remaining for each license in the Purchased Term drop-down list. If you want to apply the same license that you just

IoT Security Administrator's Guide February 2024   **112**   ©2024 Palo Alto Networks, Inc.

deactivated to another firewall, you'll notice that its remaining length of unused time will be shorter than other licenses that haven't yet been put in service. For example, if the original order contains licenses valid for three years and you used a license for one year before deactivating it, you can easily spot it because its remaining validity period will be the only one listed as just two years.

Transfer Firewalls between CSP Accounts

If you have two CSP accounts or are an MSSP managing multiple accounts, you can transfer a firewall from one account to another, perhaps because you're moving it to a different location managed by a different team with their own account. When you transfer the firewall, all its licenses are transferred along with it. To do this, log in to the CSP and click **Devices**. Find the device you want to transfer, click its serial number to open a device details pane for it, and then click **Transfer Ownership**. In the Device Transfer dialog box that appears, enter the destination email address of the owner of the account to which you're transferring the firewall.

Let the IoT Security Subscription Expire

When a firewall no longer has an IoT Security subscription because it expired (and there is no pending license renewal), IoT Security services for that firewall stop and the connection between IoT Security and the firewall is terminated. IoT Security unsubscribes from the firewall log feed. As a result, it stops receiving and processing logs from that firewall. The firewall stops receiving new policy recommendations and IP address-to-device mappings, and it clears its cached mappings after 200 minutes (about three hours). At that point, none of the device-based policy rules using Device-ID will work and should be removed from your policy set. An efficient way to remove them is to check the Source Device and Destination Device columns on the **Policies** > **Security** page and remove all rules that have entries in either of these two columns.

114

IoT Security Overview

Understand the fundamentals of IoT Security, what it does, and how it works.

- Introduction to IoT Security
- IoT Security Integration with Next-generation Firewalls
- IoT Security Portal
- Vertical-themed Portals
- Device-to-Site Mapping
- Sites and Site Groups
- Networks
- Network Visualizations
- Create a Visualization Map
- View Data in a Visualization Map
- Reports
- IoT Security Integration Status with Firewalls
- IoT Security Integration Status with Prisma Access
- Data Quality Diagnostics
- Authorize On-demand PCAP
- IoT Security Integrations with Third-party Products
- IoT Security and FedRAMP

Introduction to IoT Security

IoT Security is an on-demand cloud subscription service designed to discover and protect the growing number of connected "things" on your network. Unlike IT devices such as laptop computers that perform a wide variety of tasks, IoT devices tend to be purpose-built with a narrowly defined set of functions. As a result, IoT devices generate unique, identifiable patterns of network behavior. Using machine learning and AI, IoT Security recognizes these behaviors and identifies every device on the network, creating a rich, context-aware inventory that's dynamically maintained and always up to date.

After it identifies a device and establishes a baseline of its normal network activities, it continues monitoring its network activity so it can detect any unusual behavior indicative of an attack or breach. If it detects such behavior, IoT Security notifies administrators through security alerts in the portal and, depending on each administrator's notification settings, through email and SMS notifications.

IoT Security also uses those behaviors and device identities to automatically generate security policy rule recommendations that allow IoT devices to continue doing normal network activities and block them from doing anything unusual. Panorama or next-generation firewalls can then import these policy rules and enforce them.

> 📋 *For Panorama-managed firewalls that have an IoT Security subscription requiring Strata Logging Service – Panorama can only import policy rule recommendations if it was used to* onboard its managed firewalls to Strata Logging Service.

The firewall collects metadata from the network traffic of IoT devices, generates Enhanced Application logs (EALs), and forwards them to the logging service. The IoT Security cloud then extracts metadata from these logs for analysis and employs AI and machine-learning algorithms to detect and identify IoT devices using its patented three-tier deep-learning engine:

Tier 1: Device category—IoT Security first identifies the category to which an IoT device belongs. For example, it might identify network behaviors common to all security cameras.

Tier 2: Device profile—IoT Security next constructs a profile of the device, learning its vendor, make, and model. For example, it might discover that the camera behaves in ways that uniquely identify it, such as checking a particular server for software updates for example.

Tier 3: Device instance—IoT Security continues its analysis until it discerns behaviors unique to a specific instance of the identified security camera.

IoT Security looks at over 200 parameters in network traffic metadata, including DHCP option 55 parameter lists, HTTP user agent IDs, protocols, protocol headers, and a host of others. It matches the network traffic patterns of new devices with those of previously identified devices to identify the same types or similar types of devices, even those it is encountering for the first time.

Depending on various factors such as how much network traffic IoT devices generate and how varied their behavior patterns are, IoT Security typically identifies most IoT devices with a high level of confidence during the first day it starts accessing metadata from the logging service. After that, IoT Security continues to increase the number of confidently identified devices until it identifies all or nearly all of them. During this time, you can log in to the IoT Security portal to check that the device inventory is being populated and monitor its progress.

A confidence score indicates the level of confidence IoT Security has in its identification of a device. IoT Security has three confidence levels based on calculated confidence scores: high (90-100%), medium (70-89%), and low (0-69%).

In addition to using machine learning (ML) to observe network traffic and extract various attributes to identify devices and detect anomalous behaviors, IoT Security employs an ML-based model to check for SQL content injected into HTTP URLs, a technique commonly used in SQL vulnerability exploits. By using an ML-based model instead of a model based on rules, IoT Security can find certain patterns of injected SQL content even without specific signatures.

IoT Security Integration with Next-generation Firewalls

The IoT Security solution involves the integration of three key architectural components to process network data:

- **Palo Alto Networks next-generation firewalls** collect device data and send it to the logging service.

- **The logging service** uses a cloud-based log-forwarding process to direct the logs from firewalls to destinations like IoT Security and Strata Logging Service. Depending on the type of IoT Security subscription you have, the logging service either streams metadata to your IoT Security account and Strata Logging Service instance or just to your IoT Security account.

- **IoT Security** is an app that runs on a cloud-based platform in which machine learning, artificial intelligence, and threat intelligence are used to discover, classify, and secure the IoT devices on the network. The app ingests firewall logs with network traffic data and provides Security policy recommendations and IP address-to-device mappings to the firewall for use in Security policy rules. Administrators access the dynamically enriched IoT device inventory, detected device vulnerabilities, security alerts, and recommended policy sets through the IoT security portal.

The IoT Security app integrates with next-generation firewalls through Device-ID, which is a construct that uses device identity as a means to apply policy. The integration uses three mechanisms.

- **Device dictionary** – This is an XML file that IoT Security generates and makes available for Panorama and firewalls to import. The dictionary file provides the Panorama and firewall administrator with a list of device attributes for selection when importing recommended Security policy rules from IoT Security and when creating rules themselves. These attributes are profile, category, vendor, model, OS family, and OS version and are for both IoT and traditional IT devices. Although it's not possible to download a device dictionary file, you can see the release notes summarizing the new content added to a file that your firewall has imported. To do this, log in to the PAN-OS web portal, select **Device** > **Dynamic Updates** and then click **Release Notes** for the device dictionary file you want to learn about.

- **Policy rule recommendations** – After an IoT Security administrator creates a set of Security policy rules based on traffic from IoT devices in the same device profile, a firewall administrator can import them as recommendations for use in its policy set.

- **IP address-to-device mappings** – These mappings tell firewalls which attributes a device with a particular IP address has. When traffic to or from that IP address reaches a firewall, it checks if one of its attributes matches a policy and, if so, the firewall applies the policy. IoT Security sends IP address-to-device mappings to firewalls for both IoT and IT devices if the confidence score for device identities is high (90-100%) and they've sent or received traffic within the past hour.

The goal of Device-ID is to leverage the intelligence of IoT Security to enforce firewall policy on IoT devices.

Device-ID

PAN-OS 10.0 introduces a new concept for policy enforcement: Device-ID. Device-ID is a way to enforce policy rules based on device attributes. IoT Security provides the firewall with a device dictionary file containing a list of device attributes such as profiles, categories, vendors, and

models. For various attributes in the dictionary file, it lists a set of entries. For example, three entries for the profile attribute might be Advidia Camera, BK Medical UltraSound Machine, and Carefusion Infusion Pump Base Station.

When configuring a Security policy rule, firewall administrators have the option to select device attributes from the device dictionary. If they select **profile**, they can choose one of the profile entries: **Polycom IP Phone**, for example. The policy rule then applies to all devices that match this profile. But how does the firewall know what the profile is for a device? It knows this from the IP address-to-device mappings that IoT Security also gives the firewall. These mappings identify attributes for each device. When traffic from an IP address that's mapped to a device attribute specified in the policy rule reaches the firewall, the policy rule lookup will find a match with this rule and apply whatever action it enforces.

A firewall downloads a device dictionary file from the update server. The dictionary file populates entries in all the Device-ID attribute lists for profile, category, vendor, and so on. These attribute entries are then available for use as policy rule configuration elements. The firewall administrator next configures a firewall policy rule using the profile attribute "Polycom IP Phone". After a Polycom Trio 8800 device joins the network and IoT Security identifies it, IoT Security provides the firewall with an IP address-to-device mapping for it. The two key elements in the mapping for this example are its device profile (Polycom IP Phone profile, highlighted in yellow) and its IP address (10.1.2.3, highlighted in blue). When traffic from the Polycom Trio 8800 device at 10.1.2.3 reaches the firewall, it does a Device-ID policy rule lookup, finds that the profile for the device at this IP address matches one specified in a policy rule, and then applies the rule.

> 📄 *If a firewall becomes disconnected from IoT Security, the firewall retains its IP address-to-device mappings and continues enforcing Device-ID policy rules with them until the connection is re-established.*

Every next-generation firewall model has the same maximum of 1000 unique Device-ID objects.

The maximum of 1000 Device-ID objects is not the same as that for IP address-to-device mappings. The maximum number of IP address-to-device mappings varies based on firewall model

and is the same as the User-ID maximums listed in the + Show More sections for each firewall model on the Product Selection page

More information about the Device-ID feature is in the PAN-OS Administrator's Guide.

Device Dictionary

The device dictionary is an XML file for firewalls to use in Security policy rules. It contains entries for the following device attributes: profile, category, vendor, model, OS family, and OS version. These entries come from devices across all IoT Security tenants and are completely refreshed on a regular basis and posted as a new file on the update server. If there are any changes to a dictionary entry, a revised file will be posted on the update server so that Panorama and firewalls will automatically download and install it the next time they check the update server, which they do automatically every two hours.

IP Address-to-device Mappings

After IoT Security identifies a device, it bundles the following set of identifying characteristics about it:

- IP address
- MAC address
- Hostname
- Device type
- Device category
- Device profile
- Vendor
- Model
- OS family
- OS version
- Risk score
- Risk level

Firewalls poll IoT Security for these IP address-to-device mappings for use in policy enforcement. A firewall polls for new or modified mappings every second, and IoT Security returns mappings that it has identified with high confidence (a confidence score of 90-100%) for devices that were active within the last hour. For each IP address-to-device mapping that a firewall receives, the firewall generates an entry in its host information profile (HIP) Match log.

If IoT Security discovers duplicate IP address-to-device mappings—that is, there are two IP addresses mapped to the same device MAC address—it resolves it to the MAC address with the latest network activity.

There is no time limit for how long a firewall retains IP address-to-device mappings. It only begins deleting them when its cache fills up, starting with the oldest first.

Policy Rule Recommendations

You can generate Security policy rule recommendations based on the normal, acceptable network behaviors of the IoT devices in the same device profile and manually import them into firewalls

for enforcement. PAN-OS 8.1 and later supports the importing of IoT Security policy rule recommendations.

> 📋 *For Panorama-managed firewalls that have an IoT Security subscription requiring Strata Logging Service – Panorama can only import policy rule recommendations if it was used to onboard its managed firewalls to Strata Logging Service.*

Firewall and Panorama Communications Related to IoT Security

IoT Security communications from firewalls without Panorama management:

- Firewalls download device dictionary files from the update server at updates.paloaltonetworks.com on TCP port 443.

- Firewalls forward logs to the logging service on TCP ports 443 (for Enhanced Application logs) and 3978 (for all other firewall logs).

> 📋 *For details about the ports and FQDNs required for next-generation firewalls to communicate with the logging service, see* Strata Logging Service Getting Started.

- Firewalls retrieve IP address-to-device mappings and policy recommendations from IoT Security on TCP port 443. Depending on their region, they use one of the following edge services URLs:

  - United States: iot.services-edge.paloaltonetworks.com

  - Canada: ca.iot.services-edge.paloaltonetworks.com

  - EU: eu.iot.services-edge.paloaltonetworks.com

  - Switzerland: ch.iot.services-edge.paloaltonetworks.com

  - United Kingdom: uk.iot.services-edge.paloaltonetworks.com

  - APAC: apac.iot.services-edge.paloaltonetworks.com

  - Japan: jp.iot.services-edge.paloaltonetworks.com

  - Australia: au.iot.services-edge.paloaltonetworks.com

The following table summarizes the relationship of different data lake regions/ingestion regions with IoT Security application regions:

	Data Lake Region/Ingestion Region	IoT Security Application Region
Americas	Canada	Canada, United States*
	United States	United States
	FedRAMP	FedRAMP
European Union	France	Germany
	Germany	Germany

	Data Lake Region/Ingestion Region	IoT Security Application Region
	Italy	Germany
	Netherlands	Germany
	Poland	Germany
	Spain	Germany
	Switzerland	Switzerland, Germany*
	United Kingdom	United Kingdom, Germany*
Asia-Pacific	Australia	Australia, Singapore*
	India	Singapore
	Indonesia	Singapore
	Japan	Japan
	Singapore	Singapore

*Switzerland and the United Kingdom were added as IoT Security application regions on 7/31/2023. When onboarding IoT Security after this date to existing firewall deployments established before it, the firewalls continue to use **Germany** as the IoT Security application region. When onboarding IoT Security to new deployments in Switzerland or the United Kingdom established after 7/31/2023, the firewalls will use the local IoT Security application region for each country.

A similar situation exists in Canada, which continues to use **United States – Americas** as the IoT Security application region for deployments existing before 1/25/2023 and **Canada** for new deployments after this date. Likewise, deployments existing before 10/25/2022 in Australia still use the IoT Security application in **Singapore** while new deployments after this date use **Australia**.

- During the certificate exchange between a firewall and the edge server in front of the IoT Security cloud, they verify each other's certificates. The firewall validates the certificate it receives by checking these sites:

  - o.lencr.org
  - c.lencr.org

  Communications to these sites occur over HTTP on TCP port 80.

IoT Security communications from Panorama:

- A Panorama management server imports policy recommendations from IoT Security through the same URLs listed above that firewalls use. When validating the certificate the edge server presents, Panorama checks the same sites listed above that firewalls check.

    Firewalls under Panorama management still contact IoT Security through regional edge services URLs for IP address-to-device mappings, they still download device dictionaries from the update server, and they still forward logs to the logging service.

- A Panorama management server sends queries for logs to the logging service on TCP port 444.

IoT Security Portal

After you onboard IoT Security, activate IoT Security licenses on your firewalls, and deployed them so that they can feed data to the logging service, you're ready to access the IoT Security portal and begin using it. Log in with your account credentials for the Palo Alto Networks Customer Service Portal.

IoT Security uses Single Sign-on (SSO) to verify your login.

> *The IoT Security portal fully supports Google Chrome and partially supports Microsoft Edge, which means the portal is expected to be usable but might not look exactly as designed. It does not officially support Microsoft Internet Explorer, Apple Safari, or any other type of browser.*

The portal interface is grouped into several sections.

Navigation – The items in the left navigation menu are loosely grouped into four sections. The first section is organized around visibility: Dashboard, Assets, and Networks. The next section is security related: Alerts, Vulnerabilities, and Policy Sets. The third section is where you configure and review settings to integrate IoT Security with third-party products: Integrations. Finally, the last section is where you can check logs, reports, firewalls, and data quality, and manage administrative settings: Logs & Reports, Settings, and Administration.

Use the left navigation menu to navigate to different pages in the IoT Security portal. When there are data filters at the top of a page, use them to control the data that appears on the page by site, device type, and time period.

Under the navigation menu is a set of administrative tools:

- Give feedback – Leave feedback for IoT Security developers.
- Help – Open the Customer Support Portal.
- User name (first and last name from the user's contact information) – When you click the name, these options appear:
  - Preferences – Modify your contact information, time zone, idle session timeout, alert sound (that is, control if an audible alert sounds whenever IoT Security detects new Security alerts), and SMS and email notification settings.
  - Resource Center – See status notifications about firewall logs, and learn about IoT Security through recommended resources and useful links
  - Dark Theme/Light Theme – Switch between dark and light UI display themes.
  - Log out – Log out of your administrative session.
- App Switcher – Take a shortcut to other Palo Alto Networks applications through the hub.

Search – At the top of the page to the right of the page title bar is a search field where you find data by entering keywords to search for devices, alerts, vulnerabilities, and external destinations.

Data Filters – Below the page title bar and search field on many pages is a set of filters that control the data that the IoT Security portal displays on each page. The filter system consists of global filters and local, page-specific filters. Global filter settings persist while you navigate among different pages with various filters appearing as appropriate per page. For example, there's no time filter on the Vulnerabilities page, there are additional filters on the Devices and Security Alerts pages, and there are no filters at all on the User Accounts page. Global filters have default values but can also be customized. Modified and added filters appear in the UI as blue instead of black, so you can easily tell them apart from the default ones. If a page has a default local filter, it appears among the other global filters at the top of the page. For example, the Security Alerts page applies an Active Alerts filter by default, so this local filter automatically appears among the global filters whenever you open the Security Alerts page. In addition, there are also page filters that are only applicable to the data on a particular page. When you scroll down a page, both the global and page filters continue to remain in view in the upper right of the title bar.

Query Builder – Next to the data filters is the query builder. Use it to find information about devices, alerts, and vulnerabilities by constructing queries out of various components. A single query can combine devices and security alerts or devices and vulnerabilities. For example, you might query for all IoT devices from a particular vendor that raised a particular alert, or you can query for all IoT devices in a particular profile that have a specific vulnerability. For example, this query shows if the devices in the device profile for APC (Schneider Electric) Smart Power Supply support SNMPv1:

```
entity = device, Time Range = "month", Device Type = "All
  IoT", [device] Profile = "APC(Schneider Electric) Smart
  PowerSupply" [vulnerability] Vulnerability = "SNMPv1 Usage"
```

The results of the query show that 20 IoT devices support SNMPv1 and which ones they are.

126

The query tool uses the logic of "AND" between expressions using the operators = (equals), != (doesn't equal), and IN (includes). For example, the following query fetches data where Time Range = "week" **AND** Device Type = "All IoT" **AND** [vulnerability] Severity IN ("High", "Critical"):

```
entity = device, Time Range = "week", Device Type = "All IoT",
  [vulnerability] Severity IN ("High", "Critical")
```


You can save queries so you don't have to recreate ones used repeatedly. To save a query, click the ribbon bookmark icon to the right of the Query field, and give it a name. For example, if you regularly check the number of IoT devices running a Windows OS that were actively on the network during the past week and that have no endpoint protection or outdated protection, create this query and save it with a name such as Noncompliant Windows IoT devices:

```
entity = device, Time Range = "week", Device Type = "All IoT",
  [device] Endpoint Protection IN ("Not protected", "Outdated"),
  [device] OS = "Windows"
```


When you want to use the query again, just click the bookmark icon and then click the name in the list of previously saved queries and filters. You can also edit entries in this list and delete them.

You cannot save queries from any of the dashboards, such as the Executive Summary.

The query tool has numerous parameters you can use to find whatever nugget of data you want. For example, enter the following query to check which devices were in a vulnerability scan report:

```
Entity = device, Time Range = "1 Year", Device Type = "All IoT",
  [scanReport] Scan Report = "yes"
```


By looking at the Device Details page of devices in the results of the query and clicking **Vulnerability Report Ready**, you can download the report as a PDF to your system where you can keep and read it.

To help you get started using the query builder, IoT Security provides a collection of example templates for common queries. Study these preconfigured queries to learn query builder capabilities, use them as they are, or use them as models for building similar queries of your own.

To see the preconfigured example queries, click **Query** under the page title bar and then click the **Query Bookmarks** icon.

The preconfigured templates differ somewhat based on the vertical theme that's active on your IoT Security portal. Each vertical theme has five example templates. Here's an example for each theme:

Enterprise IoT Security Plus

- Name: [Example] This Week's Active Insecure-Login Alerts
- Query: Entity="alert", Time Range="1 Week", Alert Status="Active Alerts", Alert Type IN ("insecure login", "unsecure login", "Unsecure login")

- Summary: This queries IoT Security for all active alerts related to insecure logins over the past week.

Industrial IoT Security

- Name: [Example] Critical Risk Internet Connected Industrial Devices
- Query: Entity="device", Time Range="1 Year", Device Type="Industrial", [device] Risk = "Critical", [device] Internet Access="yes"
- Summary: This queries IoT Security to show all industrial IoT devices that had a critical risk level and Internet access within the past year.

Medical IoT Security

- Name: [Example] Risky Internet Connected IoT Devices
- Query: Entity="device", Time Range="1 Year", Device Type="All IoT", [device] Risk IN ("High", "Critical"), [device] Internet Access="yes"
- Summary: This queries IoT Security to show all IoT devices that had a high or critical risk level and Internet access within the past year.

You can edit the expressions that constitute a query template and the template name, perhaps saving a modified query with a new name to reuse later. You can also delete the example templates.

Announcements – Toggle open and closed a vertical panel on the right side of the UI with information about recent feature releases and important security announcements.

Manage dashboards – When your portal theme has multiple dashboards, such as Medical IoT Security, you can control which one is the default, which ones are available in adjacent tabs for quick access, and which ones are hidden. Recognizing that users of the IoT Security portal function in different roles, IoT Security lets you set your own preferences to best suit your needs and thereby increase efficiency and productivity.

1. To manage the display of the various dashboards, select **Dashboards** > **Manage Dashboards**.

2. In the **Manage Dashboards** drop-down menu, select the check boxes of dashboards you want to display as a tabbed dashboard for faster access. Clear the check boxes of those you don't want displayed as a tabbed dashboard.

The left-to-right order of tabbed dashboards displayed in the main window corresponds to the top-to-bottom order of dashboards listed in the drop-down menu with the pinned (preferred) dashboard appearing on the far left.

3. To set the default dashboard to display first when navigating to **Dashboards** in the left navigation panel, click the pushpin icon next to a dashboard name in the **Manage Dashboards** drop-down menu.

If you change the portal theme to a vertical that doesn't include your pinned dashboard, the default dashboard for that vertical becomes the new pinned dashboard.

4. To open a new browser tab or window showing security alerts and vulnerabilities, click **View Alerts Overview** and **View Vulnerabilities Overview**.

Vertical-themed Portals

The IoT Security portal changes to better serve users in different industries. The portal theme that users in a given IoT Security tenant see depends on two choices:

- The IoT Security product chosen upon purchase
- The theme chosen by an IoT Security tenant owner

Portal Themes

IoT Security provides four differently themed portals for enterprise, industrial, and medical verticals:

- Enterprise IoT Security Plus
- Enterprise IoT Security
- Industrial IoT Security
- Medical IoT Security

Enterprise

IoT Security offers two products for enterprise IoT: Enterprise IoT Security Plus and Enterprise IoT Security.

Enterprise IoT Security Plus is the solution for commercial enterprises and government organizations. It lets you see and secure every IoT device in your enterprise organization to meet NIST guidelines. It also helps prevent your IoT devices from becoming the target of cyberattacks. With Enterprise IoT Security Plus, you can do the following:

- Automatically classify devices with over 50 device attributes
- View, edit, confirm, and reclassify devices
- Add devices with static IP addresses
- See your IP address structure and device distribution
- See sites for firewalls and devices
- Generate reports for devices, network behaviors, and security risks
- Integrate with multiple third-party products
- See applications that devices use
- Import policy rule recommendations to firewalls
- Get security alerts for anomalous network activity
- Assess risk and device vulnerabilities
- (Optional) Retain traffic logs

The Security Dashboard, which provides quick access to information about device inventory, alerts, and risks, is shown below. It appears in the Enterprise IoT Security Plus portal as well as in portals for Industrial IoT Security and Medical IoT Security.

IoT Security Overview

> *For IoT Security customers with tenants established before 12/15/2022, you can continue using the existing Executive Summary and Inventory dashboards for a limited time. They will eventually be retired and replaced.*

Enterprise IoT Security identifies devices in enterprise networks and creates a dynamic device inventory. It does not include the security features and third-party integrations available in Enterprise IoT Security Plus, Industrial OT Security, and Medical IoT Security. Enterprise IoT Security lets you do the following:

- Automatically classify devices with 12 device attributes
- View, edit, confirm, and reclassify devices
- Add devices with static IP addresses
- See your IP address structure and device distribution
- See sites for firewalls and devices
- Generate device reports

The Devices page, shown below, is the default landing page after login to the Enterprise IoT Security portal. Unlike the other vertical-themed product portals, it does not include dashboards.

For more information, see Enterprise IoT Security Administrator's Guide.

Industrial

Industrial IoT Security is the solution for industrial corporations. It lets you see and secure every device, including specialized operational technology (OT) devices, so you can keep your operations up at all times and achieve NIST and ISA/IEC compliance. You can do the following with Industrial IoT Security:

- Get everything in Enterprise IoT Security Plus
- Detect OT device anomalies
- Use Purdue levels for device modeling and visualization (see Network Visualizations)
- Create customized rules for process integrity (see Create Alert Rules)

As in the portal for Enterprise IoT Security Plus, the Industrial IoT Security portal also includes the Security dashboard.

IoT Security Overview

It's not uncommon for industrial networks to include one or more air-gapped segments. These are areas of the network that do not allow ingress or egress connections between devices in the air-gapped network segment and devices in any other private network segment or with the public network. Through the use of next-generation firewalls configured as security telemetry gateways, you can provide IoT Security services for device in such networks.

Medical

Medical IoT Security is the solution for healthcare providers. It lets you see and secure every device on your network, including specialized medical devices, so you can deliver high-quality patient care and achieve HIPAA compliance. Use Medical IoT Security to do the following:

- Get everything in Enterprise IoT Security Plus

- Detect medical device anomalies

- Assess medical device risk leveraging FDA recalls, PHI identification, and MDS2

- Track medical device utilization

The portal for Medical IoT Security displays two pages that are relevant only to medical IoT and only appear when the Medical IoT Security theme is activated. One is for Food and Drug Administration (FDA) recalls and another is for Manufacturer Disclosure Statement for Medical Device Safety (MDS2) forms. When using the Medical IoT Security theme, the portal also includes two dashboards with data just about medical IoT devices: the Utilization dashboard and, shown below, BioMed dashboard.

For more information, see Medical IoT.

Switch Portal Themes

A tenant can only have one theme at a time for their IoT Security tenant; however, it's possible for tenant owners to switch themes. When users first log in to a tenant and a theme has already been defined by the IoT Security product that was ordered, then that theme is automatically loaded by default. However, if you purchased multiple IoT Security products with different themes (or if you have an IoT Security product purchased prior to December 15, 2022), then IoT Security prompts owners to select a theme when they initially log in to the portal. If an owner doesn't make a selection, IoT Security shows the Enterprise IoT Security Plus theme and continues to prompt owners to select a theme upon each login until one of them makes a selection. Once a selection has been made, all other users in the same tenant will also see the same theme when they access the portal.

To switch vertical themes, log in as a user with owner privileges, select **Administration** > **About** > **License**. The status indicates which theme is currently in use. (You can also see the number of subscribed firewalls and the license start and expiration dates here.) Click **Switch** next to the name of the theme that's currently in use.

License	EULA	Privacy Policy	Tenant Details

Production

IoT Security

Theme
Enterprise Plus Switch

Status	Quantity
In-Use	6

Start Date	Expiration Date
March 26, 2023	March 26, 2024

Third-party Integration Basic

Status	Quantity
In-Use	1

Start Date	Expiration Date
March 26, 2023	March 26, 2024

Select a new theme, and then click **Confirm**.

As an owner, you can switch themes for your tenant as many times as you like.

Create a Trial Enterprise IoT Security Tenant

If you have a production license for Enterprise IoT Security Plus, Industrial IoT Security, or Medical IoT Security and want to see what Enterprise IoT Security is like, you can create a one-time trial tenant and assign up to five of your firewalls to it. The trial is valid for 30 days. During that time, both the production and trial tenants consume log data that firewalls assigned to the trial tenant send to the logging service. When the trial period ends and the trial tenant is automatically deleted, the production IoT Security tenant alone continues consuming the log data from the firewalls.

1. To initiate a trial, log in to the IoT Security portal with a user account that has Owner privileges.

2. Select **Administration** > **About** > **License** and then click **Request** next to Enterprise IoT Security in the Trial section.

3. Choose up to five firewalls that you want to use for the trial and then **Save**.

A message appears explaining that a trial tenant for Enterprise IoT Security is being created, the chosen firewalls will be associated with it, and that the entire process typically takes about ten minutes.

When the process is complete, another message appears stating that the trial tenant has been created and the chosen firewalls have been associated with it. This message also includes the name of the trial tenant.

The trial tenant creation and firewall assignments are also recorded in **Logs & Reports** > **Audit Log**.

4. On **Administration** > **About** > **License**, the button next to Enterprise IoT Security in the Trial section changes from **Request** to **Enter**. To access the trial tenant portal, click **Enter**.

A login prompt appears for the trial tenant in a new browser window.

5. Log in with the same credentials you used to log in to the production IoT Security tenant.

The Enterprise IoT Security portal opens to the Resource Center and is ready for use as a trial tenant. During the 30-day trial, both the IoT Security tenant and the Enterprise IoT Security trial tenant will consume logs from the firewalls assigned to the trial tenant. You can log in to both tenants and compare the functionality of each.

6. To exit the trial tenant and return to the production tenant, navigate to **Administration** > **About** > **License** and then click **Enter** next to IoT Security in the Production section.

The trial tenant browser window remains open while the production tenant opens in a new browser window.

After the trial ends, the trial tenant is automatically deleted while the production tenant continues consuming log data from the firewalls.

*If you have a trial license for IoT Security and want to try out the Enterprise IoT Security product, log in to the IoT Security portal with a user account that has Owner privileges, select **Administration > About > License**, and then click **Manage Trial**. Select **Enterprise** and then **Confirm** your decision. To go back to the IoT Security product, return to the License page, click **Manage Trial** again, select **Enterprise Plus**, and **Confirm**.*

Device-to-Site Mapping

From March 2022, IoT Security provides existing tenants two ways to link devices to sites:

- IP address-based site assignments – IoT Security assigns devices to a site based on device IP address. This method was introduced in March 2022. It is available for existing IoT Security tenants to switch to and is the only option that new tenants (as of March 2022) can use.

- Firewall-based site assignments – IoT Security assigns devices to a site based on the location of the firewall that sends it logs. Until March 2022, this was the only method that IoT Security offered.

For the first approach, you must define one or more Classless Inter-Domain Routing (CIDR) blocks or subnets for each site at **Network** > **Subnets**. For the second approach, you must assign a site to each firewall at **Administration** > **Sites and Firewalls** > **Firewalls**. Site assignment based on firewalls works well for smaller, single-site deployments. However, an issue can arise when there are multiple sites and devices at two sites communicate with each other. When this occurs, the firewalls at both sites observe a session involving the same two devices and report them in logs to IoT Security, which cannot tell where each device is actually located. This issue doesn't occur when IoT Security assigns devices to sites based on IP address, which is the preferred method.

IP Address-based Site Assignment

This method for mapping devices to sites uses IP addresses and is the only site-mapping method available to new IoT Security tenants starting in March 2022.

If you haven't done so already, enter or upload a CSV file of the IP address blocks of your sites in CIDR notation on **Networks** > **Networks and Sites** > **Networks**. (Examples of CIDR notation: 10.55.0.0/16 and 10.197.0.0/16.) Then click **Add** > **Add a Subnet** and enter the network address in CIDR notation and a description, or click **Add** > **Upload Subnets** and upload multiple subnets using the provided template.

> You don't need to use all the subnets that belong to a site for site mapping. Instead, pick the largest subnet (IP address block) for site assignment. For example, one site might have numerous subnets such as 10.55.10.0/24, 10.55.28.0/24, and 10.55.121.0/24, all of which are within a single IP block of 10.55.0.0/16. In this case, use 10.55.0.0/16 for site mapping. IoT Security automatically assigns smaller subnets within the site-mapping IP block to the same site and assigns devices within each subnet to the same site as that of their subnet.

After adding or uploading subnets, assign them to sites on **Networks** > **Networks and Sites and Firewalls** > **Sites**. Either click the **Create Site** (**+**) icon to the upper right of the Sites table or click the three vertical dots icon at the far right of the row for a previously created site and then click **Edit Site**.

Choose the subnets you added or uploaded on **Networks** > **Networks and Sites and Firewalls** > **Networks**.

If you miss a subnet, IoT Security won't be able to link devices in the subnet to a site. When this happens, it assigns devices in this subnet to the Default site to which all the private IP ranges (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16) are assigned for the purpose of catching any unassigned subnets.

Firewall-based Site Assignment

For IoT Security tenants that onboarded before March 2022, IoT Security uses firewall-based site assignments. After you finish onboarding a firewall, it appears on the **Networks** > **Networks and**

Sites > **Networks** page assigned to the Default Site. To reassign it to another site, click the three vertical dots icon in its row on the far right and then click **Change Site**.

Choose one of the sites in the Site Name list and then click **Change**.

IoT Security maps the devices whose traffic metadata appears in the logs from this firewall to this site.

> 📋 *For information about creating sites, see* Sites and Site Groups.

If you don't assign a firewall to a site, IoT Security won't be able to link devices whose traffic appears in logs from this firewall to a site. When this happens, it assigns these devices to the Default Site.

Change Site Assignments from Firewalls to IP Addresses

> 📋 *Only a user with* owner privileges *can change from firewall-based site assignments to site assignments based on IP addresses.*

For IoT Security tenants that map devices to sites based on firewalls, IoT Security provides an option to switch to the IP address-based approach. This is a one-time change. After switching to IP address-based site assignments, you can't switch back to the firewall-based approach.

Select **Networks** > **Networks and Sites** > **Sites** and click the gear icon (⚙) in the upper right of the Sites panel.

Switch from **Firewall-based assignment** to **IP CIDR-based assignment** and then **Save**.

As the note in the dialog box says, it can take up to two days for IoT Security to transition all devices to new sites and that during this time the site assignments for some devices might be incorrect.

Read the confirmation message that appears, recalling that this switch cannot be undone later, and when you're ready, click **Yes** to continue.

After you finish setting up the IP CIDR blocks for site mapping and the new IP address-based site assignment method has had a couple days to establish device-to-site assignments, you can check **Networks** > **Networks and Sites** > **Networks** to verify the configuration and make any adjustments if necessary.

Of particular interest is the Site Mapping column. When a subnet is linked to a site and its entry in the Site Mapping column is **Yes**, this indicates that the subnet has been manually mapped to the site. When a subnet is linked to a site but its entry in the Site Mapping column is **No**, it means that the subnet is a part of a larger IP address block that is mapped to the site and this subnet inherited its site mapping.

*After switching device-to-site mapping from firewalls to IP addresses, IoT Security removes filters for **All connected sites** and **All disconnected sites**. These filters are based on the status of firewall activity at a site, and after the switch, IoT Security no longer links firewalls to sites.*

Sites and Site Groups

📋 *Only users with owner privileges can create and manage sites, organize sites into groups, and assign access to sites and site groups to other users.*

Log in as a user with owner privileges and select **Networks** > **Networks and Sites** > **Sites**. There you can add, view, edit, and delete sites with devices under IoT Security protection.

	Name	Location	↓	Devices	IoT Devices	Risk
☐	San Jose	San Jose, California, United States		5128	770	25
☐	SF-RDD	San Francisco, California, United States		545	133	22
☐	testrevert13	New Johnsonville, Tennessee, United States		754	280	24
☐	Test Site 1	Los Angeles, California, United States		313	45	13
		N/A ⓘ			37	25

Sites (53) ⊕

There are three sections on the Sites page:

- At the top is a title bar with titles for the Networks, Network Segments Configuration, and Sites tabs. There is also a global filter that controls the content displayed on the page by site and time range.

- The Organization section shows the hierarchical structure of sites in your organization.

- The Sites section is a table with useful information about individual sites.

The Default site is where IoT Security initially assigns firewalls. You can later reassign them to user-defined sites.

To add a new site, click **+** above the table. There are different settings based on the device-to-site assignment method that's in use. When assigning devices to sites by IP address, enter a site name, optionally enter a site address and description, choose either the IP prefix of an non-shared IP block or previously defined network segment, and then **Save**. When assigning devices to sites by firewall, enter a site name, optionally enter an address for the site and (if you're organizing sites into groups) choose a site group, and then **Save**.

To edit or delete a site, click the three vertical dots at the far right of a site row and then click one of the actions that appear. When assigning devices to sites based on firewalls, there are two additional options. You can assign one or more firewalls to a site or unassign a site from a group.

Before you can delete a site, you must first remove all firewalls from it or reassign them to different sites.

Organize Sites into Groups

You have the option to organize your sites into groups within a hierarchical structure and then set controls at different levels within the structure to define what administrative users see and do. For example, in the tree structure shown below, you might give a user access to data at an individual site level, or for all sites in a city, or in a state, or within a broader region.

You don't have to organize sites into groups. In fact, by default, the Organization panel is hidden on the **Networks** > **Networks and Sites** > **Sites** page. If you want, you can assign users access on a per-site basis without the use of site groups. However, if you want to see the Organization panel and use this feature, click **Show Organization** and then click **Organize Sites**.

Add Groups to the Tree and Add Sites to Groups

Only a user with owner privileges can add, edit, and delete groups and add sites to them.

There can be five levels in a group hierarchy. The root node forms the top-level group ("Acme" in the examples here) and is the group to which all sites belong by default. By default, it's the name of the tenant account and cannot be removed, but it can be renamed. All other groups below the root are completely owner-defined.

To add a group to the organization, hover your cursor over an existing group, click the **Add group** icon, and then enter a new name. To change its name, click the three dots (**...**) next to the Add group icon and then click **Rename**.

*The global filter has priority over page-level filters. When creating the tree structure, be sure that the global filter at the top of the page is set to **All Sites**. If it's set to anything else, the Organization panel will keep collapsing to show only whatever site or sites were selected in the global filter.*

Add groups and subgroups as needed to reflect the structure of your organization. After adding the groups you need, add sites to them. Select the check box for one or more sites in the Sites panel, click **Assign to Group**, and then choose the one to put them in.

You can search for a group by typing its name in the Search groups field at the top of the Assign to Group drop-down menu.

In addition to adding existing sites to groups, you can also add new sites to groups. When creating a new site (**Networks** > **Networks and Sites** > **Sites** > **+**), the Site Group option lets you assign the site to an existing group, thereby combining site creation and group assignment in a convenient one-step process.

Reassign Sites and Delete Groups

If you later want to reassign a site from one group to another, use the same process for adding it to a group but select the other group from the list.

When you assign a site to a group that also has subgroups, a node labeled Sites appears in the tree under its assigned group at the same level as the subgroups. For example, notice how the group named East Coast has two subgroups—New Jersey and Virginia—and it also has a node called Sites for two sites assigned to the East Coast group.

If you delete a group, IoT Security reassigns all its sites and child groups to its parent group. For instance, look at what happens when the Maryland group is deleted. The site that belonged to Maryland now belongs to East Coast, and its child group Annapolis becomes a child group of East Coast.

To avoid IoT Security automatically reassigning a site when its group is deleted or simply to remove it from a group, click the three vertical dots at the far right of its row in the Sites panel, and either click **Edit Site** to reassign it to another group or click **Unassign** to remove it from its current group and put it into the root node.

Use Groups to Filter and Control Access to Data

After you finish creating the organizational structure and assigning sites to groups, you can use the tree to filter what to display on the Sites page. Click any group name in the tree to display sites belonging to it in the Sites panel on the right. The sites that are displayed either belong directly to the group or are in one of its child groups. (To remove the filter, click the X to the right of its name at the top of the Sites table.)

Not only can you use groups to filter the sites displayed on the **Networks** > **Networks and Sites** > **Sites** page, but you can also filter by group on the **Devices** page.

In the drop-down list for the sites global filter, click a group name (in blue) and then click either **Select All** to see devices at all sites in the selected group or click a specific site to see devices just at that site.

You can also select a group or site when defining the scope of a report at **Reports** > **Files and Settings** > **+** and clicking **Generate a report now** or **Schedule a report for later**.

When logged in as a user with owner privileges, you can use groups to control which sites other users are allowed to access. Do this in the User Role & Access section on the user account settings for a user by clicking **Administration** > **User Accounts >** *username.*

By default, all users have access to all groups and sites. However, after a user with owner privileges gives other users access to one site or group, that's all they can access. If that site or group is ever deleted, these users won't return to having default access to everything. Instead, they won't be able to access anything; that is, until they're given access to something else. On the other hand, users with owner privileges always have access to all groups and sites in their account.

Networks

IoT Security learns about the addressing scheme on your network through several means. You can add subnets and Classless Inter-Domain Routing (CIDR) blocks manually, even specifying if a subnet contains devices that have static IP addresses. IoT Security can discover subnets by observing the exchanges between DHCP clients and servers. IoT Security can learn about subnets through third-party integrations with network switches, using SNMP for network discovery. It can also learn about subnets and CIDR blocks through IP Address Management (IPAM) integrations with BlueCat and Infoblox.

As IoT Security gathers network information, it organizes it hierarchically and displays the subnets and blocks on the **Networks** page (**Networks** > **Networks and Sites** > **Networks**). Blocks are logical partitions of IP address space that serve as an organizational tool for managing addresses. Large "parent" blocks can contain smaller "child" blocks and subnets, where devices are found. Another conceptual grouping is "remainders". These are sets of IP addresses within a block that don't belong to either a subnet or child block.

At the top of the Networks page are two panels that provide a high-level view of your network and how different types of devices are distributed throughout it. The Overview panel is divided in two sections. On the left is the overall number of "networks", which is really a collection of all the network elements (blocks, subnets, and remainders) in your network, and the total number of subnets in your network. On the right of the Overview panel is the total number of network elements at a particular level. If you don't select an entry in the Prefix column of a block in the Networks table, the current level shows the total number of blocks and subnets at the root level. For example, the following Overview panel shows that there are 342 networks (various blocks, subnets, and remainders) of which 332 are subnets. At the current (root) level, there are 24 networks (blocks and subnets) consisting of 18 subnets and 6 blocks (24-18).

If you select one of the blocks by clicking the entry for it in the Prefix column in the Networks table, the overall totals stay the same but the totals in the current level changes to show the subnets, child blocks, and remainders within the selected block.

To see the elements in a child block, select the entry in the Prefix column. To return to the root level, click **Networks (number)** in the breadcrumbs above the Networks table.

The other panel at the top of the Networks page contains a bar chart showing the distribution of device types in each subnet.

The number in parentheses after "Subnet Distribution by Device Type" is the total number of subnets with active devices during the time period set in the filter at the top of the page. The overall number of subnets in the left panel is for all subnets regardless of whether IoT Security detects device activity in them. IoT Security can learn about subnets without detecting device activity by various means:

- User-configured subnets in the IoT Security portal
- User-initiated uploads of subnet configurations in .csv files
- Third-party integration using SNMP for network discovery
- Third-party integrations with IP Address Management (IPAM) solutions from BlueCat or Infoblox
- Detection of IP endpoints but not devices in subnets
- Detection of past device activity in subnets that are inactive during a shorter time period filter set on the Networks page

The total number of subnets in the two panels might be the same if IoT Security detects device activity in every subnet of which it's aware, but most likely the totals are different.

Hover your cursor over one of the bars to see an information pop-up listing the device types in this subnet. For example, the 10.54.0.0/23 subnet shown below has one office device in a subnet that otherwise consists of only network devices. It immediately suggests that the office device might be misplaced on the network.

Click the subnet on the left of the bar chart to see the Subnet Detail panel. By default, device types are shown. To see the device categories and device profiles, click the **Category** and **Profile** tab.

To see details about one type of devices in a subnet, such as the one office device, click the number in the **QTY** (Quantity) column. IoT Security opens the **Assets** > **Devices** page filtered to show the device or devices selected. Then click the name of a particular device to see the **Device Details** page for it.

In the Networks table, IoT Security displays all the blocks and subnets it has been configured with, discovered, and learned through third-party integrations on the Networks page. When a "parent" block has other blocks and subnets nested below it, the number of its "children" is shown parenthetically. To see these blocks click the prefix of the block containing it.

Networks (15)

	Type	Prefix	VLAN	Monitored	Devices	Profiles	Static	Site	
☐	Subnet ⓘ	13.168.100.0/24	39	Yes			-	Default Site	⋮
☐	Subnet ⓘ	12.168.100.0/24	32	Yes			-	Default Site	⋮
☐	Block ⓘ	172.16.0.0/12 (2)		Yes	2	2	-	Default Site	⋮
☐	Subnet ⓘ	192.0.2.0/24	121	Yes			-	Default Site	⋮
☐	Block ⓘ	10.0.0.0/8 (268)		Yes	7320	53	-	Vermont	⋮
☐	Subnet ⓘ	11.168.100.0/24	19	Yes			-	Default Site	⋮
☐	Block ⓘ	192.168.0.0/16 (18)		No	105	13	-	Default Site	⋮
☐	Subnet ⓘ	2.2.2.0/24		Yes			-	Default Site	⋮
☐				Yes			-	Default Site	⋮

For example, if you click the **192.168.0.0/16** block in the screen capture above, IoT Security displays a list of the blocks and subnets within it.

Networks (18)

Networks (15) > 192.168.0.0/16

	Type	Prefix	VLAN	Monitored	Devices	Profiles	Static	Site	
☐	Block ⓘ	192.168.150.0/24 (2)		Yes			-	Default Site	⋮
☐	Subnet ⓘ	192.168.14.0/24	145	Yes			-	Default Site	⋮
☐	Subnet ⓘ	192.168.130.0/24	123	Yes	101	12	-	test-katherine-...	⋮
☐	Subnet ⓘ	192.168.120.0/24		Yes			-	Default Site	⋮
☐	Block ⓘ	192.168.100.0/23 (2)		Yes			-	San Jose	⋮
☐	Block ⓘ	192.168.1.0/24		Yes			-	Default Site	⋮
☐	Block ⓘ	192.168.0.0/24 (3)		Yes	2	1	-	Paris	⋮
☐	Remainder	192.168.0.0/16		No			-	Default Site	⋮

Notice how it contains 18 blocks and subnets and that some of the blocks have parenthetical numbers after them, indicating that there are other smaller blocks and subnets beneath them. You can continue to move downward to lower levels in the hierarchy by clicking the prefix of any block that has a parenthetical number after it. To move upward, click a higher level in the breadcrumb trail at the top of the page.

The Networks page mainly consists of a table presenting a hierarchical view of your network and attributes of the blocks and subnets that constitute it.

Type: There are several types of network grouping categories:

- **Subnet** – A network section with a broadcast domain and gateway.
- **IP Block** – A partition of IP address space that can logically contain other blocks and subnets.

  Shared IP Block – An IP block whose space is partitioned into at least one subnet that's reused in different shared network segments. This results in devices with overlapping IP addresses on the same network. For example, you might use the same subnet for guest traffic in multiple network segments throughout your network. In this case, you would first make a list of firewalls and where they are on the network, perhaps at the same site or at different sites. Next, you'd plan out how to group the firewalls into different network segments and assign each firewall and site to a particular segment. Finally, you'd define the IP block containing the guest subnet as a shared IP block. IoT Security can now automatically detect which network segment an IP address comes from based on its shared IP block and the firewall that sent the log containing the address.

  Non-shared IP Block – An IP block whose space is partitioned into smaller blocks and subnets that are unique in your network. IP addresses in non-shared IP blocks are used in only one network segment in your network.

- **Remainder** – All IP addresses that aren't in more specific IP blocks or subnets contained within a larger, superset block.
- **Network Segment** – A logical grouping of one or more firewalls plus an IP block. When firewalls send traffic logs, IoT Security identifies which network segment a device belongs to by the IP block that its IP address is in + the firewall that sent the log. In this way, the logs uniquely identify each device even when it's using the same IP address as another device in a different network segment.

Name, **VLAN**, and **Description**: When manually adding blocks and subnets in the IoT Security portal, you can include a name and description and, for subnets, a VLAN. IoT Security can also learn these attributes through third-party integrations. BlueCat IPAM integrations can provide a name for a block or subnet. SNMP and Infoblox IPAM integrations can provide the VLAN for a subnet. An Infoblox IPAM integration can provide a description.

> *You can later modify the VLAN and description but not the name.*

Monitored: **Yes** or **No** means a network has devices whose network activity IoT Security is monitoring or not.

Categories and **Profiles**: The number of device categories (such as Personal Computer or IP Phone) and the device profiles (such as PC-Windows and Poly IP Phone) in a subnet.

Source: There are several ways that a block or subnet can be added to IoT Security. This column shows where each block or subnet comes from. The following are the possible sources:

- **Discovered** – IoT Security discovered a subnet by observing network traffic.
- **Config** – A user manually configured an IP block or subnet.
- **Preconfig** – An IP block was preconfigured by IoT Security and cannot be removed. For example, the 10.0.0.0/8 Class A private block.
- **BlueCat IPAM** – IoT Security learned an IP block or subnet through integration with BlueCat IPAM.
- **Infoblox IPAM** – IoT Security learned an IP block or subnet through integration with Infoblox IPAM.
- **Network Discovery SNMP** – IoT Security learned an IP block or subnet by using SNMP to discover network information from switches.

IP Endpoints: IP Endpoints are devices whose IP addresses IoT Security knows but not their MAC addresses. In addition, their behaviors are not stable enough for IoT Security to confidently deduce that their addresses are statically defined. IoT Security displays the number of IP endpoints in a subnet. Click the number to download a .zip file containing a report of IP endpoints in comma-separated-value format.

DHCP and **Gateway**: When IoT Security integrates with switches using SNMP for network discovery and learns the IP addresses of the DHCP server and gateway for a subnet, it displays them in these columns. A BlueCat IPAM integration also provides the gateway for subnets.

Prefix: The network portion of an IP address for a CIDR block or subnet. If you click the entry in the Prefix column for a block, IoT Security displays the blocks, subnets and remainders within it.

If you click the entry for a subnet, IoT Security opens the Subnet Detail panel over the right side of the page. The panel includes various details about the subnet such as a VLAN ID; DHCP server IP address; the number of devices in it per device type, category, and profile; the name and details of the connected switch for the subnet; and firewall security rule details (if there are rules for this subnet learned through Cortex XSOAR integration with Panorama).

> Not every Subnet Detail panel includes the Connected Switch and Firewall Security Rules *sections. For example, IoT Security only learns about connected switches from third-party integrations with Cisco Prime, DNA Center, or Meraki or from integrations using SNMP for network discovery.*

Devices: The number of devices that IoT Security has discovered in a subnet and learned about through a third-party integration.

Static: If a subnet is defined as having static IP addresses, **Yes** appears in this column. Otherwise, a dash (-) appears here, indicating that IoT Security does not have enough data to determine if a subnet has static IP addresses or not.

Firewall Security Rules: (Requires an IoT Security Third-party Integrations Add-on license or an integration through a full-featured Cortex XSOAR server) After you configure IoT Security to communicate with Panorama through Cortex XSOAR, it can fetch any firewall security rules that reference a subnet as the source or destination. The number of rules applied to a subnet appear in the Firewall Security Rules column. When you click the subnet entry in the Prefix column, you can see the rules themselves in the Subnet Detail panel that appears.

167

> 📋 *When **0** appears in the Firewall Security Rules column, it means that a previous rule referencing the subnet has been removed and now no other rules apply to it.*

Low-confidence Devices: This is the number of devices whose identity IoT Security cannot identify confidently. Click the number for a subnet to open the Devices page with a filter applied to show only the low-confidence devices in that subnet; that is, devices with calculated confidence score of 0-69%.

> 📋 *A confidence score indicates the level of confidence IoT Security has in its identification of a device. IoT Security has three confidence levels based on calculated confidence scores: high (90-100%), medium (70-89%), and low (0-69%).*

Site Mapping: Subnets and blocks that are nested within other blocks inherit the site of the topmost block of their set. For example, if there's a 10.1.0.0/16 block at a site named "NYC" and it contains a 10.1.1.0/24 subnet or block, then this subnet or block inherits "NYC" as its site too. **Yes** or **No** indicates whether a subnet or block inherited its site in this manner or not.

Site: The site to which a block or subnet belongs can be defined manually (see Device-to-Site Mapping) or learned through an integration with Infoblox IPAM.

Devices Discovered via Integration: The number of devices learned through integration with a third-party system.

Removable: Indicates if you can remove a subnet or block. Preconfigured blocks, like 10.0.0.0/8, and those currently being used for site mapping cannot be removed.

Clicking the subnet entry in the Prefix column opens the Subnet Detail panel where you can see more information about it.

Below the Networks table is a map showing the number of devices that made connections to external destinations; that is, to destinations outside the local network. The color of the countries to which devices connected indicate how many devices made connections to them and if any of the destinations were malicious.

- Gray indicates there were no devices with connections to a country.

- Green indicates devices connected to safe destinations to a country. The darker the green the more devices connected to destinations there. (See the legend for numbers.) Hover your cursor over a country to see an information pop-up with the number of devices that connected to destinations there during the time period filter at the top of the page. If you click a country, the information pop-up remains in view until you click the country again to close it. This allows you to click two or more countries and easily compare the number of connections to each one. Click the number in the pop-up to open the **Assets** > **Devices** page with a filter set to show devices that connected to this country.

- Red indicates at least one device connected to a malicious destination there.

Click-drag your cursor to move the map. Use the scroll wheel on your mouse or the **+/-** tools in the lower right corner of the map to zoom in and out.

Network Visualizations

IoT Security monitors and analyzes network traffic to provide a data-rich, dynamically updated inventory of the devices on your network. Through its extensive monitoring and analysis of network activity, IoT Security can also expose communication patterns among devices of interest by visualizing them in user-defined network visualization maps. By focusing on different groups of devices and different facets of the network, trends, patterns, and aberrations can emerge in the visualization of device communications and in the relationship between devices and the network segments on which they operate or—for Operational Technology (OT) devices—between the OT devices and the Purdue levels to which they're assigned.

IoT Security provides various methods to group devices for visualization: by device attributes such as subnet, VLAN, vendor, category and profile, and by Purdue level. It also provides the option to create visualization maps with either one or two layers. That is, you first organize devices into groups based on a particular attribute, such as the VLAN they're in. This results in a set of device groups organized by VLAN, allowing you to see the distribution of devices across the different VLANs in your network. So far, this is a one-layer map. However, if you want, you can also organize the devices within each VLAN by another attribute such as device profile. Then, by drilling down into different VLANs, you can enter a second layer of the map and see the distribution of devices within each VLAN by profile.

Create a Visualization Map

STEP 1 | Select **Networks** > **Network Visualizations**.

Before you create your first network visualization map, the Network Visualizations page displays a map of the world. Any existing sites whose location has been defined appear in those locations on the map. Sites without a defined location appear in an Unknown Sites list in the lower left corner of the map. To define a location for a site, select **Networks** > **Networks and Sites** > **Sites**, click **N/A** in the Location column (or click the three vertical dots icon at the right end of a row and then **Edit Site**), enter a city name in Site Address, and then **Save**.

After you add and save a visualization map, it appears on this page so that you can return to view the map later by clicking **View Map**.

STEP 2 | Create a network visualization map.

> *There isn't a maximum number of visualization maps you can create, but there is a maximum of 500 nodes (subnets, profiles, devices, and so on) that a map can display. If the number of nodes exceeds 500, IoT Security hides the map and presents the information in table format only.*

1. Select **Networks** > **Network Visualizations** > **+ Create Map**, select one or more sites, and **Add to Scope**.
2. After you set the site scope, click **Next**.
3. Click **Device Grouping** to configure the method for grouping devices on the map based on your needs. You can change this later.

   The device grouping you select determines the type of map you create. First, group devices by one of the following attributes: **Category**, **Profile**, **Vendor**, **Subnet**, **VLAN ID**, or **Purdue**

Level. Then, optionally, depending on the attribute you used, organize them within each first-layer group by another type of attribute in a second layer:

First set of groups	Second set of groups (optional)
Category	–
Profile	–
Vendor	Risk Level
Subnet	Category or Profile
VLAN ID	Category or Profile
Purdue Level*	Category or Profile

* Before creating a device visualization map based on Purdue levels, you must first indicate the Purdue level to which various devices belong. You can do this by defining custom attribute rules that apply Purdue levels to devices automatically. This involves the following process:

1. Make a list of device attributes, such as profiles, for all OT devices at Purdue levels 0-3 on your network. Optionally, make a list of subnets for all other IT and IoT devices that are separate from OT and are in levels 4-5.

2. Create six filters on the Devices page, each filter listing a set of profiles or subnets for the devices at a particular Purdue level. For more information about filters, see IoT Security Devices Page.

3. Use the six pre-defined values for Purdue Levels 0-5 to create custom attribute rules to assign Purdue Levels to devices based on the filters you created (a default filter is used to assign a Purdue Level to devices based on Category). IoT Security assigns any device that doesn't match one of these rules to the "Unknown" level.

For example, if you set the first set of groups as **Subnet** and the second set of groups as **Category**, you'll create a map that first shows devices organized into various subnets. Then if you navigate to the second layer of the map by clicking one of the subnets, you'll see devices grouped by device category within it.

4. Continue to refine the map scope by entering more parameters to define the scope of the visualization map and then click **Update**.

IoT Security displays a visualization based on the scope you define. The scope must include a time range during which devices were active on the network (the past day, week, or month). The scope also typically contains at least one site; however, it's possible to make a map without specifying any specific site, in which case the map includes all sites. In addition to a time range and sites, you can optionally add numerous filters to narrow the map scope

further. Doing so lets you more easily find the types of devices you're looking for and also reduces the number of nodes that the map displays.

5. Review the visualization and, if necessary, continue adjusting the scope and device grouping until the map shows the data you want to see.

6. When you're satisfied with the content of the visualization map, click **Build Map**, and then enter the following:

   Name: Enter a name for the visualization map

   Description: Optionally enter a description of the visualization map for later reference.

   Scope: Review the filters that define the parameters of the map. Because a map can contain up to 500 nodes, define a scope that stays within this range. You can narrow the scope by filtering devices by type as well as by various device, alert, and vulnerability attributes. This filtering behaves much like the query builder.

   Device Grouping: Review the device grouping of the map. You can edit the grouping method here and while viewing a saved map.

7. Click **Confirm**.

   The map immediately becomes available to view on the Networks Visualization page.

STEP 3 | Purdue Levels Manually reassign devices individually if necessary.

After setting up the filters and letting the rules automatically assign devices to Purdue levels, periodically do spot checks of important devices to make sure they are assigned to the correct Purdue level on the visualization map. If any device isn't properly assigned, note its IP and MAC address to look it up by device ID in the IoT Security inventory. Then manually reassign it to the right level on its Device Details page.

View Data in a Visualization Map

Options for navigating a visualization map and viewing its data apply to both types of visualization methods: device attributes and Purdue levels.

Nodes (Groups and Devices)

The nodes on each level of a map are depicted as circles and the dashed lines between nodes represent network connections. A node can be a group of objects such as subnets, VLAN-IDs, device categories, device profiles, vendors, or risk levels, or a node can be a single device within one of these groups. The number that's shown within the circle of a group indicates how many devices are in it. Some groups have colored segments around the edge of their circle. These indicate the proportion of devices within it that have a particular risk severity. Critical is red, high is orange, and medium is yellow. A low risk level is the remaining gray that circumscribes the circle. (In other parts of the IoT Security portal, blue represents a low severity level; however, because blue is used to highlight nodes in visualization maps, it's not used here to indicate a low risk level.) The size of the circle for a group indicates the proportion of devices in it in relation to other groups on the map.

Highlight

The highlight tool, located at the top of a visualization map, helps you find devices with certain characteristics. To use it, enter one or more filters using query language and then click **Highlight**. IoT Security highlights (with a blue ring or partial ring) all groups and devices that match the

filters. The length of the ring denotes the proportion of items in a group matching the highlight definition. You can then drill down to the highlighted devices that match the filters.

Interactions

- **Hover**: Hover your cursor over a group of devices to see a pop-up panel with information about the groups and devices within it. You can hover your cursor over a group that contains other groups to see information about devices within all the groups or you can hover your cursor over one of the inner groups to see information just about that one. Hovering over a device displays a pop-up panel with information about that device.

- **Click once**: Clicking a group or device once puts it in focus and displays an information panel about it on the right side of the map. Clicking the **External Link** icon at the top of the device information panel opens the Device Details page for the device, where you can see relevant information.

- **Click twice**: Clicking a group twice (double-clicking or clicking on a focused group or device) drills into it to see its contents and the network connections of its contents to other groups. Clicking a device twice shows its network connections to other devices.

- **Reposition nodes**: You can also drag groups and devices to reposition them on the map. This feature only works on the main map display. When you double-click a particular group, the new group in focus always appears centered on the map.

- **Use the table and breadcrumbs**: Use links in the table to navigate through map layers by clicking links in table columns to drill down deeper into the map and clicking links in the breadcrumbs above the table to move up to higher layers.

- **Use the Back button**: In addition to clicking the breadcrumbs above the table to move back to a higher map layer, you can also click the **Back** button between the IoT Security logo and map name at the top of the page. When you're already at the top map layer, clicking the **Back** button exits the current map and returns to the visualizations landing page.

Map Name and Totals

A summary of various totals appears below the map name in the upper left of the page.

For example, the first number might be the number of subnets, the second the number of categories, and the third the number of devices on a map. If the scope contains more than 500 nodes, consider reducing the scope so the map can display them.

After creating a map and engaging with it, you might make some changes and tweaks and decide you want to save the edited map. To do that, click the **Edit Map** icon next to the map name. IoT Security displays the Update Network Visualization Map panel where you can change the map name, description, the visualization method, and scope and then **Confirm** your changes. Another option in the Update Network Visualization Map panel is Map Builder. Click **Map Builder** to view the map and make edits to the visualization method (Device Grouping) and scope. By clicking **Update** after adding or removing filters to the scope, you can see how your changes affect the contents of the map. When done, click **Update Map**, which returns you to Update Network Visualization Map. Review your modified settings and, if satisfied, **Confirm** the changes. If you aren't yet satisfied, click **Map Builder** again to return to the map and continue making adjustments as necessary.

Legend

On the left of a visualization map are zoom in (+) and zoom out (-) icons and an information icon that opens a legend of what the colors and icons mean. Click to expand it.

Basic

- When viewing an individual device, its risk level is indicated by the color at 1:00 on the circle.

- When viewing a device group, the risk level or levels of the devices within it are indicated by red, orange, and yellow around the edge of the circle. The amount of each color is the proportion of devices at that risk level in relation to the overall number of devices in the group.

- When using the highlight tool to find devices with a particular attribute, a blue ring—or segment of a ring—appears within the edge of a group, its length indicating the proportion of devices with the highlighted attribute in the group. The longer the blue segment is, the more highlighted devices there are proportionally.

Risk Level

- The color for each risk level is identified.

Icons

- A green globe indicates that one or more devices in a group have connections to normal Internet sites.

- A red globe indicates that one or more devices have connections to malicious Internet sites.

- A three-pronged yellow icon indicates that there are one or more connections to off-map devices; that is, to devices that are on the local network but aren't within the scope defined for this visualization map.

- A laptop icon indicates that one or more devices have connections to IP endpoints on the local network. An IP endpoint is the source or destination of a network connection for which IoT Security has learned an IP address but not a MAC address.

Map Management

In the Map Management section, you can control what types of devices and connections to display on the map. By selecting and clearing their check boxes, you can toggle the icons on and off on the map.

- **Inner Connection**: Select or clear the check box to show or hide inner connections, which are connections within the same device grouping. Because connections between groups are typically of more interest, this is toggled off by default. To see inner connections (connections between devices in the same group), toggle on **Inner connections**.

- Device visualization maps sometimes include **IP Endpoints**, **Off-map Devices**, and **Internet Connections** (**Normal** and **Malicious**) whenever it's necessary to show connections between devices defined within the scope of a visualization map and destinations outside that scope. Off-map devices (dark yellow shaded circles) and IP endpoints (gray shaded circles) are located in the local, private network, and Internet addresses are sites in the external public network (green shaded circles for normal sites and red shaded for malicious sites). An IP endpoint is a device for which IoT Security knows an IP address. An out-of-scope device is one for which IoT Security knows both an IP address and a MAC address but is outside the map scope. As with other device groups, you can also drill into groups of out-of-scope devices and endpoints and Internet addresses. Click the group once to put it in focus and open an information panel. Click it twice to zoom into it and view its contents.

Reports

The Reports section (**Logs & Reports** > **Reports**) is where you can view and download reports of various types:

- **Summary** provides a summary of device inventory, risk assessment, and alerts.
- **Discovery** provides a view of devices that IoT Security discovered on your network, their distribution in different subnets/VLANs, and devices with high-risk scores and pending alerts.
- **New Device** reports all the new devices detected on your network since the last report. IoT Security can generate reports on a daily, weekly, or monthly basis.
- **Risk** summarizes all risks associated with IoT devices. It reports an overall risk score, at-risk devices, alerts, vulnerabilities, risk-related trends, and the status of risk remediation efforts.
- **Inventory Gap** (when IoT Security is integrated with a CMMS) shows devices discovered by IoT Security, those in your CMMS (computerized maintenance management system) inventory, and where the two sets of devices do and do not overlap.
- **Utilization** provides data visualizations about medical IoT device operations and usage.
- **Filtered Inventory** prepares a device inventory report using a previously defined filter of your choice from the Devices page.

There are two ways to generate a report—immediately and scheduled. The scheduled reports can be generated either once or on a recurring basis.

- Discovery reports can be generated only immediately
- All the other report types must be scheduled
- Risk reports and Utilization reports can be generated both ways

Configure Reports

You can configure your reports to either generate on demand or schedule them to be generated at a later date.

Generate a Report Now

You can only generate Discovery, Utilization, and Risk reports immediately; all others must be scheduled. To generate a report immediately, click the **+** icon on the top right of the reports page and select **Generate a report now**. Provide or choose the required details and click **Generate**.

- **Report type**: Choose the report type from the drop-down list.

- **Report Name**: Enter a name for the report.

- **Sites**: You can choose all sites, an individual site, or–if you arranged your sites in hierarchical groups–a group of sites.

- **Alert Severity** and **Risk Level**: For Discovery reports, alert severity is for security alerts and risk level is for vulnerabilities. You can choose one, two, or three severity and risk levels. IoT Security will filter the devices it includes in the Discovery report based on your choices. If you leave these empty, then alert severities and risk levels aren't used to filter which devices to include in the report.

- **Device type**: Optional for Risk reports; choose all device types or one or more individual types (automotive, industrial, medical, and so on). For Discovery reports, choose from either discovered or monitored. "Discovered" are devices that IoT Security knows are on the internal network but it's not monitoring and protecting. "Monitored" devices are also on the internal network and IoT Security is monitoring their network activity to do device profiling, behavioral analysis, and risk monitoring.

- **Device Category**: For Utilization reports, the device category field for an Inventory report is limited to infusion systems and image scanners (X-ray machines, UltraSound machines, MRI machines, CT scanners, and PET scanners).

- **Subscribe**: Choose the email address to which you want to send the reports.

- **Select a time range**: Choose from the available time ranges or create a custom time range for which you want to generate reports.

The report gets generated in a few moments and is available on the reports page.

Schedule a Report for Later

Scheduled reports can be generated once or on a recurring basis. All reports can be scheduled except Discovery and Utilization reports.To schedule a report for a later date, click the **+** sign on the top right of the reports page and select **Schedule a report for later**. Provide or select the required details and click **Schedule**.

In addition to the fields described in the the previous section, fields specific to scheduled reports are described below.

- **Scope**: For Summary reports, this decides what the report will include and can either be set as Site or Device Type (automotive, industrial, medical, and so on). Choosing **All** does not filter for sites or device types.

- **Saved Filters/Queries** (optional): For Filtered Inventory reports, choose from the saved filters from the drop-down.

- **Set a recurring schedule**: You can schedule reports to be run on a daily, weekly, monthly, or custom time basis.

  - Summary reports: Weekly, Monthly on Day 1

  - Risk reports: Daily, Weekly, Monthly on Day 1, Custom schedule (set to any day and time of the week)

  - New Device and Filtered Inventory reports: Daily, Weekly, Monthly on Day 1

  The first time a scheduled report is generated, it will include data that IoT Security gathered over the time period set in the report configuration. For example, if you create a monthly recurring New Device report on 27 October to start on the first of the month, IoT Security will generate its first report on November 1 with 31 days worth of new devices starting from October 1. The same holds true for daily and weekly scheduled reports. After a scheduled report is initially generated, IoT Security continues to produce reports at the specified interval with data in each new report gathered during the time since the last one.

The report is now scheduled to be generated at the time you selected. Similar to immediate reports, this report is also available on the Reports page.

View Reports

You can view your reports in either the card view or list view. The view setting is in the upper right of the Reports page between the **+** icon and Search field.

The **card view** on the Reports page displays similar reports grouped inside a card. Cards are grouped by three parameters: Report type, Scope, and Schedule. For example:

- All Discovery reports with a scope set as "All Sites, Discovered" and a schedule set as 1 day are grouped together under a single card.

- All Summary reports with a scope set as "All Sites, Monitored" and a schedule set as "Weekly on Friday at 5 PM" are grouped together under another single card.

The **list view** on the Reports page displays all reports in a list format. You can sort the reports according to Report name, Configuration, Scope, and so on. You can delete a report only in the list view.

	Report Name	Configuration Type	Scope	Time / Schedule	Create Time ↓	Next Report On
☐	Sam's Filtered Inve...	Filtered Inventory Report	No saved filter applied	Daily at 2AM	01:00 AM, Oct. 25, 2022	02:00 AM, Oct. 26, 2022
☐	John's Filtered Inv...	Filtered Inventory Report	No saved filter applied	Daily at 3PM	01:00 AM, Oct. 25, 2022	03:00 PM, Oct. 26, 2022
☐	Sam's New Device...	New Device Report		Daily at 2AM	11:00 PM, Oct. 24, 2022	02:00 AM, Oct. 26, 2022
☐	Risk Report for Oc...	Risk Report	All Sites	1 Day	05:00 PM, Oct. 19, 2022	
☐	Discovery Report f...	Discovery Report	All Sites, Discovered	1 Day	05:00 PM, Oct. 19, 2022	
☐	Discovery Report f...	Discovery Report	All Sites, Discovered	1 Day	05:00 PM, Oct. 19, 2022	
☐	Risk Report for Ka...	Risk Report	All Sites	12 Months	04:00 PM, Oct. 19, 2022	
☐	Summary Report f...	Summary Report	All Sites, Monitored	Weekly on Friday at 5PM	04:00 PM, Oct. 19, 2022	05:00 PM, Oct. 30, 2022
☐	John's Discovery R...	Discovery Report	Paris, Discovered	1 Day	08:00 AM, Oct. 18, 2022	
☐	Beryl's Discovery ...	Discovery Report	Paris, Discovered	1 Day	07:00 AM, Oct. 18, 2022	
☐	Summary Report f...	Summary Report	All Sites, Monitored	Weekly on Monday at 12...	01:00 PM, Oct. 17, 2022	12:00 AM, Oct. 31, 2022
☐	Summary Report f...	Summary Report	All Sites, Monitored	Weekly on Friday at 11AM	10:00 AM, Oct. 14, 2022	11:00 AM, Oct. 31, 2022
☐	Risk Report for Ka...	Risk Report	All Sites	Daily at 11AM	10:00 AM, Oct. 14, 2022	11:00 AM, Oct. 26, 2022
☐	Risk Report for Oc...	Risk Report	All Sites	1 Day	11:00 AM, Oct. 12, 2022	
☐	Discovery Report f...	Discovery Report	All Sites, Discovered	1 Day	11:00 AM, Oct. 12, 2022	
☐	New Device Repor...	New Device Report		Daily at 1AM	12:00 PM, Oct. 06, 2022	01:00 AM, Oct. 26, 2022

Similar reports by multiple users are grouped together under a single card in the card view. Similar reports by multiple users can be viewed by sorting them in the list view. A report generated by a user can have multiple versions.

View Report History (*n*) and View Report (*m*) of (*n*)

View report history (*n*) is different from **View report (*m*) of (*n*)** as can be seen in the following sample use case:

John and Beryl generate the same report: Discovery report with a scope set as "Paris, Discovered" and a schedule set as 1 day. The reports of both are placed inside a single card. Each time John clicks **Generate now** to generate his report, the **View report history (n)** increases incrementally. For example, when you see View report history (3) for John's report, it means that John has generated his report three times. So, **View report history** indicates the different versions of the same report by the same user. The **View report 1 of 2** at the bottom indicates that this report (which is also the latest) was generated by John and that there's a second, similar type of Discovery report generated by some other user.

To view the second user's report, click **>** at the bottom of the card, the number now changes to **View report 2 of 2** (see image below). Report 2 was generated by Beryl. Beryl generated her report four times, so we now see **View report history (4)**.

View Reports in a Browser

Because Summary, Connectivity, Discovery, Risk, Inventory, and Utilization reports are generated as HTML, they can be viewed in a browser. To view your reports, click on **View Report History (n)** or **View Report (m) of (n)** on the report. You can also print and download them as PDF files.

Download Reports to View

Because New Device reports and Filtered Inventory reports are generated as .csv files, they can only be viewed in a spreadsheet reader or editor after downloading them. To download your reports, click on **View Report History (n)** or **View Report (m) of (n)** on the report.

	hostname	ip address	internet access	profile type	profile vertical	category	profile	last activity	long description	vlan	site name	risk score	risk level	in use	subnet	number of critical al	number of warning a	number of caution al	number of info alerts	first seen date
1	hostname	ip address	internet access	profile type	profile vertical	category	profile	last activity	long description	vlan	site name	risk score	risk level	in use	subnet	number of critical al	number of warning a	number of caution al	number of info alerts	first seen date
2	10.5.34.190	10.5.34.190	2022-10-24T09:12:0	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T09:11:27.000Z			8923	Default Site		20	Low		10.5.32.0/20	3	0	0	0	2020-12-26T
3	10.6.129.61	10.6.129.61	2022-10-24T09:13:4	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T09:13:49.000Z			8923	Default Site		20	Low		10.6.128.0/20	3	0	0	0	2021-01-08T
4	10.5.104.235	10.5.104.235	2022-10-24T09:17:4	IoT	Network Devices		Palo Alto Networks D	2022-10-24T09:16:59.000Z			8923	Default Site		20	Low		10.5.96.0/20	3	0	0	0	2020-12-29T
5	10.46.168.198	10.46.168.198	2022-10-24T09:31:1	IoT	Network Devices	Network Security Equi	Palo Alto Networks D	2022-10-24T09:18:07.000Z			76	Default Site		20	Low		10.46.168.0/22	4	0	0	0	2021-02-03T
6	BLRIOPEXL637	10.193.204.208	2021-12-23T12:04:0	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T09:23:36.000Z			8923	Default Site		20	Low		10.192.0.0/14	3	0	0	0	2021-01-16T
7	10.6.129.30	10.6.129.30	2022-10-24T09:26:3	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T09:26:34.000Z			8923	Default Site		20	Low		10.6.128.0/20	3	0	0	0	2021-01-02T
8	10.2.127.154	10.2.127.154	2022-10-24T09:43:5	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T09:43:18.000Z			8923	Default Site		20	Low		10.2.64.0/18	3	0	0	0	2021-01-06T
9	10.5.76.10	10.5.76.10	2022-10-24T09:48:4	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T09:48:29.000Z			8923	Default Site		20	Low		10.5.64.0/20	3	0	0	0	2020-12-19T
10	10.5.136.41	10.5.136.41	2022-10-24T09:51:5	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T09:51:55.000Z			8923	Default Site		20	Low		10.5.128.0/20	3	0	0	0	2020-12-25T
11	10.5.76.11	10.5.76.11	2022-10-24T10:16:4	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T10:16:18.000Z			8923	Default Site		20	Low		10.5.64.0/20	3	0	0	0	2020-12-29T
12	10.2.127.31	10.2.127.31	2022-10-24T07:59:4	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T10:16:45.000Z			8923	Default Site		20	Low		10.2.64.0/18	1	0	0	0	2021-01-08T
13	LAP-01-5175	10.193.206.141	No	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T10:22:11.000Z			8923	Default Site		20	Low		10.192.0.0/14	2	0	0	0	2022-06-21T
14	10.3.28.7	10.3.28.7	2022-10-24T10:45:0	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T10:41:52.000Z			8923	Default Site		20	Low		10.3.0.0/18	3	0	0	0	2020-12-29T
15	potato1	10.2.127.60	2022-10-24T10:43:1	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T10:43:13.000Z			8923	Default Site		20	Low		10.2.64.0/18	3	0	0	0	2020-12-29T
16	10.2.132.153	10.2.132.153	2022-10-22T03:53:3	IoT	Network Devices	Network Security Equi	Palo Alto Networks D	2022-10-24T10:54:45.000Z			8923	Default Site		20	Low		10.2.128.0/18	1	1	0	0	2021-02-26T
17	10.2.132.220	10.2.132.220	2022-10-11T03:58:2	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T10:54:45.000Z			8923	Default Site		20	Low		10.2.128.0/18	1	0	0	0	2021-01-06T
18	10.2.132.241	10.2.132.241	2022-09-22T20:35:1	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T10:54:45.000Z			8923	Default Site		20	Low		10.2.128.0/18	1	0	0	0	2021-01-08T
19	10.2.132.30	10.2.132.30	2022-10-22T03:54:4	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T10:54:45.000Z			8923	Default Site		20	Low		10.2.128.0/18	3	0	0	0	2021-
20	10.2.132.100	10.2.132.100	2022-10-21T12:32:0	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T10:54:46.000Z			8923	Default Site		20	Low		10.2.128.0/18	2	0	0	0	2020-
21	10.46.35.93	10.46.35.93	2022-10-24T11:03:4	IoT	Network Devices	Network Security Equi	Palo Alto Networks D	2022-10-24T11:03:16.000Z			8923	Default Site		20	Low		10.44.0.0/14	2	0	0	0	2021-02-25T
22	10.2.238.196	10.2.238.196	2021-12-23T06:09:1	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T11:04:04.000Z			8923	Default Site		20	Low		10.2.192.0/18	1	0	0	0	2021-01-09T
23	10.2.238.201	10.2.238.201	2022-10-24T08:04:3	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T11:04:04.000Z			8923	Default Site		20	Low		10.2.192.0/18	3	0	0	0	2020-12-26T
24	cranberry2	10.5.34.171	2022-10-24T11:21:0	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T11:21:02.000Z			8923	Default Site		20	Low		10.5.32.0/20	3	0	0	0	2021-01-07T
25	10.3.220.204	10.3.220.204	2022-09-24T13:31:3	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T11:25:48.000Z			8923	Default Site		30	Low		10.3.192.0/18	1	0	0	0	2021-02-03T
26	10.3.220.195	10.3.220.195	2021-03-17T07:59:3	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T11:25:49.000Z			8923	Default Site		20	Low		10.3.192.0/18	1	0	0	0	2020-12-29T
27	10.3.220.250	10.3.220.250	2021-03-24T18:05:2	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T11:25:49.000Z			8923	Default Site		20	Low		10.3.192.0/18	1	0	0	0	2020-12-24T
28	SJCWN014F5F0	10.47.125.94	No	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T11:28:16.000Z			8923	Default Site		30	Low		10.44.0.0/14	1	0	0	0	2022-07-20T
29	10.4.0.100	10.4.0.100	2021-10-12T18:49:3	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T11:29:00.000Z			8923	Default Site		30	Low		10.4.0.0/20	1	0	0	0	2021-02-03T
30	10.4.0.176	10.4.0.176	2022-10-13T04:01:1	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T11:29:08.000Z			8923	Default Site		20	Low		10.4.0.0/20	1	0	0	0	2022-06-21T
31	10.5.35.140	10.5.35.140	2022-10-24T11:29:1	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T11:29:19.000Z			8923	Default Site		20	Low		10.5.32.0/20	3	0	0	0	2021-01-01T
32	LAP-01-5735	10.130.177.135	2022-10-07T11:03:5	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T11:35:56.000Z			8923	Default Site		30	Low		10.128.0.0/14	4	0	0	0	2022-06-22T
33	10.5.12.16	10.5.12.16	2021-12-19T05:20:3	IoT	Network Devices		Palo Alto Networks D	2022-10-24T11:53:49.000Z			8923	Default Site		10	Low		10.5.0.0/20	1	0	0	0	2021-02-03T
34	10.5.17.228	10.5.17.228	2021-12-21T09:51:3	IoT	Network Devices		Palo Alto Networks D	2022-10-24T11:54:05.000Z			8923	Default Site		20	Low		10.5.16.0/20	2	0	0	0	2021-02-03T
35	10.3.255.196	10.3.255.196	2022-10-24T11:54:0	Non_IoT	Traditional IT	Personal Computer	PC-Windows	2022-10-24T11:54:06.000Z			8923	Default Site		30	Low		10.3.192.0/18	3	0	0	0	2021-02-03T
36	10.6.129.125	10.6.129.125	2022-10-24T12:05:0	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T12:04:55.000Z			8923	Default Site		20	Low		10.6.128.0/20	4	0	0	0	2021-01-03T
37	10.6.142.110	10.6.142.110	2022-10-24T12:11:4	Non_IoT	Traditional IT	Personal Computer	PC-Linux	2022-10-24T12:11:42.000Z			8923	Default Site		20	Low		10.6.128.0/20	3	0	0	0	2021-01-08T

Edit, Copy, and Disable Reports

Click the Action menu icon (...) on a report to edit, copy, and disable it.

Edit Reports

Edit a report to adjust settings. For example, you might want to increase or decrease the frequency of a scheduled report or add or remove subscribed email addresses.

You can also generate a scheduled report on demand instead of waiting for the scheduled time. When you click **Edit** > **Generate Now** for scheduled reports, IoT Security generates a report based on the period of time-daily, weekly, or monthly-that's set in the configuration, going back a day, a week, or a month from the moment you clicked **Generate Now** and generates a report for till that moment. For example, if you have a New Device report scheduled to be generated every Monday at 3:00 AM and you click **Generate Now** on Wednesday at 10:00 AM, you'll then get a report for a full week (7 days) from the previous Wednesday at 10:00 AM up to the moment you clicked **Generate Now**.

Copy Reports

     184

Copy a report to keep your original report but create another one based on it with some changes; for example, you might want to regularly generate two New Device reports-one report each day for a daily sanity check and another each month for a monthly team report.

Disable Reports

Disable a report to suspend its scheduled generation, perhaps during a scheduled network maintenance. However, you can still view existing reports while it's disabled. You can enable the report again later when you want to resume its use.

IoT Security Integration Status with Firewalls

The Firewalls page (**Administration** > **Firewalls** > **Firewalls**) provides an overview of firewall connectivity and activity, the status of logs that firewalls send and the requests they make for policy rule recommendations and IP address-to-device mappings, and individual firewall details.

Firewalls	Certificate Generation	On-demand PCAP					Search devices, vulnerabilities, ip endpoints and exte...

Filter Query 1 Day ⌄ ☰ Add Filters ⌄

Overview

30	4.1K	51 ☁	4.1K ⊘	6.5K
Sites	Firewalls	Receiving logs	Not receiving logs	System Alerts

● Firewall Request Status ⓘ

Firewall Log Type Status

Type	Status	Latest Log	Log Events	Log Bytes	Ave. Latency	Max Latency
Total	● Live Data	Feb 7, 2024, 18:55	3.1B	4.51TB	1 min	26 min
EAL	● Live Data	Feb 7, 2024, 18:55	172.4M	290.39GB	1 min	26 min
DHCP	● Live Data	Feb 7, 2024, 18:55	746.8K	1.38GB	1 min	4 min
DHCP ACK	● Live Data	Feb 7, 2024, 18:55	96.6K	206.22MB	1 min	7 min
Traffic	● Live Data	Feb 7, 2024, 18:55	2.9B	4.15TB	1 min	12 min
Threat	● Live Data	Feb 7, 2024, 18:55	22.2M	44.11GB	1 min	1 min
ARP	● Live Data	Feb 7, 2024, 18:55	23.6M	31.39GB	1 min	3 min
GTP logs ⓘ	● No Live Data	-	0	0B	0	0

Firewalls (4,131)

	Status ↓	Serial N...	Hostname	IoT Devices	EAL	DHCP	Traffic	ARP	Software ...	Applicati...	First Con...	Latest Ac...	License E...	Instance Type
☐	☁	01630100...	sfo-cfw01	31	2.7M	—	3.4M	389.6K	11.0.2	8806-8548	15:47, Nove...	18:54, Febr...	08:47, Nove...	NGFW
☐	☁	01250100...	tlv-cfw02	768	10.7M	60.9K	86.3M	1.3M	11.0.2	8806-8548	23:18, Nove...	18:54, Febr...	16:17, Nove...	NGFW
☐	☁	01640100...	syd-cfw01	3490	2.9M	20.1K	3.7M	197.9K	11.0.2	8806-8548	13:51, Nove...	18:54, Febr...	06:50, Nove...	NGFW
☐		01250100	tlv-cfw01	4067	720.3K	—	—	719.1K	11.0.2	8806-8548	22:56, Nove...	18:49, Febr...	15:56, Nove...	NGFW
						8.2K	115.5M						16:17, Nove...	NGFW

The overview at the top of the page shows how many sites are under IoT Security management, how many firewalls are subscribed to IoT Security, how many firewalls are active and from which IoT Security is receiving logs, how many firewalls IoT Security isn't receiving logs from, and

how many system alerts there are. Click the system alerts number to view the list of alerts at **Administration** > **System Events**.

IoT Security considers a firewall to be active if it received a log from it within the past 30 minutes, and if it doesn't receive a log during this time, it automatically generates an alert. The Firewalls page also shows how many log events firewalls sent to IoT Security over the past 7 days, 24 hours, or hour (depending on the time filter you set), the time the last log was received, and the connectivity status of the firewalls.

📋 *IoT Security coordinates data received from all the firewalls at the same site. Not every firewall needs to send logs to IoT Security as long as other firewalls do and their logs capture network traffic data from all the IoT devices that you want IoT Security to monitor.*

Hover your cursor over the **Firewall Request Status** icon to see if IoT Security is receiving requests from firewalls for policy recommendations and IP address-to-device mappings.

Firewall Request Status			
Status	Request Type	Latest Request	# of Request
● Pulling/not Pulling ⓘ	Policy Recommendations	Nov 10, 2023, 09:21	0
● Live	IP address-to-device Mappings	Feb 7, 2024, 20:38	260.6K

When IoT Security has received requests for one of these within the past 30 minutes, the status icon is green. Otherwise it's red.

For firewalls in an active/passive HA pair that have sent log events to IoT Security within the past 30 minutes, the status of the active firewall is shown as **Receiving logs**. The status of the passive firewall is usually shown as **Not receiving logs** except for a period of 30-60 minutes after it reboots. During this time, its status changes to **Receiving logs** before returning to **Not receiving logs** again. This is true for a passive firewall with physical interfaces and a passive firewall with aggregate interfaces without Link Aggregation Control Protocol (LACP) passive pre-negotiation configured. If the passive firewall has aggregate interfaces with LACP passive pre-negotiation configured, it always appears as **Receiving logs** because it continually sends learned ARP entries to IoT Security.

📋 *If you upgrade your firewalls from PAN-OS 9.x to 10.0 or later and notice that passive firewalls in HA pairs that appeared as Active in IoT Security now appear as Inactive, check if they have aggregate interfaces and if they have LACP passive pre-negotiation configured.*

Firewalls send log events to the logging service, which streams them to IoT Security for analysis and, depending on your IoT Security subscription type, sends them to Strata Logging Service for storage. IoT Security then processes and analyzes the raw metadata it receives from the logging service and retains the data generated from its analysis for the following lengths of time:

• One month of data retention for device network traffic behavior

- One year of data retention for the following:

  - Device identity

  - Security alerts, risks, and vulnerabilities

  - (Medical IoT) Device utilization

  📋 *The above retention periods are for IoT Security. For more information about IoT Security data retention, see* IoT Security Privacy Sheet. *For information about Strata Logging Service data retention, see* Strata Logging Service Privacy Sheet.

The Firewall Log Type Status section shows whether or not IoT Security has received log events from the logging service for EAL, DHCP, DHCP ACK, ARP, traffic, and threat logs within the past 30 minutes. If it has, the status is **Live Data**. If it hasn't, the status is **No Live Data**.

Firewall Log Type Status						
Type	Status	Latest Log	Log Events	Log Bytes	Ave. Latency	Max Latency
Total	● Live Data	Feb 7, 2024, 20:50	📊 16.5B	23.87TB	📊 1 min	43 min
EAL	● Live Data	Feb 7, 2024, 20:50	📊 1.1B	1.89TB	📊 1 min	26 min
DHCP	● Live Data	Feb 7, 2024, 20:50	📊 4.5M	8.26GB	📊 1 min	6 min
DHCP ACK	● Live Data	Feb 7, 2024, 20:50	📊 543.2K	1.13GB	📊 1 min	7 min
Traffic	● Live Data	Feb 7, 2024, 20:50	📊 15B	21.50TB	📊 1 min	43 min
Threat	● Live Data	Feb 7, 2024, 20:50	📊 140.2M	280.15GB	📊 1 min	2 min
ARP	● Live Data	Feb 7, 2024, 20:50	📊 146M	193.76GB	📊 1 min	5 min
GTP logs ⓘ	● No Live Data	-	📊 0	0B	📊 0	0

When the status is **Live Data**, it does not mean that all active firewalls have sent log events to the logging service within the past 30 minutes. Although that is possible, you can only safely deduce that at least one active firewall has done so and that the logging service has then streamed whatever log events it received to IoT Security. However, if the status is **No Live Data**, you can safely conclude that within the past 30 minutes the logging service has received no log events from any firewalls.

📋 *The status of firewalls in the overview section as **Active** or **Inactive** is not real-time data. It's updated every 30 minutes on the hour and half hour. On the other hand, the firewall log type status is close to real time. Every time you refresh the page, the Firewall Log Type Status shows the current status of these four log types. As a result, a temporary mismatch can sometimes occur between the two status indicators.*

Hover your cursor over the graph icon in the Log Events column to see a panel pop-up with information about each type of log.

The panel contains a graph that shows the total number of log events that IoT Security received. When you set the time filter at the top of the page to **1 Week**, the data is displayed in seven 24-hour intervals covering the last 7 days. When you set the filter to **1 Day**, the data is displayed in six 4-hour intervals covering the last 24 hours. And when you set it to **1 Hour**, the data is displayed in six 10-minute intervals for the last 60 minutes. Hover your cursor over various data points to see a tooltip with more information about it.

Hover your cursor over the graph icon in the Average Latency column to see a panel popup with information about the latency between the time that a firewall uploads logs to the logging service and the time that IoT Security receives them.

When you set the time filter at the top of the page to **1 Week**, the average latency is displayed for each of the last 7 days. When you set the filter for **1 Day**, the average latency is displayed in six 4-hour intervals covering the last 24 hours. When you set it to **1 Hour**, the average latency is displayed in six 10-minute intervals for the last 60 minutes. Hover your cursor over various data points to see a tooltip with more information about it.

The remainder of the Firewalls page contains a table with all the firewalls subscribed to IoT Security services. You can use the column control tool (icon with three gray bars that appears above the table) to customize the data that appears in the table. In addition to the status, hostname, serial number, IP address (not visible by default) of a firewall and the version of PAN-OS running on it, the table displays several other data points. There are columns for the IoT dictionary version, application content version for App-ID, and firewall license type— Prod (Production), Eval (Evaluation), or Lab. You can also see the number of different types of log events from each firewall, the site where it's located (not visible by default), when it first connected with IoT Security, and when it was last active.

By default, the IoT Dictionary Version column is visible in the table. If it does not appear in the table, click the column visibility icon (three vertical bars) and select **IoT Dictionary Ver**. This column shows the version-build number of the dictionary file on each firewall. A new version is released every two weeks, with the build number incrementally increasing across versions. For example, version was 1-218 released, and then two weeks later (and two internal builds later) version 2-221 was released.

All firewalls should have the same IoT dictionary version; that is, the latest version. If a firewall is using an outdated dictionary—most likely because it cannot reach the update server—it cannot use Device-ID to enforce Security policy rules with complete accuracy. Take steps to restore its connectivity to the update server so the next time the firewall automatically checks its IoT dictionary version against the one on the server, which it does every two hours, it will detect a new version and download it.

> *Only firewalls running PAN-OS version 10.0 or later support Device-ID and IoT dictionaries. For firewalls running earlier versions of PAN-OS, a dash appears in this column.*

The application content version determines the type of protocol data in the logs a firewall sends to IoT Security. A low version might not generate the IoT protocol logs that IoT Security needs.

The far-right column provides options to move a firewall from one site to another.

If you click a firewall serial number, a pop-panel appears with information about the logs from this firewall. You can see if IoT Security is currently receiving live data, timestamps of the latest logs, the number of events received, the average latency, and the maximum latency within the time filter specified for the Firewalls page (**1 Week**, **1 Day**, or **1 Hour**).

Log Status for 012501000789

Type	Status	Latest Log	Log Events	Ave. Latency	Max Latency
Total	● Live Data	Feb 7, 2024, 21:10	73M	1 min	3 min
EAL	● Live Data	Feb 7, 2024, 21:10	1.2M	1 min	1 min
DHCP	● Live Data	Feb 7, 2024, 21:10	5.6K	1 min	1 min
DHCP ACK	● Live Data	Feb 7, 2024, 21:10	756	1 min	1 min
Traffic	● Live Data	Feb 7, 2024, 21:10	71.8M	1 min	3 min
Threat	● Live Data	Feb 7, 2024, 21:10	10.4K	1 min	1 min
ARP	● Live Data	Feb 7, 2024, 21:10	53.4K	1 min	1 min
Tunnel	● No Live Data	-	0	0	0

Done

192

IoT Security Integration Status with Prisma Access

In the IoT Security portal, the Sites and Firewalls pages provide the status of next-generation firewalls with active IoT Security subscriptions. They show the total number of firewalls at each site, the connection status of each firewall, the total number of log events they've forwarded to logging services, and the types of logs they're sending. However, when Prisma Access subscribes to IoT Security through the IoT Security add-on, the information displayed on these pages is unlike that shown for next-generation firewalls.

Sites

When Prisma Access is using an IoT Security add-on, the site name for it on the **Networks** > **Networks and Sites** > **Sites** page is simply "Prisma Access". Whether a single Prisma Access instance is protecting one or a hundred remote sites, IoT Security remains unaware of their number. From the perspective of IoT Security, the numbers of devices and IoT devices come from a single Prisma Access entity regardless of how many remote sites it protects.

The following screen capture shows a mixed deployment of Prisma Access and several sites with on-premises next-generation firewalls for comparison.

The Sites page contains the following types of information for Prisma Access:

Status: A green cloud means that IoT Security is connected to Prisma Access and is receiving logs. A red cloud with a line through it means that IoT Security does not detect logs forwarded from Prisma Access to Strata Logging Service.

Name: Prisma Access

Location: This is the site location, if a location was previously defined.

Devices: This is the total number of devices that IoT Security identified across all remote sites under Prisma Access protection.

IoT Devices: This is the total number of IoT devices that Prisma Access identified across all its remote sites. This is a subset of the total shown in the Devices column.

Risk: This is the overall risk score calculated for all IoT devices protected by Prisma Access.

Subnets: These are the subnets across all Prisma Access remote sites. Because IoT Security has no visibility into how many sites Prisma Access is protecting, this might come from a single site with a single subnet, a single site with multiple subnets, multiple sites each with a single unique subnet, multiple sites with multiple subnets, or any combination of these scenarios.

Group: This indicates the group within the hierarchical site organization where the site is positioned.

Source: If IoT Security has a third-party integration with BlueCat IPAM or Infoblox IPAM and learns site names from there, the name of the integration appears here. When a third-party integration isn't how IoT Security learned of a site, a dash appears here.

Firewalls

This page (**Administration** > **Firewalls** > **Firewalls**) is not particularly applicable to Prisma Access. If you are using IoT Security exclusively with Prisma Access, the top of the page shows a total of two sites, one for Prisma Access and one for the default site, which is where IoT Security initially assigns on-premises firewalls. The Active and Inactive status will be 1 or 0 depending on whether IoT Security detects any logs from Prisma Access to Strata Logging Service in the last 30 minutes.

IoT Security displays the number of system alerts relating to Prisma Access. These pertain to the reception of requests from Prisma Access for policy recommendations and IP address-to-device mappings. For example:

```
IoT Security hasn't received any requests for policy recommendations in
the past 30 minutes.
```

```
IoT Security is receiving requests for IP address-to-device mappings
again.
```

Click the number of system alerts at the top of the Firewalls page to open **Administration** > **System Events** to see them. The source for Prisma Access system alerts is always `All firewalls`.

The rest of the Firewalls page doesn't have any data relevant to Prisma Access.

If your deployment includes a mix of Prisma Access and on-premises next-generation firewalls, then this page contains the information mentioned above for Prisma Access and much more information about firewalls and the logs they provide.

Data Quality Diagnostics

The quality of the network data that firewalls process and forward to the logging service directly impacts the quality of analysis that IoT Security is able to make. The **Administration** > **Data Quality** page is where you can see the quality of data that IoT Security has to work with. Two key factors are IP endpoints and low-confidence devices.

IP endpoints are devices without a unique identifier, making them untrackable over time. When IoT Security cannot locate a unique device identifier for a device, it categorizes it as an IP endpoint. This typically happens when IoT Security knows the IP address but not MAC address of a device through DHCP or ARP, and when IoT Security knows the IP address of a device but its device profile isn't stable enough to classify it as a static IP device. In the first case, the MAC address is the unique identifier for a DHCP client. In the second case, the IP address is the unique identifier for a static IP device if its profile is stable enough to show that the IP address isn't shifting among different DHCP clients.

Low-confidence devices are devices that IoT Security can identify with a confidence level under 70%. One of the fundamental services that IoT Security provides is identifying network-connected devices and assigning device profiles to them. It considers a host of factors throughout this process and creates a confidence score for each identification. The score is a number between 0-100, with 100 being the most confident. The confidence level is important because IoT Security only sends a firewall an IP address-to-device mapping if the confidence score for a device identity is high (90-100%), and if it has sent or received traffic within the past hour.

📄 *A confidence score indicates the level of confidence IoT Security has in its identification of a device. IoT Security has three confidence levels based on calculated confidence scores: high (90-100%), medium (70-89%), and low (0-69%).*

When firewalls forward fewer data logs to the logging service for IoT Security to analyze, it tends to identify devices less confidently. On the other hand, when firewalls forward more logs to the logging service, the more confidently IoT Security can identify devices and the more thoroughly it can baseline their behaviors. This results in higher device identity confidence scores.

This page shows the number of IP endpoints and low-confidence devices on the network and the percent of devices that fall into these two categories in relation to the overall number of devices on the network. You can infer the quality of device data that IoT Security is receiving from these numbers, which are taken from all devices over the last 30 days.

Each deployment has its unique characteristics and your reason for using IoT Security will determine the acceptable percent of IP endpoints and low-confidence devices on the network. For example, if your goal is to discover, identify, and protect only IoT devices, you might only use IoT Security with one or two firewalls near them. In this case, an acceptable percentage of IP endpoints and low-confidence devices would be fairly close to the percentage of non-IoT devices on the network. In short, consider what your goal is and use the data here to see how close you are to it. If there are more IP endpoints and low-confidence devices than you would like on your network, consider the recommendations offered on the page and follow those you think will reduce these numbers.

It's good practice to check Data Quality Diagnostics weekly for the first few months after deployment to make sure IoT Security is getting the data it needs to identify devices and, if not, make adjustments as needed. After you're satisfied, return periodically for spot checks and as follow-up whenever there are changes to the network.

Authorize On-demand PCAP

The On-demand Packet Capture (PCAP) feature for next-generation firewalls allows you to authorize the IoT Security Research Team to perform packet captures and automatically upload the captured packet files to IoT Security for offline analysis. The IoT Security Research Team takes packet captures only when necessary, such as when an unknown device or an unknown application appears on your network and the information required to assess the situation can be obtained no other way. The scope of such packet captures is limited so that they don't affect normal firewall operations.

PCAP files are securely stored and only accessed by IoT Security Research Team members. The files will be deleted either manually after an analysis is complete or automatically after 30 days elapse.

For the IoT Security Research Team to use PCAP to collect network traffic metadata from a firewall, you must first authorize the firewall to allow packet capturing.

> *To support PCAP on firewalls, they must be running:*
> - *PAN-OS 11.0.4 or later 11.0 releases*
> - *PAN-OS 11.1.0 or later*

STEP 1 | Log in to PAN-OS and install the openconfig plugin.

    1. Select **Device** > **Plugins** and search for `openconfig`.

    2. **Download** version 2.1.0 or later and then **Install** it.

STEP 2 | Log in to the IoT Security portal with a user account with administrator or owner privileges.

STEP 3 | Authorize PCAP on one or more firewalls.

    1. Select **Administration** > **Firewalls** > **On-demand PCAP** and then click the **Add** (**+**) icon.

    2. Choose the firewall either by its serial number or by the concatenation of its serial number and name.

    3. Set the time period to authorize PCAP on the firewall, which can be for 1 month, 3 months, or an unlimited length of time.

       When an authorization period expires, PCAP is no longer authorized on the firewall. If you want, you can reauthorize PCAP on it. You can then see the new PCAP authorization period in the list of authorized firewalls.

    4. **Confirm** the authorization.

    5. To authorize PCAP on additional firewalls, repeat these steps.

STEP 4 | Unauthorize PCAP on one or more firewalls.

When you want to deauthorize PCAP on firewalls.

1. Select one or more firewalls in the Authorized Firewalls list.

2. **Unauthorize** the selected firewalls.

> *If you want to deauthorize PCAP on just one firewall, you can also click the* ***Reauthorize*** *icon for it in the Actions column.*

198

IoT Security Integrations with Third-party Products

After IoT Security identifies IoT devices on your network and discovers if they pose any security threats, it works with next-generation firewalls—and also with Prisma Access—to protect your devices and network. In addition, you can integrate IoT Security with third-party products to expand the use of their specific features to include IoT. For example, when a network access control (NAC) solution integrates with IoT Security, it can allow or deny network access to IoT devices whose identity it would otherwise be unaware of. IoT Security users can also send a NAC system or a wireless LAN controller commands to quarantine IoT devices that have vulnerabilities or for which there are security alerts. Sometimes an integration works in one direction with IoT Security sharing its device information with a third-party product, and sometimes it works the other way with IoT Security learning device information from a third-party product. Other integrations enhance IoT Security functionality, such as its integration with third-party vulnerability scanners.

There are two options for integrating IoT Security with third-party systems and a third option for integrating Cortex XSOAR with IoT Security through its API:

- IoT Security public cloud with a cohosted, limited-featured Cortex XSOAR instance (requires the purchase of an IoT Security Third-party Integrations Add-on, which comes with an automatically generated, cohosted XSOAR instance at no extra charge)

  > *An IoT Security third-party integrations add-on does not require the purchase of a full Cortex XSOAR product. After you enable the add-on, IoT Security automatically generates a cloud-hosted XSOAR instance with limited functionality (in contrast to a full Cortex XSOAR product) to assist IoT Security with the integrations it supports.*

- IoT Security with an on-premises, full-featured Cortex XSOAR server

- Full-featured Cortex XSOAR instance with access to the IoT Security API

For information about the third-party integrations that IoT Security supports, see the IoT Security Integration Guide.

IoT Security and FedRAMP

Federal Risk and Authorization Management Program (FedRAMP) is a U.S. government program that promotes the use of secure cloud services by the federal government. Cloud computing systems categorized at the Moderate security impact level in accordance with the FIPS Publication 199 security categorization are authorized to store and process government data. The Palo Alto Networks IoT Security cloud is FedRAMP Moderate authorized.

The IoT Security FedRAMP Moderate solution is intended for use by U.S. government agencies requiring a standardized approach to the security assessment, authorization, and continuous monitoring of cloud products and services. It is also intended for use by commercial entities that do business with the U.S. government. The IoT Security FedRAMP Moderate solution operates as a separate and distinct entity.

The IoT Security commercial solution and the IoT Security FedRAMP Moderate solution have the following differences:

- You must purchase an additional SKU to get an IoT Security FedRAMP Moderate solution: PAN-IOT-FEDRAMP-MOD.

- The IoT Security FedRAMP Moderate solution permits only FedRAMP-authorized personnel access to data.

- Because Palo Alto Networks enforces strict incoming security policy rules for FedRAMP tenants, you must provide Palo Alto Networks customer services with a list of IP addresses for the administrative users who will be accessing your IoT Security portal. When user traffic to the portal passes through a NAT device on a perimeter firewall, edge router, or VPN gateway, provide the IP address to which NAT translates the users' original IP addresses. After you submit a support ticket with these addresses, customer services will create an allow list for the addresses you provided, which will let users log in from these addresses and access the portal.

- When integrating with third-party products, use a full on-premises Cortex XSOAR server. FedRAMP recommends running on-premises components of the solution using a vendor-approved FIPS version that complies with the FIPS 140-2 standard.

    > *Using an on-premises Cortex XSOAR server for IoT Security third-party integrations does not require the purchase of an IoT Security Third-party Integration Add-on license.*

When using next-generation firewalls in FIPS mode, IoT Security policy rule recommendations and enforcement based on Device-ID are not currently supported. In fact, because the primary purpose of the IoT Security FedRAMP Moderate solution is the discovery and identification of devices on the network and not the enforcement of Security policy rules, IoT Security does not even send firewalls the IP address-to-device mappings that are essential for Device-ID to work.

For more information about Palo Alto Networks IoT Security FedRAMP authorization, visit these websites:

- Official website for FedRAMP

- Palo Alto Networks solutions on FedRAMP Marketplace

- Palo Alto Networks website for FedRAMP Authorized Services

Discover IoT Devices and Take Inventory

IoT Security uses multiple methods to discover IoT devices and create a dynamic inventory.

- IoT Device Discovery
- IoT Security Devices Page
- IoT Security Device Details Page
- Create Multi-interface Devices
- Devices with Static IP Addresses
- Upload a List of Static IP Devices
- Add a Static IP Device Configuration
- Upload a List of Subnets with Only Static IP Addresses
- Add a Subnet with Only Static IP Addresses
- IP Endpoints
- Discover Mobile Device Attributes
- Custom Attributes
- Tag Management

IoT Device Discovery

Unlike IT assets that are generally multi-purpose hardware, IoT devices are purpose-built systems. These devices are designed to perform a few tasks on a very repetitive basis, and the IoT Security solution provides deep visibility into normal and suspicious network behaviors.

Each IoT device exhibits unique characteristics on the network. When an unknown device joins the network, one or more Palo Alto Networks firewalls log its network traffic and then send the logs to the logging service. These logs include session logs, containing metadata about traffic flow, and enhanced application logs, containing data from packet payloads. IoT Security accesses the data from the logging service and uses its advanced machine-learning algorithms and three-tier profiling system to analyze network behaviors and form a baseline for the device. It then compares that baseline with the behaviors of other known devices (for more information, see IoT Security Overview). By doing so, it determines the unique personality of the device and creates a profile for it consisting of device type, category, vendor, model, operating system, and many more. IoT Security automatically builds a behavioral profile for the device, including a baseline of acceptable behaviors and communication patterns with other devices.

IoT Security continuously learns and maintains a rolling baseline of device behaviors. The time required for building an initial profile depends on several factors:

- How active are the devices on the network? IoT Security can profile a device that produces a lot of traffic faster than a device that produces a little because it has more data to analyze.

- How many devices of the same type are there on the network? The more devices of the same type there are the faster the profiling works because it can aggregate knowledge learned from multiple devices simultaneously.

- How complicated is the behavior of an individual device? For example, IoT Security learns the behavior of a network-connected thermostat much faster than that of a surgical robot in a hospital.

The devices that IoT Security discovers on the network and identifies appear on the Devices page in the IoT Security portal.

IoT Security Devices Page

This page (**Assets** > **Devices**) is where you can see an inventory of all the devices that were discovered or are being monitored and the device profiles applied to them. There are three sections on this page: filters to control what data appears on it, a high-level summary of the devices on your network, and the device inventory table.

At the top of the page are filters to control the data displayed by site, monitoring status (Monitored Devices or Discovered Devices), device type, and time period. This same set of global

filters is at the top of the Devices page and the Dashboard. Whatever global filters you set in one section persists when you navigate to the other. These filters control what to display and what to download. Whatever is currently active is what you save when you click the **Download** icon (⬇️) > **Download**. For each device in the report, IoT Security includes whatever data it has for all the inventory table columns, whether they are currently visible or not at the time of the download.

There are two other options in the Download menu. Clicking **Create report** opens a new browser window or tab in which you can configure one of the following types of scheduled reports: Summary, Risk, New Device, and Filtered Inventory. Clicking **Download change log** and selecting two dates generates a CSV-formatted file that compares changes in your device inventory on your two selected dates. IoT Security checks and reports changes in data fields such as category, profile, profile vertical, OS group, device model, IP address, and subnet.

Clicking the pie chart or clicking content in the table lets you view device data at multiple levels of granularity.

At the top of the inventory table is a search tool, which allows you to search for device names. You can search for a full or partial match. If you employ a naming convention that identifies all devices by function, location, or some other characteristic, this allows you to search by that part of the name shared by all the devices in a particular grouping.

There is also a tool for creating custom filters that control what is IoT Security displays in the Inventory table. To create and apply a new filter or to apply a previously created filter, click the **Filter** icon (≡).

In the Filters dialog box that appears, select a previously defined and saved filter or click in the **Create a new filter** field and choose a device characteristic by which you want to filter devices.

Enter a value for the characteristic you want to use to filter devices.

Decide if you want to include global filters in your custom filter or not. When you select **Include global filter data**, you can control the global filters for sites, device types, and time whenever you apply the custom filter you are defining. Your custom filter can use either the current global filters or, if you modify them in the settings, the modified global filters. If you do not select **Include global filter data**, your custom filter will use whatever global filters happen to be in effect at the time you apply it.

Click the star icon to save the filter for future use. Click **Apply** to use it to filter the contents of the Inventory table now.

You can rearrange the columns in the device inventory table by click-dragging column headings into different locations.

You can also change which columns appear in the table. Click the **Columns** icon (three vertical bars), select the names of the columns you want to see, and clear the ones you want to hide. The columns with selected check boxes appear and those with cleared check boxes do not. Use the search tool to find column headings quickly.

To return to the default set of columns, **Reset to default**.

If you select the check box for one or more devices, the Download and Edit buttons appear.

When you click **Edit**, a dialog box opens where you can change the device type between IoT and Traditional IT and define other device characteristics: category, profile, vendor, model, OS family, OS version, location, asset tag, serial number, user tags, and description.

When you edit a device manually and change any of its attributes, your change is considered definitive and won't be overwritten. Therefore, be careful when manually editing a device because you're locking in your edits.

Whenever you manually edit a device, the modifications are fed into machine learning. If IoT Security determines the input is valid, it retrains its models with the added or modified data and propagates the results to all its customers. IoT Security then applies its revised models to other devices of the same type in all customer environments.

If you type something in the category field, and there isn't an existing category, a "Request New Category" option appears.

Use this option to request that IoT Security create a new category for the device. If the request is validated, then the category is added—not just for the person requesting it but for all IoT Security customers.

208

When you select multiple devices to edit, a table appears at the bottom of the dialog box for convenience. It displays the current values for your selections. If you mistakenly selected one that you don't want, you can spot it here.

IoT Security Device Details Page

To see details about a device, click its device name. The IoT Security portal then displays the device details page, with content grouped into the following sections:

- Identity
- Active Directory Attributes (appears when Cloud Identity Engine integration is enabled)
- Security (summary)
- Risks
- Alerts
- Security
  - Network Traffic
  - Applications
  - Software Components
  - Network Usage
- MDS2 (for medical IoT devices)

Identity: The Identity section at the top of the page provides identifying data such as the category and profile of a device, its vendor and model, its OS, and various network-specific details.

> *The IoT Security portal only shows a field if it has a value for it. You might see more or fewer details than shown here, depending on the amount of information IoT Security has.*

Active Directory Attributes (appears when Cloud Identity Engine integration is enabled)

If you have on-premises Active Directory (AD) synchronized with Cloud Identity Engine(CIE) and have a CIE tenant in the same tenant service group (TSG) as your IoT Security tenant, you can integrate IoT Security with CIE. Through this integration, you can identify devices discovered by IoT Security that are part of your AD and collect some AD attributes for display on the Device Details page. To view only devices that are in Active Directory, you can filter and search for devices in your inventory by their AD join status.

To integrate IoT Security with CIE, log in to the IoT Security portal as a user with owner privileges, select **Integrations** > **Cloud Identity Engine Integration**, and toggle the integration on. The toggle is in the upper right of the page.

> *The External Link icon ( ) opens the portal of your CIE tenant.*

Because IoT Security learns from the hub if a CIE tenant is part of its TSG, it will either let you enable integration if IoT Security and CIE are both tenants in the same TSG, or the toggle will be inoperable if they are not. Assuming you can enable integration, IoT Security will do an immediate retrieval of Active Directory attributes only if it's the first time or if the last sync was more than 24 hours ago and then do a daily retrieval every 24 hours going forward. (Toggling the integration off and back on won't trigger a new sync if it's less than 24 hours since the last one.) When you enable the toggle, IoT Security connects with your CIE and starts matching devices against the CIE/AD database to identify which ones are in your AD. The matching process compares the device name in IoT Security with the Common Name in AD. For devices that are in AD, IoT Security also retrieves the following attributes for display on the Device Details page:

Device attributes learned from Active Directory

AD Domain	OS
Common Name (IoT Security looks for Common Names in Active Directory that match Device Names in IoT Security. When it finds a match, IoT Security then retrieves device attributes from Active Directory.)	OS Version

Distinguished Name	OS Service Pack
Security Accounts Manager (SAM) Account Name	Serial Number
AD Groups	Last Login (This is the last time a device authenticated to AD. It comes from the AD lastLogon attribute.)

When CIE integration is enabled, these attributes are displayed in columns on the **Assets > Devices** page and in an Active Directory Attributes section on Device Details pages. IoT Security displays the source for attributes learned from Active Directory through CIE integration as **On-prem AD via CIE**.

For most device attributes, IoT Security uses the latest value it learns regardless of whether it's discovered through network traffic or through an integration. However, there are eight attributes for which a value learned through network traffic has priority even if IoT Security later learns of a different value through integration:

Device attributes whose values when learned through network traffic have priority over values learned later through integration

Model	Firmware
Vendor	Serial number
OS group	Wired or wireless
OS version	VLAN

If IoT Security learns a conflicting value for one of these attributes, it prioritizes the value learned through network traffic first and then through an integration (including CIE integration) second. The basic logic is as follows:

- Whatever new value is learned through network traffic replaces a value learned previously by any means.
- A new value learned through integration will replace a previously learned value learned through the same type of integration. It won't replace a value learned through network traffic or through another type of integration.

Security (summary): The information in the next section relates to security and includes the individual risk score for the device and whether baseline modeling is complete or still in progress. The current behaviors diagram shows evaluations for five types of behavior ranging from normal (near the center) to anomalous (near or beyond the edge).

When the Device Details page is for a medical device for which IoT Security has an MDS2 file, it displays information about device capabilities and operational states learned from the file such as the following:

- Remote service and patch support
- Personal health information (PHI) types and transmission support
- Antivirus installability and patchability
- Data storage and encryption
- Whether unnecessary applications and ports have been disabled
- Device communications both within and outside its local network
- Support for external user authentication

IoT Security uses the attributes listed in the MDS2 file to adjust the baseline risk level of the device. Risk factors based on MDS2 attributes contribute to a portion of the overall device risk score.

Risks: The Risks section contains the alerts, vulnerabilities, and anomalies that occurred to the device during the time range set at the top of the page. The events are displayed along a timeline and in a list with detailed information about each one.

1 Month ▼   Risk Score Trend | Alert History

3 Risks Associated

Ⓥ Uncontrolled New Internet Destinations   (Action ▼)

Status New   Assignee Assign   Risk Events

An abnormally high number of new Internet destinations can cause anomalous network behaviors. Making uncontrolled connections to a large number of new Internet destinations poses a security risk for ... More

Impact on Risk Score	▬ ▬ ▬
Uncommon URL count	1209
Normal Average	13

Uncommon Components   [blurred IP addresses]   More

Direction   Outbound
Zone   External

Ⓥ New Internet Application Usage   (Action ▼)

Status New   Assignee Assign   Risk Events

Starting multiple new or unauthorized Internet applications on this IoT device is not normal. It deviates from the baseline and poses a security risk.

Impact on Risk Score	▬ ▬ ▬
Uncommon service count	29
Normal Average	1

When IoT Security has recommendations for responding to a risk, it displays **More Insights**. Click it to expand the section and read more about how the impact of the risk on the device and network and what you can do to address it.

Alerts (3)   Hide ∧

☰ Active (3)   ∿ 1 month alerts ▾

⚠ **NetBSD tnftp Url Fetching Command Execution Vulnerability (CVE-2014-8517)**   ⊕ Action ▾
NetBSD tnftp is prone to a command execution vulnerability while parsing certain crafted HTTP responses. The vulnerability is due to the lack of proper checks HTTP
responses, leading to an … More

		Alert Events
Impact on Risk Score	▬▬ ▬	● Alert Detected
Alert Type	Vulnerability	03:47, June 07, 2020
device profile	Tridium Controller	
client port	31486	
threat ID	37835	
threat category	code-execution	
threat type	vulnerability	
number of occurrences	1	
CVE	CVE-2014-8517	
reference	reference	
alert source	Firewall	
firewall name	21542-bisma-p5250-m	
firewall action	Terminated the session and sent a TCP reset to both sides of the connection	
firewall inbound interface	ethernet	
firewall outbound interface	ethernet	

More Insights ∧

Recommendation

Install software updates in a timely manner to prevent the exploit of known vulnerabilities.
Check network traffic coming to and from the device on the device details page and enable trusted
behavior by applying an ACL (access control list) to allow only essential traffic to and from resources
at specific IP addresses.
Take the device 00:01:f0:90:2d:78 offline.

⚠ **NetBSD tnftp Url Fetching Command Execution Vulnerability (CVE-2014-8517)**   ⊕ Action ▾
NetBSD tnftp ... ... vulnerability while parsing certain crafted HTTP responses. The vulnerability is due to the ...

For medical IoT devices with MDS2 risks that were summarized near the top of the page, the risks are also listed with a few more details here. IoT Security displays them after any other detected vulnerabilities.

220

Alerts: This section contains only the alerts that the device raised during the specified time range. Alerts are a subset of risks, and IoT Security generates them when it detects irregular behavior and activity matching an alert rule. You can see when alerts occurred along a timeline, read details about them, and take action to resolve them.

Security: The Security section contains three subsections that show how a device connects to other devices on the network and which applications it's using.

- **Network Traffic**: View a conceptual network topology displaying the nodes with which the device has formed connections. Use filters to display inbound or outbound connections; nodes

with various alert levels; connections to nodes within the same VLAN, same intranet, or in the Internet; and so on.

If you click **Explore Topology**, a new browser window opens with an informative display of internal and external connections from the device in focus. You can interact with the

information, viewing details about each node and clicking different ones to put them in focus and see their connections.

📋 *Any node with "S" on it is a server.*

To learn more, watch a pair video explanations of the Topology Explorer. Part 1 covers navigation, information pop-ups, zoom, device category filters, and SMB filters. Part 2 looks at the information panel, how to explore the topology, and how to start a new path. Each video is about two to three minutes long.

- **Applications**: This section shows the applications the device uses, their risk levels (a 1-5 scale with numbers closer to 5 indicating increased risk), and how many other devices and device profiles use the same application. Click a number in the Used by Devices column to open the Devices page with its contents filtered by the corresponding application. Hovering your cursor over the blue text of an entry in the Profiles column displays a list of all profiles that use that application.

Applications

Name	Risk Level	Used by Devices ▾	Profiles
ping	2	7960	Access Control Device and 96 more
dns	3	6848	Access Control Device and 78 more
ldap	2	6106	Android and 29 more
web-browsing	4	5014	Airtame Wireless Presentation Device and 65 more
ntp	2	4712	Airtame Wireless Presentation Device and 68 more
dhcp	2	3633	Airtame Wireless Presentation Device and 52 more
ipsec-esp-udp	2	1804	Android and 28 more
unknown-udp	1	1783	Android and 31 more

36 Applications

1 to 8 of 36   ‹   ›

- **Software Components**: Most software makes use of various third-party software components such as libraries, modules, binaries, compilers, executables, files, and source code. The details of these components are increasingly being documented within Software Bills of Materials thanks to the Software Component Transparency initiative led by the U.S. National Telecommunication and Information Administration and with the participation of numerous manufacturers. A Software Bill of Materials (SBOM) is a comprehensive record detailing all the bits and pieces of software within a system or device and their relationships with each other. It's essentially a nested inventory of software components and subcomponents, including firmware and embedded software. It also typically includes licensing, author, and version

information plus other metadata. The purpose is to provide as much transparency as possible into the software contents running on devices so that we can better protect them from attack.

Some exploits specifically take advantage of the very lack of transparency and target vulnerabilities that occur in software components such as Spring4Shell, Urgent/11, Ripple20, and Log4j 2. Knowing which software components are on a device can expedite vulnerability detection, risk analysis, and remediation efforts. For example, the Log4j 2 vulnerability affects specific versions of the Apache Log4j 2 Java logging library, an open-source Java-based logging framework used by Java applications around the world. Attackers can exploit the vulnerability to launch denial-of-service attacks or gain remote control of target devices. The first step in responding to this threat is to identify which devices use the Log4j 2 Java logging library and, if so, if it's a vulnerable version. With IoT Security, you can search your inventory for devices using this particular library and version–or for devices vulnerable to one or more of the related CVEs–in just seconds and save days or even weeks of response time.

IoT Security primarily learns SBOM information from traffic inspection of, for example, the user agent field in HTTP headers and to a lesser degree from other sources like FTP banners and HTTP URL information. It then shows the software components and version numbers identified in the SBOM for a device in the Software Components column on the Devices page. IoT Security also shows the software component name, version number, and any related CVEs in the Software Components section on the Device Details page.

You can download a device inventory report from the Devices page. The report includes a list of software component names and version numbers for all devices with software libraries detected by IoT Security.

You can also download the software library details for an individual device in Software Package Data Exchange (SPDX) format, which is one of the most common data standards for capturing

SBOM data. To download the SPDX file, click **Download SBOM** at the bottom of the Software Components section. You can then open and read the SPDX file with any standard text editor.

> 📋 *The amount of data IoT Security learns is limited to whatever SBOM information devices send over the network and by what can be extracted from network traffic.*

- **Network Usage**: The last section shows a Sankey diagram with lines indicating network connections. The red line indicates it's involved in an alert of high severity. Click one of the blue bars and then click the **Create Policy** option that appears to create a policy with the following fields in the Policy Editor auto filled: ("Group #1" = source, and "Group #2 = destination).

MDS2 (for medical IoT devices)

Medical device vendors often list the security-related features of their products in Manufacturer Disclosure Statement for Medical Device Safety (MDS2) forms, which they share with their customers. Vendors issue these MDS2 documents for each version of a medical device and include valuable information such as whether a device processes PHI (personal health information); if it stores PHI and, if so, if it's encrypted; and if antivirus software is installed on the device.

Manufacturer Disclosure Statement for Medical Device Security – MDS²			
DEVICE DESCRIPTION			
Device Category	Manufacturer	Document ID	Document Release Date
Patient Monitor	Philips Medizin Systeme Boeblingen GmbH		May-2017
Device Model	Software Revision		Software Release Date
MX550, MX500, MX450, MX430, MX400, XG50, MX100, MMS X3	M.0		May-2017
Manufacturer or Representative Contact Information	Company Name Philips Medizin Systeme Boeblingen GmbH Representative Name/Position productsecurity@philips.com	Manufacturer Contact Information Philips Medizin Systeme Boeblingen GmbH, Hewlett-Packard-Str. 2, 71034 Boeblingen	

Intended use of **device** in network-connected environment:

The monitors are indicated for use by health care professionals whenever there is a need for monitoring the physiological parameters of patients. The monitors are intended to be used for monitoring and recording of, and to generate alarms for, multiple physiological para⁓⁓⁓⁓⁓ ⁓⁓s, and neonates. The monitors are intended f⁓⁓ ⁓⁓ ⁓⁓re professionals in a hospital ⁓⁓⁓P30/MP40/MP40⁓ ⁓ ⁓⁓itionally intended f⁓

⁓⁓rs are only f⁓ ⁓ome use. ⁓⁓
prescription use only.

MANAGEMENT OF PRIVATE DATA		
Refer to Section 2.3.2 of this standard for the proper interpretation of information requested in this form.	Yes, No, N/A, or See Note	Note #
A	Can this **device** display, transmit, or maintain **private data** (including **electronic Protected Health Information [ePHI]**)?		
B	Types of **private data** elements that can be maintained by the **device**:		
B.1	Demographic (e.g., name, address, location, unique identification number)?		
B.2	Medical record (e.g., medical record #, account #, test or treatment date, **device** identification number)?		
B.3	Diagnostic/therapeutic (e.g., photo/radiograph, test results, or physiologic data with identifying characteristics)?		
B.4	Open, unstructured text entered by **device** user/operator?		1
B.5	**Biometric data**?		
B.6	Personal financial information?		
C	⁓⁓ta - Can the **device**:		
	⁓⁓rarily in volatile me⁓⁓		

Over time, healthcare providers can collect thousands of MDS2 documents for thousands of medical devices. When used as intended, MDS2 documents can greatly enhance your security posture and incident response (IR). However, absorbing the details from these documents for the specific version of the software running on their connected devices is a daunting task. As a result, MDS2 files often go unused.

IoT Security simplifies the management and use of the MDS2 files you have. If you upload an MDS2 file for a device to IoT Security, it then includes this data along with other environmental factors when assessing the risk to the device. For example, if the software version of a device specified in an MDS2 file has a known vulnerability, IoT Security more precisely identifies it as a vulnerability instead of just a potential vulnerability. IoT Security supports MDS2 files in 2004, 2008, 2013, and 2019 formats.

To upload an MDS2 file for one of your medical devices, click the MDS2 button on the device details page, click the upload icon in the lower right corner, and then navigate to your MDS2 document (its format must be PDF) and upload it.

A prompt appears to apply the MDS2 file to all devices sharing the same model, vendor, and profile. To apply the MDS2 file to all devices with the same attributes, click **Yes**. To apply it to just this particular device, click **No**.

📋 *To upload MDS2 files and automatically apply them to all devices with matching model, vendor, and profile attributes, use the upload option on **Administration > MDS2**. For more information, see* MDS2.

An entry for the uploaded MDS2 file appears in the MDS2 section on the Device Details page with some upload details, device manufacturer name, and software revision number (if available). In addition, if you selected **Yes** when prompted to apply the MDS2 file to other devices with the same model, vendor, and profile and there are such devices, then IoT Security applies the uploaded MDS2 file to them as well.

The upload date shows when this file was uploaded to IoT Security.

📋 *The timestamp uses the time zone specified on the Preferences page (🛡 > Preferences).*

The source of an uploaded MDS2 file is always **Directly Uploaded**, which means that a user manually uploaded the file to IoT Security.

The status of an uploaded file indicates one of the following states:

- **Matched** – The uploaded file is a PDF containing correctly formatted fields
- **Cannot Extract Data** – The file is a PDF with incorrectly formatted fields
- **Unsupported File Type** – The uploaded file is not a PDF

If the file status is either of the last two states, hover your cursor over the table row with the MDS2 file and then click the Delete icon that appears on the far right (🗑).

To see more details about the device and MDS2 file, expand the row.

A manufacturer might release an updated MDS2, perhaps to add more models to the Device Model list, change its Manufacturer Contact Information, or for some other reason. If so, delete the first MDS2 file and then upload the new file.

To see a preview of an MDS2 file, hover your cursor over its table row, which causes the preview icon to appear ( ). Either click the icon or hover your cursor over it to see the file in a pop-up preview window.

Use the viewing options to scroll through the file and zoom in and out.

To view the file itself, click the filename. IoT Security downloads the PDF file so you can open and view it locally.

IoT Security uses several fields in MDS2 forms for risk detection:

• Can this device display, transmit, or maintain private data?

• What types of private data elements can be maintained by the device?

• Can security patches or other software be installed remotely?

The wording for these questions varies in different versions of MDS2.

This information can help IoT Security assess risk. For example, if an MDS2 file states that a device doesn't support remote servicing and IoT Security detects an inbound connection from an external source, it will flag this as anomalous behavior and generate a security alert. Similarly, if an MDS2 file states that a device cannot be remotely patched, any attempted inbound file transfer from an external location will also be treated as anomalous and trigger an alert.

Create Multi-interface Devices

Some devices have multiple network interfaces. These can be networking and security devices like L3 switches and firewalls with multiple network ports or physical endpoint devices, such as printers, that can connect to both wired and wireless networks.

Because each interface on a multi-interface device has its own MAC address and IP address, IoT Security initially considers each interface as a separate single-interface device. This can result in duplicate devices in your asset inventory and duplicate vulnerabilities. When IoT Security detects two or more devices that share common attributes, such as hostname or serial number, it provides a recommendation for you to group them as different interfaces on the same multi-interface device. In addition to accepting the recommendation as is, you can modify or ignore the recommendation and merge other devices instead. The merge process involves assigning one "device" as the primary interface and the others as secondary interfaces. When you do this, IoT Security applies the device-level attributes of the primary interface to the entire multi-interface device while retaining the network-level attributes for each interface.

Device-level attributes originally learned from the device assigned to be the primary interface and then applied to all merged interfaces	Network-level attributes originally learned on each previously unmerged device and retained for interfaces on the merged device
Category	IP address
Device name	MAC address
Endpoint protection (vendor)	OUI vendor (NIC vendor)
Model	Site
OS group	Status (network connectivity)
OS combined (OS group + OS version)	Subnet
Patient health information support (Medical IoT only)	Switch
Profile	Tags
Risk level	Wireless access point
Risk score	VLAN
Serial number	All network attributes except those for CMMS (computerized maintenance management system), EDR (endpoint detection and response), and External Inventory

Device-level attributes originally learned from the device assigned to be the primary interface and then applied to all merged interfaces	Network-level attributes originally learned on each previously unmerged device and retained for interfaces on the merged device
Type	All traffic attributes except the following: Software, Software Components, and Restricted Traffic.
Vendor	–

These attributes are assigned to a multi-interface device at the time individual *devices* are merged and become *interfaces* on a single device. After the merge, they can continue to change based on the network behaviors that IoT Security observes. IoT Security also merges vulnerabilities, security alerts, risk scores, and reports of the previously separate devices as they become interfaces on one device.

Merge Devices into a Multi-interface Device

You can merge one or more devices into a single multi-interface device based on IoT Security recommendations or create your own multi-interface device without recommendations. When IoT Security has recommendations, it displays a notification above the Inventory table on the **Assets** > **Devices** page.

STEP 1 | View the groups of two or more single-interface devices that IoT Security recommends be merged into multi-interface devices.

1. To see the list, click **View All Recommendations** above the Inventory table.

    A panel opens on the right of the Devices page showing all the devices that IoT Security recommends merging together and the reason for each recommendation.

2. Click the arrow to the left of a recommendation to see the individual devices to be merged.

    IoT Security displays the name and profile of each single-interface device that it recommends merging into one multi-interface device.

Clicking **Create** starts the merge process. Clicking **Dismiss** permanently dismisses the recommendation. However, if a dismissed recommendation changes—a device is added to the original recommendation or removed from it—IoT Security will make a revised recommendation.

STEP 2 | Merge individual devices into a single multi-interface device.

    1. Click **Create** for the multi-interface device you want to create.

       This launches a three-step process, the first of which is the selection of devices to merge. The devices that IoT Security chose appear in a Selected Devices section above the rest of the devices in the All Devices section.

2. Keep the IoT Security-recommended devices selected if you want to include them in the multi-interface device, clear any you want to exclude, and add more from the All Devices table if you think they should also be included.

Any devices that you select in All Devices are also shown in Selected Devices.

You can't add a previously merged multi-interface device to another multi-interface device.

3. When you're satisfied, click **Next**.

4. Select the primary interface of the multi-interface device.

While all interfaces retain their network-specific attributes (IP address, MAC address, subnet, and VLAN), the merged device will use the physical device attributes from the primary interface. You might consider choosing the interface that processes the most traffic because IoT Security most likely has the most data from this interface and, therefore, the most accurate device identification and risk analysis. If you have a dedicated management

subnet and VLAN on your network, another option is to choose the interface in that subnet and VLAN.

5. After you've selected the primary interface for the device, click **Next** and then expand different sections to review the merged attributes.

You can click **Expand All** to view all six sets of attributes at once and then **Collapse All** to close them together. You can reduce the height of expanded sections by clicking **Hide**

Empty Fields. To see all fields-–both those with data and those without—**click Show Empty Fields**.

> 📋 *You can also see this information later in the Attributes section on the **Device Details > New device page** after you create the multi-interface device.*

6. When you're satisfied and want to complete the merge process, click **Create**.

7. To see the merged device on the **Assets > Devices** page, add a filter to show multi-interface devices.

   The newly created multi-interface appears in the Inventory table with the multi-interface device icon ( ) after its device name.

8. Click the multi-interface device icon (**M**) to see its interfaces with the primary interface identified at the top, and to access the **Edit** and **Unmerge** options.

STEP 3 | (Optional) Edit a multi-interface device.

After creating a multi-interface device, you can later change the primary interface, merge more devices as interfaces into it, remove one or more interfaces from it, or unmerge all interfaces.

To change the primary interface on a multi-interface device:

1. Select **Assets** > **Devices**, click the multi-interface icon (M) to open the Interfaces panel for the device whose primary interface you want to change, and then **Edit**.
2. Click **Next** to advance to the step where you select a primary interface.
3. Select the interface that you want to make the new primary interface and then click **Next**.
4. Review the settings to make sure the new primary interface is the one you want it to be and then **Create**.

To add one or more interfaces to an existing multi-interface device:

1. Select **Assets** > **Devices**, click the multi-interface icon (M) to open the Interfaces panel for the device to which you want to merge one or more single-interface devices as interfaces, and then **Edit**.
2. Select one or more devices in the All Devices table that you want to convert from single, separate devices to interfaces on the multi-interface device and then click **Next**.
3. Either keep the previously selected primary interface in its role or make another interface the primary if you want and then click **Next**.
4. **Create**.

To remove one or more interfaces—but not all—and return them to the inventory as individual single-interface devices while keeping the multi-interface devices:

1. Select **Assets** > **Devices**, click the multi-interface icon (M) to open the Interfaces panel for the device whose interfaces you want to remove, and then **Edit**.
2. Clear the selection of the interfaces that you want to remove from the multi-interface device and then click **Next**.
3. Either keep the previously selected primary interface in its role or make another interface the primary if you want and then click **Next**.
4. **Create**.

To unmerge all interfaces:

1. Select **Assets** > **Devices**, click the multi-interface icon (M) to open the Interfaces panel for the device whose interfaces you want to unmerge, and then **Edit**.
2. **Confirm** the unmerge operation and return of each interface to an individual single-interface device.

Devices with Static IP Addresses

While most network-connected devices receive their IP addresses dynamically through DHCP, it's common to reserve part of the network address space for use as static IP addresses for devices such as routers, printers, FTP servers, and DHCP servers. Beyond this common practice, there are some industries and facilities that use static IP addresses predominantly; for example, manufacturing, utilities, oil and gas, warehouses, order fulfillment centers, and processing and distribution centers. Because most automation and control applications use the IP address directly in their programs, it's important that robotic devices and controllers in assembly lines and processing centers have static IP addresses, which is why static addressing is so prevalent in these areas.

IoT Security can be deployed in networks where DHCP dynamically assigns IP addresses to devices, where network administrators manually configure devices with static IP addresses, and where there's a combination of both. IoT Security uses multiple techniques for detecting and monitoring network activity and correlating it to individual devices. By examining the DHCP traffic logs that firewalls provide, it associates dynamically assigned IP addresses with device MAC addresses and adds these devices to its inventory. By looking at ARP logs, IoT Security also learns IP address-to-MAC address mappings and adds devices with static IP addresses, which might not otherwise be discovered through DHCP, to its inventory as well. However, by the very nature of ARP broadcasts, this only works for devices within the same Layer 2 broadcast domains as the reporting firewalls. For devices with static IP addresses beyond Layer 2 boundaries, IoT Security uses machine learning to discover network activity patterns indicating the likely presence of such devices. You also have the option of manually providing IoT Security with static IP address assignments through static IP device and subnet configurations.

> *Providing IoT Security with a static IP address configuration by itself is not enough to add a device to the inventory. IoT Security must also detect network traffic to or from a device with a configured static IP address. Then it adds the device to its inventory.*

Use one of the following methods to add static IP devices and subnets to the IoT Security inventory:

- Upload a List of Static IP Devices
- Add a Static IP Device Configuration
- Upload a List of Subnets with Only Static IP Addresses
- Add a Subnet with Only Static IP Addresses

IoT Security then uses the IP addresses of these devices (rather than their MAC addresses) to identify and track them.

Upload a List of Static IP Devices

If you have a list of the static IP addresses for your devices, enter them in a CSV (comma-separated values) file and upload it to IoT Security.

> *There is a limit of 10,000 static IP devices for each uploaded CSV file. If you need to upload more than 10,000, upload multiple CSV files.*

STEP 1 | Navigate to the User-Defined Static IP Devices page (**Assets** > **Devices** > **User-Defined Static IP Devices**) and then click **Add** > **Upload Static IP Devices**.

STEP 2 | Click the link to download the CSV template.

STEP 3 | Fill out the template with static IP device information or create a new file in the same format as the template and fill that out.

Enter the static IP address of each device you want to upload. Optionally enter its MAC address, vendor, and model in the columns indicated in the template. IoT Security accepts any of the following MAC address formats:

`aa:bb:cc:00:11:22`	`AA:BB:CC:00:11:22`
`aa.bb.cc.00.11.22`	`AA.BB.CC.00.11.22`
`aa-bb-cc-00-11-22`	`AA-BB-CC-00-11-22`
`aa bb cc 00 11 22`	`AA BB CC 00 11 22`
`aabbcc001122`	`AABBCC001122`

IoT Security uses IP addresses rather than MAC addresses to identify and track static IP devices. The additional user-configured attributes provide extra information when referring to entries on the User-Defined Static IP Devices page later. However, only the uploaded IP addresses and (if provided) MAC addresses will ever appear on the Devices and Device Details pages.

STEP 4 | Return to the User-Defined Static IP Devices page, click **Add** > **Upload Static IP Devices**, choose or drag the completed CSV file onto the space in the dialog box, and then **Upload**.

If IoT Security previously detected network activity from one of the uploaded IP addresses, it's considered a match. The Device Matches counter at the top of the page increases incrementally and "matched" appears in the Result column for this IP address. IoT Security then adds the static IP device to its inventory and displays it on the Devices and Device Details pages. It takes several minutes for IoT Security to check for potential matches with existing data and then update the inventory and static IP device list accordingly.

If IoT Security has not yet detected network activity for one of the uploaded IP addresses, it's considered "not found". In this case, the Devices Not Found counter increases incrementally and a dash appears in the Result column. If IoT Security later discovers network activity for this IP address, it moves it from "not found" to "matched", adds the static IP device to its inventory, and begins displaying it on the Devices and Device Details pages.

📋 *If a user-defined MAC address is different from a MAC address IoT Security detects on the network, the detected MAC address overrides the user-defined one.*

Add a Static IP Device Configuration

Instead of uploading a CSV file with a list of static IP devices (see Upload a List of Static IP Devices), you can add them individually.

STEP 1 | Navigate to the User-defined Static IP Devices page (**Assets** > **Devices** > **User-Defined Static IP Devices**) and then click **Add** > **Manually Add a Static IP Device**.

STEP 2 | Define a static IP device and then click **Add**.

IP Address: Enter the static IP address of the device you want to add to your inventory. The IP address is what IoT Security uses to track user-defined static IP devices.

MAC Address (optional): If you want, add the MAC address of the device in hexadecimal notation. IoT Security accepts any of the following MAC address formats:

`aa:bb:cc:00:11:22`	`AA:BB:CC:00:11:22`
`aa.bb.cc.00.11.22`	`AA.BB.CC.00.11.22`
`aa-bb-cc-00-11-22`	`AA-BB-CC-00-11-22`
`aa bb cc 00 11 22`	`AA BB CC 00 11 22`
`aabbcc001122`	`AABBCC001122`

> 📋 *If the user-defined MAC address is different from the MAC address IoT Security detects on the network, the detected MAC address overrides the user-defined one. If IoT Security does not detect a MAC address, the user-defined MAC address appears on the Devices and Device Details pages.*

Vendor (optional): Enter the vendor for this device.

Model (optional): Enter the device model.

The vendor and model attributes provide extra information when referring to entries on the User-Defined Static IP Devices page later. However, they do not appear on the Devices and Device Details pages.

STEP 3 | Click **Add** to add the configuration to IoT Security and then click **OK** to close the confirmation message that appears.

After you add the static IP device, IoT Security initially treats it as "not found". It incrementally increases the Total User-Defined Static IP Devices counter by one and the Devices Not Found counter by one. Although it adds an entry for it to the User-Defined Static IP Devices list, the Result column remains empty--there isn't a "matched" entry, indicating that IoT Security

detected network activity for this IP address, or a dash, indicating that no such activity was detected.

Because IoT Security periodically compares entries in the user-defined static IP devices list with those in its inventory and its internal database of detected IP addresses without accompanying MAC addresses, the page can remain in this initial state for several minutes.

If a match is found, the Device Matches counter increases by one and the Devices Not Found counter decreases by one. Also, "matched" now appears in the Result column.

252

If IoT Security does not find a match, it eventually displays a dash in the Result column.

You might have to reload the User-defined Static IP Devices page to see the updated data.

Upload a List of Subnets with Only Static IP Addresses

In the case where an entire subnet consists of static IP addresses, it's more efficient to add a subnet and define it as having static IP addresses than adding numerous static IP devices individually. When you have multiple subnets with static IP addresses, you can upload all of them in a CSV file at once.

There is a limit of 10,000 subnets for each uploaded CSV file. If you need to upload more than 10,000, upload multiple CSV files.

After you provide IoT Security with a subnet configuration specifying that it has static IP addresses and then IoT Security detects traffic from a device in that subnet, it considers the device a static IP device. Using the IP address as the device ID (instead of a MAC address), it adds

the device to its inventory. IoT Security adds static IP devices to its inventory in this manner only for devices that are not discovered through other detection mechanisms such as ARP logs.

> *If you later remove a static IP subnet after IoT Security added static IP devices for this subnet to its inventory, IoT Security reverses this action and automatically removes them from its inventory.*

STEP 1 | Navigate to the Networks page (**Networks** > **Networks and Sites** > **Networks**) and then click **Add** > **Upload Subnets**.

STEP 2 | Click the link to download the CSV template.

STEP 3 | Fill out the template with subnet information.

Enter the following for each subnet you want to upload:

`prefix`: Enter the IP address of the subnet in dot-decimal notation and its netmask in CIDR notation (example, 10.1.1.0/24). This appears on the Subnets page, and for a device in this subnet, the subnet and netmask appear on the Devices and Device Details pages.

`vlan`: (Optional) Enter a VLAN ID. If entered, this also appears on the Subnets page, and for a device in this subnet, it appears on the Devices and Device Details pages.

`description`: (Optional) Enter a description of the subnet/VLAN, perhaps noting the type of devices for which it's intended. These special characters are not allowed in the description field: ~ ` ! # $ % ^ & * + = { } [ ] | \ < > ? This description only appears on the Subnets page.

`static`: Enter **yes** to define this as a subnet containing static IP addresses. When IoT Security discovers a device from a user-configured static IP subnet in a different L2 domain from the firewall and adds it to its inventory, the Source column on the Devices page shows **User-Configured**. (Leave it blank if you do not want the subnet to be static).

`monitored`: Enter **yes** if you want IoT Security to provide device profiling, behavioral analysis, and risk monitoring of the devices in this subnet. Leave it blank if you only want IoT Security to detect the devices in the subnet and perform a simplified device identity analysis. Based on this field, the IoT Security portal displays **Yes** or **No** in the **Monitored** column in

the Networks table on **Networks** > **Networks and Sites** > **Networks** tab and provides the appropriate level of device monitoring, analysis, and protection.

> 📋 *The **Add a subnet** option does not provide an option to specify a subnet as monitored or unmonitored. IoT Security automatically classifies an added subnet as monitored. However, you can change its classification after you add it by selecting the subnet and clicking **Stop Monitoring**above the Networks table. You can also make multiple selections to stop monitoring multiple subnets at the same time. Later you can select unmonitored subnets and click **Start Monitoring** above the table.*

Networks (61)

	Type	IP Prefix	Networ...	VLAN	Static ↓	Monitored
☑	Subnet ⓘ	17.17.17.0/24	default	17	Yes	Yes
☐	Subnet ⓘ	126.1.2.0/24	default	5	Yes	No
☐	Subnet ⓘ	5.3.9.0/24	default	32	Yes	Yes

STEP 4 | Upload the CSV file.

On the Networks page, click **Add** > **Upload Subnets**, choose or drag the completed CSV file onto the space in the dialog box, and then click **Upload**.

If IoT Security previously detected network activity from an IP address in one of the uploaded subnets, it now considers it a static IP address and automatically adds a static IP device to the inventory on the Devices page. Similarly, if IoT Security later detects traffic from an IP address in one of these subnets, it automatically adds an entry to the inventory at that time.

> 📋 *It can take several minutes for new entries to appear on the Devices page.*

Add a Subnet with Only Static IP Addresses

Instead of uploading a CSV file with a list of static IP subnets (see Upload a List of Subnets with Only Static IP Addresses), you can add them individually.

STEP 1 | Navigate to the Networks page (**Networks** > **Networks and Sites** > **Networks**) and then click **Add** > **Add a Subnet**.

STEP 2 | Define a subnet and then **Save**.

Type: Select **Subnet**.

Prefix: Enter the IP address/netmask of the subnet you want to add. Enter the IP address of the subnet in dot-decimal notation and its netmask in CIDR notation (example, 10.1.1.0/24).

Name (optional): Enter a name for the subnet

VLAN ID (optional): Enter the VLAN ID for the subnet.

Description (optional): Enter a description of the VLAN/subnet, such as the type of devices for which it's intended. These special characters are not allowed in the description field: ~ ` ! # $ % ^ & * + = { } [ ] | \ < > ?

Mark this subnet as static: Select.

It can take several minutes for new entries to appear on the Networks page. You might have to reload the page to see the updated data.

IP Endpoints

When IoT Security receives sufficient network traffic metadata, it uses AI and machine learning to identify the devices generating the traffic. However, there are times when it doesn't receive enough to identify devices uniquely. For example, IoT Security might be aware that there is traffic to and from a specific IP address but, because the device is in a different Layer 3 domain from the firewall logging the network traffic metadata, it never learns its MAC address. The device might be behind a router, a NAT device, or a wireless tethering device, so the firewall only gets its IP address. If DHCP is providing network settings to network devices, it's possible that different devices use the same IP address at different times. As a result, the network behavior associated with the IP address will keep changing as different types of device take turns using it. When IoT Security is aware of an IP address that is the source and destination of traffic but it doesn't know its MAC address and the network behavior isn't stable enough to deduce that it's a statically assigned IP address, IoT Security categorizes it as an IP endpoint.

Another way that IoT Security can learn about IP endpoints is through third-party integrations. IoT Security can receive device data by integrating with a network management or asset management solution and by using SNMP to query network switches about the devices connected to them.

If IoT Security observes stable traffic patterns associated with an IP endpoint and there are no changes to any of its major device attributes for seven days, it moves it to the Devices page. There are eight major device attributes that IoT Security watches for changes: device profile, category, vendor, model, OS, hostname, serial number, and site ID. A change to any of these attributes indicates that the device using the IP address has changed, so if they all remain unchanged for seven days, it's reasonable to assume that the device identity is stable.

After adding the IP endpoint to the Devices page, IoT Security continues tracking its attributes on a daily basis. If there's a change to any of its device attributes later, IoT Security immediately moves it to the Identified IP Endpoints table where it continues tracking these attributes. You can see a total of all IP endpoints discovered on the network or learned from integrated third-party products and a total and a list of all identified IP endpoints on **Assets** > **Devices** > **IP Endpoints**.

At the top of the page are data filters for sites, device types, and time periods (1 Day, 1 Week, and 1 Month). The sites filter controls the data displayed for IP endpoints and identified IP endpoints per site, per site group, or for all sites. The filter for device types controls the display of data by types such as Industrial, Medical, Office, Traditional IT, All IoT, and All Devices. The time filter displays data that IoT Security discovered or learned within the past day, week, or month.

You might wonder why the device type filter affects the total number of IP endpoints. After all, IoT Security is not yet able to identify what type of device an IP endpoint is. However, for some of them, it already has an approximate idea—enough to distinguish an IT device from an IoT device, for instance. That's why you might see a different total number of IP endpoints when the filter is, say, **All Devices** and when it's **All IoT**.

To see the history of an identified IP endpoint, click its IP address. For example, the history below shows that IoT Security initially identified this IP endpoint as a Windows PC and then revised that to a Windows tablet. IoT Security maintains a history of up to 10 changes over the past 30 days.

If the behavior of an identified IP endpoint eventually settles to a consistently stable pattern again and there are no further changes to its major device attributes for seven consecutive days, IoT Security moves it back to the Devices page. You can also see the historical record of the last ten changes on its Device Details page.

The relationship between the internal database of IP endpoints, the Devices table, and Identified IP Endpoints table is shown below.

IP endpoints (internal database)	Devices table (IoT Security portal)	Identified IP endpoints (IoT Security portal)
IoT Security maintains an internal database of IP endpoints. When an IP endpoint has a stable identity for seven days, it's published to the Devices table.	The identity of an IP endpoint published to the Devices table should be stable. However, if any of its main attributes (device profile, vendor, model…) changes, IoT Security immediately moves it to the Identified IP Endpoints table.	IoT Security continues to track the identity of the IP endpoint on a daily basis. If its identity is stable for seven days, it's moved back to the Devices table.

Discover Mobile Device Attributes

IoT Security can learn mobile (cellular) device attributes, add the devices to its inventory, and track them by the IMEI numbers. You can then see various mobile device attributes for them on the **Assets** > **Devices** and **Device Details** pages. You can also use the mobile device attributes when creating custom alerts. However, because they are classified as Traditional IT, IoT Security doesn't make policy rule recommendations or send firewalls IP address-to-device mappings for mobile devices.

Set up PAN-OS to Send IoT Security Mobile Device Attributes

> *This assumes that IoT Security is already* onboarded *on your firewall, it has the* required licenses and certificates*, and logging is enabled.*

STEP 1 | Enable GTP Security on the firewall.

1. Log in to PAN-OS, select **Device** > **Setup** > **Management**, and then click **Edit** (the gear icon) for General Settings.
2. Select **GTP Security** and then click **OK**.
3. **Commit** your changes and then select **Device** > **Operations** > **Reboot Device**.

STEP 2 | Create a Log Forwarding profile that includes GTP logging.

1. Log back in and select **Objects** > **Log Forwarding** > **Add**.
2. Enter a name for the log forwarding profile like `Mobile Device Logging`, select **Enable enhanced application logging to Strata Logging Service**, and then click **OK**.

STEP 3 | Create a Mobile Network Protection profile for the types of mobile devices on the network.

The following are the recommended settings that enable the correlation of user IDs and equipment IDs to user equipment IP addresses (UEIP) for different mobile devices. For details about each setting, see the Mobile Network Protection Profile help in PAN-OS.

- **5G mobile devices with RADIUS**

    1. Select **Objects** > **Security Profiles** > **Mobile Device Protection** and then click **Add**.

    2. Enter a name for the profile such as `RADIUS Correlation`, click **Correlation**, and then enter the following:

        UEIP Correlation: (select)

        Mode: **Loose**

        User Plane with GTP-U encapsulation: (clear)

        Source: **RADIUS**

        Log At Ueip Start: (select)

        Log At Ueip End: (select)

    3. Click **GTP Inspection** > **GTP-U**, and then enter the following to perform validity checks of the Information Element (IE) in GTP headers and generate alerts if any irregularities are found:

        Alert: (select)

        Reserved IE: (select)

        Order of IE: (select)

        Length of IE: (select)

        Spare Flag in Header: (select)

        Unsupported message type: (select)

        GTP-in-GTP: **alert**

- **5G mobile devices with Packet Forwarding Control Protocol (PFCP)**

    1. Select **Objects** > **Security Profiles** > **Mobile Device Protection** and then click **Add**.

    2. Enter a name for the profile such as `PFCP-5G Correlation`, click **Correlation**, and then enter the following:

        UEIP Correlation: (select)

        Mode: **Loose**

        User Plane with GTP-U encapsulation: (clear)

        Source: **PFCP**

        Log At Ueip Start: (select)

        Log At Ueip End: (select)

    3. Click **GTP Inspection** > **GTP-U**, and then enter the following to perform validity checks of the IE in GTP headers and generate alerts if any irregularities are found:

Alert: (select)

Reserved IE: (select)

Order of IE: (select)

Length of IE: (select)

Spare Flag in Header: (select)

Unsupported message type: (select)

GTP-in-GTP: **alert**

- **3G and 4G mobile devices with GTP-C**

    1. Select **Objects** > **Security Profiles** > **Mobile Device Protection** and then click **Add**.

    2. Enter a name for the profile such as `GTP-C-3G4G Correlation`, and then enter the following in the **GTP-C** tab to use stateful inspection, perform validity checks of the IE in GTP headers, and generate alerts if irregularities are found:

        GTPv1-C

        Stateful Inspection: (select)

        Alert: (select)

        Reserved IE: (select)

        Order of IE: (select)

        Length of IE: (select)

        Spare Flag in Header: (select)

        Unsupported message type: (select)

        GTPv2-C:

        Stateful Inspection: (select)

        Alert: (select)

        Reserved IE: (select)

        Length of IE: (select)

        Spare Flag in Header: (select)

        Unsupported message type: (select)

    3. Click **GTP-U**, and then enter the following:

        Alert: (select)

        Reserved IE: (select)

        Order of IE: (select)

        Length of IE: (select)

        Spare Flag in Header: (select)

        Unsupported message type: (select)

GTP-in-GTP: **alert**

Log at GTP-U session start: (select)

Log at GTP-U session end: (select)

GTP-U Content Inspection: (select)

STEP 4 | Create Security policy rules to log mobile device traffic and forward the logs to the logging service.

Create Security policy rules to log mobile device traffic and forward logs to the logging service for IoT Security to analyze. The rules you create depend on the generation of mobile devices on the network and whether the network uses RADIUS or PFCP.

- **5G mobile devices with RADIUS**

    1. Select **Policies** > **Security** and then click **Add**.
    2. Create a universal Security policy rule with the following settings:

        Allow **radius** as the application from any source to any destination.

        In the Actions tab, choose **Profiles** as the Profile Type, choose the Mobile Network Protection profile you created previously for the RADIUS correlation, select **Log at Session Start** and **Log at Session End**, and choose the Log Forwarding profile you previously created.

        Click **OK**.

    3. Click **Add** and then create a universal Security policy rule with the following settings:

        In the Actions tab, choose **None** as the Profile Type, select **Log at Session Start** and **Log at Session End**, and choose the Log Forwarding profile you previously created.

        Allow any application from any source to any destination.

        Click **OK**.

    4. If necessary, reposition the first rule above the second in the ruleset.

- **5G mobile devices with PFCP**

    1. Select **Policies** > **Security** and then click **Add**.
    2. Create a universal Security policy rule with the following settings:

        Allow **pfcp** as the application from any source to any destination.

        In the Actions tab, choose **Profiles** as the Profile Type, choose the Mobile Network Protection profile you created previously for the PFCP 5G correlation, select **Log at Session Start** and **Log at Session End**, and choose the Log Forwarding profile you previously created.

        Click **OK**.

    3. Click **Add** and then create a universal Security policy rule with the following settings:

        Allow **gtp-u** as the application from any source to any destination.

        In the Actions tab, choose **Profiles** as the Profile Type, choose the Mobile Network Protection profile you created previously for the PFCP 5G correlation, select **Log at**

Session Start and **Log at Session End**, and choose the Log Forwarding profile you previously created.

Click **OK**.

4. Click **Add** and then create a universal Security policy rule with the following settings:

Allow any application from any source to any destination.

In the Actions tab, choose **None** as the Profile Type, select **Log at Session Start** and **Log at Session End**, and choose the Log Forwarding profile you previously created.

Click **OK**.

5. If necessary, reposition rules so that the first and second rules are above the third in the ruleset.

- **3G and 4G mobile devices with GTP-C**

  1. Select **Policies** > **Security** and then click **Add**.
  2. Create a universal Security policy rule with the following settings:

     Allow **gtpv1-c** and **gtpv2-c** as the application from any source to any destination.

     In the Actions tab, choose **Profiles** as the Profile Type, choose the Mobile Network Protection profile you created previously for the GTP-C 3G and 4G correlation, select **Log at Session Start** and **Log at Session End**, and choose the Log Forwarding profile you previously created.

     Click **OK**.

  3. Click **Add** and then create a universal Security policy rule with the following settings:

     Allow **gtp-u** as the application from any source to any destination.

     In the Actions tab, choose **Profiles** as the Profile Type, choose the Mobile Network Protection profile you created previously for the GTP-C 3G and 4G correlation, select **Log at Session Start** and **Log at Session End**, and choose the Log Forwarding profile you previously created.

     Click **OK**.

  4. Click **Add** and then create a universal Security policy rule with the following settings:

     Allow any application from any source to any destination.

     In the Actions tab, choose **None** as the Profile Type, select **Log at Session Start** and **Log at Session End**, and choose the Log Forwarding profile you previously created.

     Click **OK**.

  5. If necessary, reposition rules so that the first and second rules are above the third in the ruleset.

STEP 5 | **Commit** the configuration

View Mobile Device Attributes in IoT Security

After the firewall begins logging mobile device traffic, it forwards the traffic metadata in GTP logs to the logging service, which in turn streams it to IoT Security. To check the status of the GTP logs, log in to the IoT Security portal and select **Administration** > **Firewalls**. There you can see if

IoT Security is receiving GTP logs, the time of the latest log, and how many GTP log events and bytes it's received.

To see mobile device attributes in the device inventory on the Devices page, select **Assets > Devices**. Because the Mobile Device columns are hidden by default, click the icon with three vertical bars to open the column selection panel, and select all the columns you want to see. All the columns displaying mobile device attributes are available in the Mobile section:

- **Mobile Equipment Identity** – The 15-to-17-digit code assigned to every mobile device to uniquely identify it International Mobile Equipment Identity (IMEI)

- **Mobile Subscriber Identity** – A unique identifier issued on a Subscriber Identity Module (SIM) card. In 2G, 3G, and 4G networks, this identifier is referred to as International Mobile Subscriber Identity (IMSI). In 5G networks, it is called Subscription Permanent Identifier (SUPI).

- **Mobile Subscriber ISDN** – The Integrated Services Digital Network number is a mapping of a cellular telephone number to a mobile subscriber

- **Mobile APN** (Access Point Name) – Term used to identify the external Packet Data Network (PDN) to which mobile devices connect through the 2G, 3G, or 4G cellular network. In a 5G network, it refers to the Data Network Name (DNN).

- **Radio Access Technology** – The underlying connection method mobile devices use for wireless radio communications; for example, Bluetooth, Wi-Fi, UMTS, LTE, or 5G NR

- **Mobile Base Station Code** – The identification number that uniquely identify a cellular base station

- **Mobile Area Code** – The area code of the user's location

- **Mobile Network Code** (MNC) – A two-digit (European standard) or three-digit (North American standard) number identifying the Public Land Mobile Network (PLMN) of the mobile subscriber

- **Mobile Country Code** (MCC) – A three-digit number identifying the country of the mobile subscriber

- **Mobile TAC** (Type Allocation Code) – An eight-digit number that identifies the manufacturer of a mobile device

- **Network Slice** – The logically discrete section of network operating over a common infrastructure

- **Mobile Device** – The end user device operating on a wireless network

In addition to showing columns with these attributes in the inventory table, you can also use them in filters and queries at the top of the Devices page. They are displayed on the Device Details page of mobile devices and are available for use when creating custom alert rules.

Custom Attributes

IoT Security provides a large number of attributes for the devices it discovers and learns. A few of these are the device model, vendor, OS, VLAN ID, risk level, and location. For the full list, see the columns for the inventory table on the Devices page. When viewing the devices in your inventory, you can sort and filter by these device attributes, making it easier to find and track those of interest. However, if these attributes don't accommodate all your needs, you can create custom attributes that better align with the device attributes you use. IoT Security allows up to 50 custom attributes per tenant.

Create a Custom Attribute and Apply It Automatically

You can configure a custom attribute with a conditional statement, so that IoT Security automatically applies a value when a condition is met.

Before you start, make sure you've already created and saved one or more data filters on the **Devices** page (**Assets** > **Devices**). You will use a filter in the "IF" clause of each IF/THEN statement, indicating the condition that's required for IoT Security to apply the value in the "THEN" clause to devices.

The automatic assignment of custom attributes through the use of simple IF/THEN statements provides an efficient approach to their application. For example, when devices are managed by different departments in an organization, custom attributes can indicate which department manages which device. To do this, first create a data filter that groups together all the device profiles that a particular department manages. Then create another data filter for all the device profiles that another department manages. Continue as necessary until all the devices are divvied up by profile among the various departments that manage them. Then create a custom attribute with conditional statements that say IF a device matches <filter-1>, then apply <name of department-1> to it; IF another device matches <filter-2>, then apply <name of department-2> to it; and so on. After you're done, you can then sort the devices in your inventory on the Devices page by the departments that manage them.

STEP 1 | (Optional) Create a filter to use in the attribute.

If you don't already have a filter to use in the custom attribute, log in to the IoT Security portal and select **Devices**. Define a data filter at the top of the page and then save it.

STEP 2 | Create a custom attribute that IoT Security will apply to devices automatically.

1. Select **Settings** > **Custom Attributes** > **+** (Create Custom Attribute).

2. Enter the following in the Create Custom Attribute pop-up panel that appears:

   Attribute Name: Enter a name for the custom attribute. It cannot contain special characters and it cannot be longer than 50 characters.

   Default Value (Optional): Enter a value for IoT Security to apply by default to all the devices in your inventory. If you don't include a default value, IoT Security will enter `N/A` in the field for this attribute.

   Value Automation (Optional): **Add** an IF/THEN conditional statement to determine when a value is applied to the device attribute. Choose a previously defined filter for the **IF a device matches this filter** field and then enter a value in the **THEN apply this value to the attribute**

field. You can add more IF/THEN statements (up to five). The logical relationship between them is "or" and their order is important because IoT Security checks the conditions from the top down and will apply the value of the first match it finds.

STEP 3 | **Save** the custom attribute configuration.

IoT Security searches through its inventory for any devices that match the condition—or one of several possible conditions—in the Value Automation section and then applies the prescribed value. This search can take several minutes to complete. Going forward, IoT Security applies the value to any device whose condition matches that in the attribute configuration.

Manually Apply Custom Attribute Values to a Device

In addition to creating custom attributes that IoT Security automatically applies to devices based on specified conditions, you can create custom attributes and manually apply values to them per device yourself.

STEP 1 | Create a custom attribute whose value you will manually apply to devices.

1. Log in to the IoT Security portal and select **Settings** > **Custom Attributes** > **+** (Create Custom Attribute).

2. Enter the following in the Create Custom Attribute pop-up panel that appears:

   Attribute Name: Enter a name for the custom attribute. It cannot contain special characters and it cannot be longer than 50 characters.

   Default Value (Optional): Enter a value for IoT Security to apply by default to all the devices in your inventory. If you don't include a default value, IoT Security will enter N/A in the field for this attribute.

   Value Automation (Optional): Do not configure this section.

STEP 2 | **Save** the custom attribute configuration.

STEP 3 | Apply the custom attribute to a device.

1. Select **Assets** > **Devices** and use the search, filter, and sort tools to display the devices in the inventory to which you want to apply the attribute you just created.

2. Click the device name, which opens the Device Details page.

3. Click **Edit** next to Custom Attributes.

4. Remove any you don't want to apply to the device and edit or add any you do.

5. **Save** your configuration changes.

View Devices by Custom Attribute

After applying custom attributes to devices, you can then show custom attribute columns on the Devices page. Click the column icon (▦) and select one or more custom attributes whose columns you want to display on the page.

The selected columns appear in the inventory section of the Devices page.

To hide the column, click the column icon again and clear the check boxes for custom attributes you no longer want to see.

Edit Custom Attributes and Delete Them from Devices

To edit or delete a custom attribute, select **Settings** > **Custom Attributes**, click the three vertical dots at the far right of the a custom attribute, and then click either **Delete** or **Edit**.

Attribute Name	Default Value	Value Rule	Modified By	Modified Time	
RiskyDevice	N/A	–	jun_owner@abc.com	Jun 23, 2022, 06:24	⋮
FDA Classification	N/A	–	jun_owner@abc.com	Jun 23, 2022, 06:24	Delete
Laptop	N/A	–	jun_owner@abc.com	Jun 23, 2022, 06:25	Edit

Custom Attributes (3)

Items per page   25 ⌄   1 - 3 of 3 rows   1 ⌄   of 1 page   ‹ ›

271

Tag Management

The **Settings** > **Tag Management** page contains a list of all tags that you can apply to the devices in your inventory. There are two tabs on this page: **System Tags** with predefined system tags and **Custom Tags** with user-defined custom tags.

You can create your own custom tags and use them to add meaningful labels to your devices. IoT Security creates system-defined tags based on the types of devices detected in your environment. If manufacturing devices are found, for example, then it creates system tags for Purdue levels 1 through 5.

Tag Type	Tag Value	Tag Rule	Tagged Devices	Create Date	
Owner	Default1	test save 7	7	11:11, March 08, 2023	⋮
Aruba ClearPass	In Scope	risk and 2 more	100,002	15:36, January 03, 2021	⋮
Cisco ISE	In Scope	Risk: Critical	7	15:36, January 03, 2021	⋮
Cisco ISE with pxGrid	In Scope	Risk: High 1 and 1 more	177	15:36, January 03, 2021	⋮
Forescout	In Scope			15:36, January 03, 2021	⋮

Follow these procedures to manage device tags:

- Create a Custom Tag and Apply It Automatically
- Manually Apply Tags to One or More Devices
- Manually Apply Tags to an Individual Device
- Remove Tags from Devices

272

Create a Custom Tag and Apply It Automatically

STEP 1 | Define a tag.

To create a custom tag, click the **+** icon in the upper right corner of the Custom Tags tab.

The Create custom tag window opens with fields for a tag type and a tag value. The type is optional and the value is required.

Optionally select or create a tag type, define a tag value, and then click **Next: Tag Rule**.

STEP 2 | Optionally define a tag rule to apply the tag automatically.

A tag rule defines a condition for applying a tag. When a device matches the filter or filters in a tag rule, IoT Security automatically applies the specified tag. IoT Security not only does this when you initially define a tag rule but it also applies the tag if it later finds new matching devices in the future. Conversely, if a device no longer match the filter, IoT Security automatically removes the tag from it.

If you want to apply the tag when a device matches a filter, choose a previously saved filter from the list. You can also add one or more filters to apply the tag to more devices. If there are multiple filters, IoT Security applies a tag to a device if it matches any one of them.

If you want to apply the tag manually to one or more devices on the Devices page or individually on the Device Details page instead of automatically through a tag rule, don't select or add any filters.

When done, **Save** the tag.

> *You can create a maximum of 1000 unique tags and manually apply them to a maximum of 100,000 devices. A single device can have a maximum of 100 tags applied to it.*

Manually Apply Tags to One or More Devices

There are two ways to apply tags to devices:

- Manually apply tags to one or more devices on the **Devices** page or to an individual device on the **Device Details** page
- Automatically apply tags through the use of special filters called tag rules on the **Settings** > **Tag Management** page

The quickest way to tag your devices is to do so manually through the device inventory on the **Devices** page.

STEP 1 | Filter the devices to tag.

Open the **Assets** > **Devices** page and use the filter tool to refine the devices in the list.

After the correct devices are listed, click the **Tag** icon (🏷) to tag the filtered devices. In fact, you're not only tagging this set of filtered devices but you're tagging the filter itself. If IoT Security detects devices matching this filter in the future, it will tag them as well.

STEP 2 | Confirm the filtering parameters before applying tags.

Check that the filters are the ones you want to use. If not, **Cancel** and modify the filters before tagging the devices.

To include the global filters for site and device type at the top of the Devices page, select **Include global filters for site and device type in this filter**. Clear the check box to exclude the site and device type global filters.

When done, click **Next: Select tag**.

STEP 3 | Select one or more tags and apply them.

The Apply tags to this filter window opens with fields for a tag type and a tag value. The type is optional and the value is required.

Optionally select or create a tag type, and select or create a tag value.

To apply more than one tag, click **+ Add Tag**.

When done, **Save and Apply**.

IoT Security tags the filtered devices, and if it detects new devices that match your filters in the future, it will automatically tag them as well. Likewise, if any tagged devices no longer match filters, IoT Security will automatically remove tags from them.

The initial tagging process can take a few moments to complete, depending on how many devices IoT Security must tag.

Manually Apply Tags to an Individual Device

In addition to tagging devices on the Devices page, you can tag an individual device from its Device Details page.

STEP 1 | Open the Manage Tags window for an individual device.

From the Devices page, click a device name to open the Device Details page for this device.

Click the Action menu icon (⋮) in the upper right of the page and then click **Manage Tags**.

STEP 2 | Apply one or more tags to the device.

Optionally select or create a tag type and select or create a tag value.

To apply additional tags, click **+ Add Tag**. You can apply a maximum of 100 tags to a single device.

When done, **Save**.

Apply tags to b4:47:5e:b1:db:03

Division	123123	✕
Enter tag type: Purdue	Enter tag value *: Level0	✕
Enter tag type: division	Enter tag value *: West	✕
Enter tag type: Follow Up	Enter tag value *: Q2 2021	✕

+ Add Tag

Cancel   Save

Remove Tags from Devices

It's possible to remove tags from individual devices and from the IoT Security system.

To remove a manually applied tag from an individual device:

1. Navigate to the Device Details page.
2. Click the Action menu icon (⋮) and then click **Manage Tags**.
3. Click the **X** next to a tag entry to remove it and then **Save**.

> *You can only remove manually applied tags from an individual device because IoT Security would reassign any removed tags that are automatically assigned as a result of tag rules. To remove automatically assigned tags, you must remove them completely from the system.*

To remove a tag from the entire IoT Security system:

1. Select **Settings** > **Tag Management**.
2. Click the Action menu icon (⋮) in the far right column and then click **Delete Tag**.

When you delete a tag on the Tag Management page, IoT Security removes it from all devices. This operation cannot be undone, so remove tags with caution.

278

Discover IoT Device Applications

IoT Security uses machine learning to discover the applications that IoT devices on your network use.

- IoT Device Applications Discovery

IoT Device Applications Discovery

Knowing which applications your network-connected IoT devices use and how many devices use them can prove useful, especially when defending against a potential threat. For example, if you know a widely used application was recently compromised, you can check which devices use it and respond in proportion to how critical the application is. If it's non-essential for business, you can create policy recommendations for firewalls to block that application. If it is essential and there is a new version, you can assign operations the task to upgrade all devices that use it. And if it is essential and there isn't a new version yet, segment all devices that use it and restrict access to them only to people and resources that are necessary for them to function. Having visibility into the applications on your network allows you to take swift action to safeguard your assets when danger threatens.

On the **Networks** > **Applications** page, IoT Security displays all the applications that have been spotted in use by the IoT devices on your network.

	Application	App Risk ⓘ	Number of Devices ↓	Category	Subcategory	Technology	Profiles	Evasive
☐	dns-base	3	2,049	networking	infrastructure	network-protocol	Advantech Industrial PC a...	No
☐	ssl	4	1,754	networking	encrypted-tunnel	browser-based	Advantech Industrial PC a...	No
☐	ntp-base	2	1,707	networking	infrastructure	network-protocol	Advantech Industrial PC a...	No
☐	ping	2	1,644	general-internet	internet-utility	network-protocol	Advantech Industrial PC a...	No
☐	ssh	4	1,370	networking	encrypted-tunnel	client-server	APC Smart PowerSupply a...	No
☐	snmp-base	2	1,343	networking	infrastructure	client-server	AIC Device and 30 more	No
☐	traceroute	2	1,265	general-internet	internet-utility	network-protocol	Advantech Industrial PC a...	No
☐	lpd	3	1,221	business-systems	management	client-server	Arista Network Switch and...	No
☐	unknown-tcp	1	1,211	unknown			AIC Device and 23 more	No
☐	web-browsing	4	1,050	general-internet	internet-utility	browser-based	Advantech Industrial PC a...	No
☐	paloalto-updates	2	729	business-systems	software-update	client-server	Aruba UXI Sensor and 6 m...	No
☐	snmpv3	1	665	networking	infrastructure	client-server	APC Smart PowerSupply a...	No
☐	pan-db-cloud	1	563	business-systems	general-business	client-server	Aruba UXI Sensor and 6 m...	No
☐	gnutella	5	525	general-internet	file-sharing	peer-to-peer	Arista Networks Device an...	Yes
☐	paloalto-dns-securi...	1	505	business-systems	general-business	client-server	DTEN Display Board PC M...	No
☐	snmpv2	2	474	networking	infrastructure	client-server	APC Smart PowerSupply a...	No
☐	dhcp	2	444	networking	infrastructure	network-protocol	APC Smart PowerSupply a...	No
				general-internet			...Industrial PC a...	No

The Applications page shows the total number of unique applications detected for IoT devices matching the site and time-range filters set at the top of the page.

> 📋 *The IoT Security portal disregards the device-type filter on this page and always shows applications for "All IoT" devices, as indicated by the blue icon at the top of the page.*

281

Although IoT Security displays devices and networks as soon as it discovers and identifies them, it collects data about detected applications over the course of a day and then compiles a list. It then displays that list on the Applications page until it compiles the next daily list of applications detected on the network. When you start using IoT Security, you might notice that it begins showing data on the Devices and Networks page before showing anything on the Applications page. This can happen because IoT Security hasn't generated a list of applications yet. After it does, it will continue doing that every day thereafter.

If you set the time-range filter for **1 Day**, **1 Week**, or **1 Month**, the Applications page shows numbers for the time range you set. However, because IoT Security organizes the applications it detects into daily lists, the time-range filter for **1 Hour** shows the same set of unique applications as **1 Day**, which is the smallest list of applications you can see. In addition, IoT Security doesn't maintain application details for more than a month. Therefore, the time-range filter for **1 Year** shows the same set of unique applications as **1 Month**, which is the largest list of applications you can see.

IoT Security provides data from Applipedia about each of the applications it monitors. When a new application appears, you can use this data to determine if it's expected or not and also to see the level of risk it introduces to your network. For example, the following shows the application description, characteristics, and security information that IoT Security retrieves from Applipedia for DNS:

Here's the same information about DNS presented in Applipedia:

The following summarizes the different characteristics and types of security information that IoT Security retrieves from Applipedia and displays for each application.

Application Characteristics	
Category	A broad application type to which an individual application belongs
Subcategory	A more specific application type for an individual application
Risk Level	The level of risk that's inherent in an application as determined by the characteristics listed in the next table, on a scale of increasing risk from 1 to 5
Standard Ports	The protocol and standard service port numbers that the application uses
Technology	How an application functions: network-protocol, client-server, peer-to-peer, or browser-based

Application Security Information	
Evasive	Yes = The application uses a port or protocol for something other than its originally intended purpose with the intention of evading firewall policy enforcement.
Excessive Bandwidth	Yes = The application consumes at least 1 Mbps on a regular basis through normal use.
Prone to Misuse	Yes = The application is often used for nefarious purposes or is easily set up to expose more than the user intended.
Capable of File Transfer	Yes = The application has the capability to transfer a file from one system to another over a network.
Tunnels Other Applications	Yes = The application can transport other applications inside its protocol.
Used by Malware	Yes = Malware has been known to use the application for propagation, attack, or data theft, or the application has been distributed with malware.
Has Known Vulnerabilities	Yes = The application has at least one publicly reported vulnerability. (Web-based

Application Security Information	
	applications are always set to Yes because HTTP always has vulnerabilities.)
Widely Used	Yes = The application likely has more than 1,000,000 users.
SaaS	Yes = The application is cloud based and provided through Software as a Service (SaaS). No = The application is hosted on premises.

Many of these explanations come from the KB article "How to Determine Risk Level of Application, Spyware, and Anti-Virus". *There you can read more about the information that Applipedia provides and how risk scores are calculated.*

To see data from Applipedia about applications on the Applications page, either click or hover your cursor over an application name to view a pop-up with information about the application taken directly from Applipedia.

In addition, use the column picker to show information from Applipedia in columns on the Applications page.

Click a number in the Number of Devices column to open the Devices page with a filter applied to show only devices that use the corresponding application.

Clicking or hovering your cursor over the blue text of an entry in the Profiles column displays a list of all profiles that use that application.

Discover IoT Device Applications

	Application	App Risk ⓘ	Number of Devices ↓	Category	Subcategory	Technology		Evasive
☐	dns-base	3	2,049	networking	infrastructure	network-protoc	a...	No
☐	ssl	4	1,754	networking	encrypted-tunnel	browser-based	a...	No
☐	ntp-base	2	1,707	networking	infrastructure	network-protoc	a...	No
☐	ping	2	1,644	general-internet	internet-utility	network-protoc	a...	No
☐	ssh	4	1,370	networking	encrypted-tunnel	client-server	a...	No
☐	snmp-base	2	1,343	networking	infrastructure	client-server		No
☐	traceroute	2	1,265	general-internet	internet-utility	network-protoc	a...	No
☐	lpd	3	1,221	business-systems	management	client-server	Arista Network Switch and...	No
☐	unknown-tcp	1	1,211	unknown			AIC Device and 23 more	No
☐	web-browsing	4	1,050	general-internet	internet-utility	browser-based	Advantech Industrial PC a...	No
☐	paloalto-updates	2	729	business-systems	software-update	client-server	Aruba UXI Sensor and 6 m...	No
				networking			...rt PowerSupply a...	No

Profiles: ssl (52)

Advantech Industrial PC
Amazon Device
Arista Network Switch
Arista Networks Device
Aruba Instant AccessPoint
Aruba Networks Device
Aruba UXI Sensor
Axis Communications Video Surveillance
Cisco Meraki Device
Cisco Systems Device
View all

Detect IoT Device Vulnerabilities

IoT Security uses machine learning to detect vulnerabilities and assess risk. It bases its detection and assessment on the network traffic behaviors of IoT devices and dynamically updated threat feeds.

- IoT Device Vulnerability Detection
- Vulnerability Overview Dashboard
- Vulnerabilities Page
- Vulnerability Details Page
- IoT Risk Assessment

IoT Device Vulnerability Detection

A vulnerability refers to an intrinsic flaw built into the software or hardware of a device that is often well-known and can be exploited in some way. A risk, on the other hand, considers environmental, configuration, behavioral, and security policy-related factors in addition to one or more underlying vulnerabilities. This distinction is important because some risks appear in the device details page but not on the Vulnerabilities page, and yet they can influence the severity level that IoT Security assigns to a vulnerability.

IoT Security considers a vulnerability to be potential when it applies to a specific device type, model, and version number and one or more devices match the specified device type but their model and/or version number are unknown. Similarly, a device is considered to be potentially vulnerable for the same reason.

A vulnerability can also be considered potential if it only applies to devices with certain serial numbers and there are devices whose serial numbers are unknown but match the vulnerability description in all other regards.

> *The IoT Security app detects vulnerabilities for IoT devices only. It does not provide vulnerability detection, alerts, policy recommendations, and network behavior analysis for IT devices. For IT devices, the IoT Security app provides device identification only.*

Vulnerability Overview Dashboard

The Vulnerability Overview dashboard (**Vulnerabilities** > **Vulnerability Overview**) lets you customize how information about vulnerabilities and vulnerability instances is presented so you can view their impact on your devices from different perspectives. By setting filters, you determine the scope of the information displayed, and by defining queries and settings, you control the types of vulnerabilities and the types of devices you want to see.

> *The filters you set at the top of the page do not affect the Vulnerabilities of Interest section. The vulnerabilities displayed there are determined by the settings you configure within that section itself.*

The dashboard consists of four main sections to help you easily see key statistics, identify top vulnerabilities of interest, gain insights into their distribution among different groups of devices, and track vulnerability instances trends.

Detect IoT Device Vulnerabilities

At the top of the page is a summary of key vulnerability statistics within the parameters defined by the filters for sites, device category, and time range.

- **Vulnerabilities to Date** – This is the total number of vulnerabilities detected since you began using your IoT Security tenant.

  Although IoT Security retains security alerts in its database for up to one year, it does not impose this time limit on vulnerabilities. If you've been using IoT Security longer than a year, it will continue showing vulnerabilities detected more than a year before.

- **New Vulnerabilities in** <time range> – This is the total of vulnerabilities that were detected within the time range specified in the data filter at the top of the page.

- **Top Priority Vulnerabilities in** <time range> – This is the total number of vulnerabilities that IoT Security prioritized as "Top". (There are also "Medium" and "Low" priorities.) It's followed by the number of instances for these vulnerabilities and the number of critical assets they affect. If you click one of the links here, IoT Security opens **All Vulnerabilities** with filters applied to show only top-priority vulnerabilities within the site, device type, and time range set on **Vulnerabilities Dashboard**.

- **Aged Vulnerabilities** – This is the total of all vulnerabilities that remain unresolved beyond the specified time range (30, 60, 90, or 180 days).

- **Instances I resolved in** <time range> – This is the total of vulnerability instances that were assigned to the person currently logged in and which were resolved during the time range specified in the data filter at the top of the page.

Vulnerabilities of Interest – Define criteria for vulnerabilities that matter most to you. IoT Security will then display the top ten vulnerabilities in response to your query with the most severe CVSS scores and those affecting the most device profiles displayed first. For example, if you want to see vulnerabilities for a specific vendor or profile that were detected within the last week, click the gear icon (⚙) and configure a query to show the vulnerabilities that interest you. IoT Security then displays the ten most severe vulnerabilities with the broadest impact that match your terms.

By default, IoT Security uses the predefined "Risky Vulnerabilities" query to search for confirmed critical vulnerabilities for which a proof of concept (PoC) is publicly available. You can edit this query to define other attributes of interest and then click the bookmark icon (🔖) to save it for reuse.

Vulnerability Instances Distribution – The Sankey chart lets you see the distribution of vulnerability instances across different groupings of devices. Reading the chart from left to right, you start off on the left with all the vulnerability instances that match the site and device category filters at the top of the page. (Regardless of the time range filter set for the page, this chart shows all vulnerability instances to date.) The chart then relates these instances to a type of grouping in the middle and relates these again to another type of grouping on the right. The choices for these groupings are **Severity**, **Vulnerability Type**, **Status**, **Device Type**, **Device Category**, **Profile**, **Vendor**, **Exploit Status**, **Attack Vector** (the type of access required to exploit a vulnerability, as defined in a CVE), and **Vulnerability Priority** (Top, Medium, Low). Vulnerability instances are distributed vertically in the chart by severity (when Severity is the chosen grouping), priority (when Vulnerability Priority is chosen), or by instance count (for all other types). Those groupings with the highest severity, highest priority, or most instances at the top of the chart. When there are more than five groupings, the Sankey chart shows the top five and then gathers everything else in an "Others" group. Hover your cursor over **Others** to see a list of the next ten groupings, and click **View all** to see a pop-up panel with a complete list.

When you use **Profile** to group instances and then hover your cursor over an area on a post for a particular profile, IoT Security displays an Action pop-up panel that lets you create a set of recommended policy rules with this profile as the source.

When you click **Create Policy**, IoT Security opens **Assets** > **Devices>** *profile-name* **> Create New Policy Set**. From there, you can modify the automatically generated policy set if necessary, save it, and then activate it for firewalls to import.

For example, to see the ratio of vulnerability instances among different device profiles and different vulnerability types, choose **Profile** for the middle post and **Vulnerability Type** for the right post. The gray bands between the left and middle posts show how many instances pertain to each of the top five device profiles, and the gray bands between the middle and right posts show how many instances in each profile belong to different vulnerability types. Each band is labeled and shows the total number of vulnerability instances per profile (on the left) and for that profile per vulnerability type (on the right). The width of the bands lets you see at a glance the relative quantities of vulnerability instances. Hovering your cursor over a section of a post shows the percent of instances for the adjacent bands.

> *Colors only convey meaning to denote vulnerability severity levels: red = critical, orange = high, yellow = medium, and blue = low. For other types of groupings, semi-transparent shades of gray are used solely to distinguish one band from another.*

To download the data from the Sankey chart for your records or reports, click the download icon (⬇) in the upper right above the chart. IoT Security saves it as an .xlsx file with vulnerability instance distribution information on the first sheet and a complete list of vulnerability instances on the second.

Vulnerability Instances Trend – The Instance Trend chart displays a cumulative count of vulnerability instances over the specified time period and a daily noncumulative count of resolved instances. This visually shows vulnerability instance trends to help vulnerability management teams see if the number of vulnerability instances has been increasing or decreasing over time. You can view data presented either by vulnerability priority (Top, Medium, Low priorities) or CVSS score (Critical, High, Medium, Low). Use the toggle on the right above the chart to switch between the two views. When using the CVSS score view, the chart also displays data for resolved vulnerability instances, which can help teams gauge their progress in regard to vulnerability resolution. Hover your cursor over different points on the chart to see the number of vulnerability instances with different priorities or CVSS scores for different dates.

To download data from the Instance Trend chart for reports or records, click the download icon (⬇) in the upper right above the chart. IoT Security saves it as an .xlsx file with the number of vulnerability instances to date and resolved instances over the specified period of time.

Vulnerabilities Page

The Vulnerabilities page (**Vulnerabilities** > **Vulnerability Overview** > **All Vulnerabilities**) lists the vulnerabilities that IoT Security has detected or learned about through a third-party integration .

You can search for a text string in any of the columns, download the list of vulnerabilities, create a filter to show only the vulnerabilities you want to see, and control which columns you want to show and hide.

📋 *Although the Severity column in the table shows only icons, you can still search by the severity level words Critical, High, Medium, and Low.*

You can also set the number of rows you want to see on each page (from 5 to 200) and navigate among multiple pages.

CVSS Score Range	Severity Level
9.0 – 10.0	Critical
7.0 – 8.9	High
4.0 – 6.9	Medium
< 4.0	Low

While a severity level in the IoT Security system reflects a Common Vulnerability Scoring System (CVSS) score, there isn't always a direct correlation between the two. For example, a hard-coded password in a device might have a CVSS score of 10.0, but an IoT Security severity level of High rather than Critical. This can happen when there isn't proof that the device can be accessed from the Internet or by an unauthorized user. While the National Institute of Standards and Technology (NIST) assigns a CVSS score to a vulnerability generically, IoT Security assigns a "risk severity" level to vulnerabilities based on the specifics of each case.

For example, although the first vulnerability has a CVSS score of 9.8, its risk severity is High instead of Critical. IoT Security bases the severity level not only on the CVSS score but on other determining risk factors as well.

The Vulnerabilities table columns are organized into five categories: Risk, Basic, Vulnerability Metrics, Threat Metrics, and Impact Metrics. While the Risk and Basic categories each contain a single column, the three metrics categories each contain a group of columns. You can click-drag columns to rearrange them within their respective groups or click-drag the groups to rearrange their order on the table. However, you can't click-drag columns outside of their groups while grouping is enabled. To disable grouping, click the three vertical dots icon above the table on the right and click **Ungroup columns**. With the columns ungrouped, you can reposition them so that they mingle with columns that were previously separated into other groups.

As with other tables in the IoT Security portal, you can control which columns are shown. Click the three vertical dots above the table on the right, click **Edit columns**, and then select the columns you want to see and clear the ones you want to hide.

Risk – The risk is a sort of ranking of the potential danger a vulnerability poses. It's the result of various factors that, when combined, help you prioritize which vulnerabilities to watch and address.

- **Priority** – IoT Security determines a ranking of Top, Medium, and Low by weighing various factors that calculate the likelihood of an attack and the impact it would pose on your

resources. Hover your cursor over a priority to see a summary of the indicators of risk that contribute to its ranking.

When **Processing...** *is displayed here, it indicates that IoT Security is still determining the priority of a vulnerability. Because IoT Security runs a service to determine priority on a daily basis, it can take up to 24 hours to determine the priority of a device.*

IoT Security automatically assigns a high asset criticality level to industrial and medical devices and a medium level to all other devices by default. It does this through the system-defined Asset Criticality attribute, which you can see in **Settings** > **Custom Attributes**. You can also define filters on **Assets** > **Devices** and add them to the Asset Criticality attribute to assign different asset criticality levels to devices based on attributes such as device category, profile, or vendor. For example, you might first define a filter for patient monitor profiles on **Assets** > **Devices** and then add a rule to the system-defined Asset Criticality attribute, the rule stating

that if a device matches the filter for patient monitor profiles, then IoT Security will apply an asset criticality level of Critical to it.

> You can also edit asset criticality for an individual device on its **Device Details** page.
> Click **Edit** in the Custom Attributes section and change the Asset Criticality field to the level you want.

Basic – This is the name of a vulnerability.

- **Vulnerability Name** – The name or Common Vulnerabilities and Exposures (CVE) number of a vulnerability. This links to the Vulnerability Details page.

Vulnerability Metrics – These metrics are about vulnerabilities and the attacks that exploit them

- **Severity** – The severity level of a vulnerability: critical, high, medium, or low.
- **CVSS** – The CVSS (Common Vulnerability Scoring System) score of a vulnerability.
- **Vulnerability Type** – (Not shown by default) This identifies the type of vulnerability, such as Info Leak, Overflow, or Code Execution.
- **Vulnerability Source** – (Not shown by default) The source that identified the device vulnerability: IoT Security, a third-party integration (Rapid7, Qualys, Tenable), or IoT Security Device Software Library.

- **Attack Vector** – (Not shown by default) Also referred to as "Access Vector", this is the type of access an attacker must have to exploit a vulnerability. The metric values are defined in the CVE. The vulnerability score increases as the possible distance from the target increases:

    - **Physical** – An attacker must physically touch or control the vulnerable device.
    - **Local** – An attacker must launch an exploit locally or use social engineering to dupe a user into helping launch it.
    - **Adjacent** – An attacker must have access to the same physical or logical network as that of the vulnerable device.
    - **Network** – An attacker can launch an exploit remotely from anywhere on a network that can access the vulnerable device.

        When an attack vector is not defined, it's classified as "Unknown".

- **Attack Complexity** – (Not shown by default) Indicates the level of complexity required to exploit a vulnerability "Low" or "High".
- **Privilege Required** – (Not shown by default) Indicates the level of administrative privilege necessary to execute an exploit on the vulnerability, which can be "None", "Low", or "High".
- **User Interaction** – (Not shown by default) Whether a user, other than the threat actor, must participate in exploiting a vulnerability in some way. The values shown here are "None" or "Required".
- **Confidentiality Impact** – (Not shown by default) Whether sensitive information would be accessible to an attacker exploiting the vulnerability and the degree of sensitivity that might be disclosed. The values are "None", "Low", and "High".
- **Integrity Impact** – (Not shown by default) Whether protected information might be altered in any way. The values are "None" (no loss of data integrity), "Low" (a small amount of data can be modified), and "High" (any or all data can be modified).
- **Availability Impact** – (Not shown by default) Whether an exploit of the vulnerability makes data or devices inaccessible. The values are "None" (no loss of availability), "Low" (either poor performance or occasional loss of accessibility), and "High" (complete loss of accessibility).

Threat Metrics – These metrics focus on the threat that vulnerabilities pose to the security of your network and the devices on it.

- **EPSS** – The Exploit Prediction Scoring System (EPSS) provides a daily estimate of the probability that a vulnerability will be exploited within the next 30 days. To learn more about EPSS, see the EPSS Model.
- **Exploit Status** – (Not shown by default) Shows if a proof of concept has been identified for the vulnerability ("POC") or not ("Unknown").
- **APTs** – (Not shown by default) Shows if there have been any known exploits of a vulnerability for Advanced Persistent Threats (APTs). The values are "Yes" or "No".
- **Covered by Threat Prevention** – (Not shown by default) Indicates if a vulnerability is covered by the Palo Alto Networks Threat Prevention application ("Yes") or not ("No").

Impact Metrics – These metrics provide insight into how extensive and severe the impact would be of an exploited vulnerability.

- **Confirmed** – Indicates if a vulnerability is confirmed to apply to one or more devices. An empty field indicates that it is a potential vulnerability.

- **Confirmed Instances** – The number of devices to which a vulnerability is confirmed to be applicable. This number links to the Vulnerability Details page.

- **Potential Instances** – The number of devices to which a vulnerability might be applicable but has not been confirmed. This number also links to the Vulnerability Details page.

- **Addressed Instances** – (Not shown by default) The number of instances of the vulnerability that have been addressed.

- **Critical Assets Impacted** – The number of assets categorized as critical that the vulnerability impacts.

- **Vulnerable Profiles** – The number of device profiles to which a confirmed or potential vulnerability applies.

When you hover your cursor over an entry in the Vulnerability column, a panel pops up with showing its description and impact.

Clicking the name of a vulnerability entry opens the Vulnerability Details page.

Vulnerability Details Page

Clicking the CVE (Common Vulnerabilities and Exposures) link in the Vulnerability column or a number in the Confirmed Instances or Potential Instances column on the Vulnerabilities page opens the Vulnerability Details page for that vulnerability (**Vulnerabilities** > **Vulnerability Overview** > **Vulnerabilitiy Details**). Here you can read a description of the vulnerability, see details about it, and learn which device profiles it affects. You can also see which devices the vulnerability affects or potentially affects.

At the top of the Vulnerability Details page are several important attributes:

- The CVE ID links to a page about the vulnerability in the National Institute of Standards and Technology (NIST) database. For example, clicking CVE-2022-4436 opens https://nvd.nist.gov/vuln/detail/CVE-2022-4436.

CVE-2022-4436 ☑   🅅 High (v3)

- The CVSS (Common Vulnerability Scoring System) score ranks the vulnerability on a scale of 0-10, where 0 is the least severe and 10 is the most.

- The IoT Security rating system, which is based on the CVSS, categorizes a vulnerability score into one of several severity levels. There are two CVSS versions and both are presented:

Severity	CVSS (v2)	CVSS (v3)
Critical	—	9.0 - 10.0
High	7.0 - 10.0	7.0 - 8.9
Medium	4.0 - 6.9	4.0 - 6.9
Low	0.0 - 3.9	0.1 - 3.9
None	—	0.0

Next is a section describing what the vulnerability is, how it was detected, and the source of its discovery. It also explains the impact the vulnerability can cause if exploited and recommended actions you can take to remediate it. Finally, there is a chart that shows the total number of affected devices grouped by profile and the relative sizes of each group.

Summary

- The **Description** summarizes the vulnerability.

- The **Impact** section explains how attackers can exploit the vulnerability and the threat it poses.

- **Detection Reasons** explain how confirmed vulnerability instances were detected. When you click **View Details**, a panel appears over the right side of the page showing each detection reason, the device profiles it applies to, and the number of vulnerability instances that were detected for different profiles. (Reasons for the detection of potential vulnerabilities are not shown.)

- The **Vulnerability Type** identifies the category of the vulnerability, such as code execution, info leak, overflow, and denial of service.

- The **Vulnerability Source** identifies where the vulnerability was detected. One source of detection is IoT Security when it's based on device attributes such as firmware, model, and OS. Another source is IoT Security Device Software Library when the detection is based on the software and applications running on a device. Yet another source is one of the third-party vulnerability scanners with which IoT Security integrates: Qualys, Rapid7, or Tenable.

- IoT Security lists any identified software patches that can remediate the vulnerability.

  > *We recommend that you don't apply patch updates identified by IoT Security to your devices until your security or vulnerability management team or the product vendor has qualified them to ensure there are not any unexpected results or side effects.*

- On the right side of the Summary section is a list of recommendations. It typically includes various options you can take to reduce the risk that the vulnerability poses or even remediate the issue.

Vulnerability Metrics – In this section, you can see CVE submetric scores, which provide additional insight to the vulnerability severity level to help you prioritize remediation efforts. For example, vulnerabilities that can be exploited remotely might require a more urgent response than other vulnerabilities, even if the others have a higher CVSS score.

- The exploitability metrics include the attack vector (Network, Adjacent, Local, or Physical), attack complexity (High or Low), what privileges are required to launch an attack, and whether human action—other than that of the attacker—is required during the exploit.

- The impact metrics indicate what areas an exploit might affect—confidentiality, integrity, and availability—and what the impact is in these areas—none, low, or high.

- The scope metric indicates if the effects of an exploited vulnerability are limited to the impacted component (Unchanged) or they can extend to other components as well (Changed).

Threat & Compensating Metrics – In this section, you can see information about the likelihood that the vulnerability will be exploited, the types of exploits that are known to have occurred, and if there are options to compensate the threat through the Palo Alto Networks Threat Prevention application.

- The Exploit Prediction Scoring System (EPSS) percentile is a daily estimate of the probability that the vulnerability will be exploited within the next 30 days. To learn more about EPSS, see the EPSS Model.

- The exploit status can be one of the following:

  - **Unknown** – There is no known or weaponized malware exploiting this vulnerability.

  - **POC** – There's known code to exploit the vulnerability to demonstrate a security weakness.

  - **Weaponized** – There's a known exploit that is malicious or works consistently against targets.

  - **Exploited in the Wild** – An exploit of the vulnerability has been publicly reported in the wild, either by threat actors or in the Known Exploited Vulnerability (KEV) catalog.

- By clicking **View Details** for Exploits Identified, you can see a list of known POC and Weaponized exploits (but not any whose status is Unknown or Exploited in the Wild). For each one, there is a URL (source) where you can learn more, the exploit status, and the date that the exploit was published.

- Advanced Persistent Threat (APT) indicates if any exploits are known to have been used by an APT. By clicking **View Details**, you can see a list of APTs. For each one, there's the name of the APT, a description about it, the countries they've targeted, the known CVEs they've exploited, and the tactics and techniques they've employed.

- Threat prevention coverage indicates if a vulnerability is covered by the Palo Alto Networks Threat Prevention application or not. By clicking **View Details**, you can see the name of the vulnerability, its unique threat ID number, the minimum PAN-OS version that supports it, the dates of its first release and latest update, and a URL (reference) where you can learn more.

Impact View – In this section, you can see the number of devices that the vulnerability affects and their various levels of criticality: critical, high, medium, low. The level of criticality helps you assess the level of impact your organization would incur if they were compromised.

- **Asset Criticality** – In the Asset Criticality tab, a chart and accompanying table show the total number of assets (instances) affected by the vulnerability and the number and percent of affected assets at each level of criticality. The chart provides a visual representation of the data contained numerically in the table. By selecting and clearing the Critical, High, Medium, and Low check boxes in the table, you can show and hide the corresponding segments in the chart.

- **Confirmed** – In the Confirmed tab, a chart and table show the total number of assets that are confirmed to be vulnerable and those that are potentially vulnerable but not yet confirmed. In addition to the total, they also show the percent of assets that are confirmed and unconfirmed to be vulnerable. You can select and clear the check box for each row to show or hide the corresponding segment in the chart.

- **Profiles** – In the Profiles tab, a chart shows the total number of affected devices grouped by profile and the relative sizes of each group. When you hover your cursor over a section in the chart, a pop-up appears identifying that profile and the number of devices in it. This is particularly helpful when a vulnerability affects numerous device profiles.

At the bottom of the Vulnerability Details page are two tabs—Active Instances and Addressed Instances. On each tab a table shows all vulnerable and potentially vulnerable devices, which are referred to as instances. Here's an example to clarify the difference between these two types of devices. If a vulnerability only affects devices running a specific software version and IoT Security identifies the version running on one device as having this vulnerability but it can't identify which software version is on another, then the first device is considered as having a confirmed vulnerability but not the second one. (If **Yes** appears in the Confirmed column, a device is confirmed as vulnerable. If the Confirmed column is empty, a device is potentially vulnerable but it's not confirmed.)

A vulnerability instance initially appears in the Active Instances tab.

	Instance	Status	Confirmed ↓	IP Address	MAC Address	Site
☐	Polycom_64167f0a4...	Detected ⌄		10.50.147.90	64:16:7f:0a:46:65	test-1117
☐	Polycom_64167f619...	Investigating ⌄		10.50.146.55	64:16:7f:61:9c:40	test-1117
☐	Polycom_64167f03d...	Detected ⌄		10.50.147.33	64:16:7f:03:d6:fb	test-1117
					03:c0:14	test-1117

Active Instances(302)   Addressed Instances(3)

As soon as you change the status of a vulnerability instance to **Resolved**, IoT Security moves it from the Active Instances tab to the Addressed Instances tab.

Active Instances(302)   Addressed Instances(3)

	Instance	Status	Confirmed ↓	IP Address	MAC Address	Site	Detected Time	Vulnerability R...
☐	Polycom_64167f9f73...	Resolved ⌄		10.50.147.110	64:16:7f:9f:73:24	test-1117	Dec 5, 2019, 15:59	Resolved
☐	Polycom_64167f9f71...	Resolved ⌄		10.50.147.171	64:16:7f:9f:71:c4	test-1117	Dec 5, 2019, 15:59	Resolved
☐	Polycom_64167f619...	Resolved ⌄		10.55.20.251	64:16:7f:61:9a:e7	test1	Jan 24, 2020, 15:5	Resolved

If you later change a resolved instance to **Detected**, it's automatically moved back to the Active Instances tab.

To see more information about a device, click the device name in the Instance column to open the Device Details page for it in a new browser window or tab.

The status of a vulnerability instance begins in the Detected state. You can leave it there or set it to a different state to reflect where it's in the remediation process:

- **Detected**: This is the state of a newly detected vulnerability instance. It makes sense to keep it in this state if no action has been taken to investigate, remediate, or resolve it.

- **Investigating**: Consider setting a vulnerability instance in this state after preliminary work on it has started and it's being verified, researched, and its impact analyzed.

- **Remediating**: Consider setting an instance in this state while action is being taken to remediate it but has not yet completed.

- **Resolved**: An instance becomes resolved either by mitigating the issue or by ignoring and accepting it.

To change the state of a vulnerability instance, click the entry in the Status column and choose another state.

When you resolve a vulnerability instance, IoT Security prompts you to provide a reason for its resolution.

To assign a vulnerability instance to someone to work on, select the check box for the instance, and then click **More** > **Assign**. Enter the username or email address of a user and then click **Assign**.

The person to whom you assign a vulnerability instance must have an IoT Security user account so that it can send a message to the appropriate email address.

The user then receives an email message that states that a vulnerability was assigned to him or her and provides a link to the vulnerability for investigation.

To add a note about a vulnerability instance or the work being done on it, select the check box for the instance, and then click **More** > **Add notes**. Enter the note and then click **Add**.

The Vulnerability Responses column displays Added Notes.

To read the note and any previous status changes that were made, hover your cursor over "Added Notes". An historical record about the response to the vulnerability instance appears in a pop-up window.

IoT Risk Assessment

Assessing risk is a continuous process of discovering vulnerabilities and detecting threats. During this ongoing process, IoT Security measures risk and assigns a score for the amount of risk it observes. In fact, IoT Security measures and scores risk at four levels, starting from individual IoT devices and expanding in scope to device profile, site, and finally organization. The different scores provide a simple means to check the risk posed at various points and areas of your network.

When assessing risk, IoT Security uses both static and dynamic factors. Static risks form a baseline and include the following:

- All MDS2 risks (for medical equipment)

- Intrinsic risk factors specific to a profile such as OS, applications, roles, environment

- Trending threats that are hard to mitigate

- The usage behavior specific to a profile or a device

Dynamic risks are added on top of the baseline risk:

- Threats detected in real time (example: alerts)

- Behavioral risks (anomalies, user practice issues) which also trigger alerts

- Vulnerabilities (discovered through passive analysis and detections and through vulnerability scans using integrated third-party vulnerability scanning engines like Qualys and Rapid7)

By collecting and modeling data and analyzing vulnerabilities and threats, IoT Security calculates risk on a daily basis. The risk scores it generates consists of alerts, vulnerabilities, behavioral anomalies, and threat intelligence. When calculating the risk scores of device profiles, sites, and organizations, IoT Security considers not only the scores of individual devices within a particular group but also the percent of risky devices in relation to all devices in the group.

The following sections provide more information about the risk scores that IoT Security generates for these four levels: device, device profile, site, and organization.

Device Risk

IoT Security displays the risk score for each device in the Risk column on the Devices page (**Assets > Devices**). It generates risk scores for devices on a daily basis.

Also see the **Device Details** page (**Assets** > **Devices** > *device-name* > **Device Details**) where the device risk score is listed twice—at the top and in the Security summary section. The Risks section includes a graph that charts changes in the risk score over the specified period of time: day, week, month, year, or all to date. The graph lets you see how the risk score trends over time. Hover your cursor over a marker on the line to see a list of alerts for that point in time. Click a marker to see a list of alerts below the graph.

Detect IoT Device Vulnerabilities

Device Profile Risk

IoT Security displays risk scores for device profiles in the Risk column on the Profiles page (**Assets > Profiles**).

IoT Security uses the scores of individual at-risk devices (that is, those with a risk score of 40 or higher) in the same profile to calculate the risk score for the entire device profile. However, it's not as simple as averaging the risk scores of all the devices in the profile. The computation takes other factors into consideration such as the number of risky devices in the profile.

For example, if five devices in the same profile have individual risk scores of 42, IoT Security would calculate the risk score for the profile to be 89. In this case, because all of the devices in the profile are at risk, the profile score becomes higher than you might have expected at first.

Consider another example, again with five devices in the same profile. One device is at high risk with a score of 98. The other four devices are at normal risk each with a score of 30. In this case, IoT Security calculates the risk score for their profile to be 64. In such a small set, the one high-risk device has a much greater impact on the profile score than it would if the scores of more devices had been involved in the calculation.

Site Risk

See the Risk Score column in the Risk column on the Sites page (**Networks > Networks and Sites > Sites**).

The formula that IoT Security uses to calculate the risk score for a site uses a weighted average of device profile risk scores, the weight for each profile being determined by the number of devices in the profile and the profile risk level.

Organization Risk

See the Risk Score in the Risk panel on the **Dashboards** > **Security Dashboard**.

IoT Security uses the same method to calculate the risk score for an organization as it does for sites.

312

Risk Scores and Severity Levels

The following explains how the severity of a risk score is ranked:

Risk score	Risk severity	Notes
< 40	Low	This is a normal risk level.
40-69	Medium	There might be a few anomalous network behaviors, medium-level alerts, and vulnerabilities with CVSS (Common Vulnerability Scoring System) scores between 4.0 and 6.9.
70-89	High	There might be multiple highly anomalous behaviors, high-level alerts, and vulnerabilities with CVSS scores between 7.0 and 8.9.
90-100	Critical	There might be multiple extremely anomalous behaviors, critical alerts (such as a malware attack), and vulnerabilities with the highest CVSS score of 10.

Adjust Device Risk Scores

It's possible to adjust how much individual risks contribute to the overall risk score of a device. On the **Vulnerabilities** > **Vulnerability Overview** > **All Vulnerabilities** page, click a number in either the Confirmed Instances or Potential Instances column to see details of a vulnerability including which devices it affects or potentially affects. Then click a device name in the Instance column to open the Device Details page for it.

IoT Security categorizes CVE-based risks differently based on their source. When IoT Security discovers them through its internal vulnerability-matching logic (Source = IoT Security Device Software Library) or as a result of a vulnerability scan, it categorizes them as vulnerabilities. When a firewall applies Threat Protection and reports them to IoT Security (Alert Source = Firewall), IoT Security categorizes them as alerts. The Adjust option only appears in the Action menu for vulnerabilities; or, in other words, for risks not categorized as alerts.

In the Vulnerabilities section, expand the Actions menu for a vulnerability and then click **Adjust**.

Take the severity of this risk and its impact on the organization into account and adjust how much you think it should contribute to the overall risk score of the device. Choose whether it makes a low, medium, or high contribution.

Note that the influence of the change you make on the overall score depends on the number and severity of other risk factors. If there are lots of risks, adjusting how much a single risk contributes to the score might not affect it much if at all. On the other hand, if there are only a few risks, adjusting the contribution of one can change the score significantly.

Alerts for Risk Score Changes

When the increase of a risk score causes it to cross a threshold separating one risk level from another, IoT Security generates a risk change alert. (Crossing a risk level threshold as the result of a risk decrease does not trigger an alert.) A risk increase triggers an alert with differing severity levels depending on the new severity of the risk:

- **Warning** when the risk level increases from high to critical
- **Caution** when the risk level increases from medium to high

> *To reduce the overall number of alerts generated, no alert is triggered when the risk level increases from low to medium.*

In addition to risk scores changing because of a manually adjusted risk factor, they can also change for the following reasons:

Increased risk

- A daily risk refresh discovers new vulnerabilities or increased CVSS risk scores.

Decreased risk

- A user resolves a risk factor.
- A daily risk refresh discovers reduced vulnerabilities or decreased CVSS scores or mitigated risks.

Resolve Risks

You can resolve vulnerabilities and security alerts through a workflow built into the IoT Security portal. Essentially, you resolve them by either mitigating or ignoring the vulnerability or alert. As a result, the device risk score might be lowered depending on other contributing factors such as the severity of the risk and the number and severity of other risks. Resolving a vulnerability or alert on a device might similarly affect its profile, site, and organization risk scores depending on how big of an impact the change makes in relation to the number and risk levels of other devices in the

same group. For information about resolving vulnerabilities and security alerts, see Vulnerability Details Page and Act on Security Alerts.

Respond to IoT Security Alerts

Learn about sections of the IoT Security portal that relate to security alerts and how to use them effectively when detecting alerts and responding to them.

- Security Alert Overview
- Create Alert Rules
- Learn about Security Alerts
- Act on Security Alerts
- Routine Security Alert Management

Security Alert Overview

All security alerts that IoT Security generates are based on one of these mechanisms:

- Machine-learning algorithms that automatically learn normal device behavior and can, therefore, detect abnormal behavior.

- Detection of specific traffic patterns—without the use of machine-learning algorithms. For example, IoT Security generates alerts if devices connect to websites that site-reputation services have associated with malware.

- User-defined Security alert rules specifying activity or a state that generates one or more configured actions—a Security alert, user notification, device quarantine. A few examples would be when a specific activity is observed, or when it's not observed, or when a device or group of devices goes offline for two hours. (This period of time isn't configurable.)

- Threats on an IoT device detected by a Palo Alto Networks next-generation firewall are reported to IoT Security in the threat log.

IoT Security examines network traffic in real time, analyzing communications from and to every device on the network. It generates alerts if it detects irregular behavior or activity matching a policy rule.

> *IoT Security generates alerts for IoT devices only. It does not provide alerts, vulnerability detection, policy recommendations, and network behavior analysis for IT devices. For IT devices, IoT Security provides device identification only.*

The Alerts and Alert Details pages in the IoT Security portal provide an overview of all generated alerts and detailed information about individual alerts for analysis and follow-up. IoT Security retains security alerts up to a maximum of one year.

Security alerts pertain to device settings and network behavior that indicate possible security breaches:

- Unsecure device settings (example: devices using the default username and password)

- Suspicious behavior (example: excessive DNS lookup failures)

- Reconnaissance or exploits (examples: port sweeps and EternalBlue SMB exploit attempts)

The Security Alerts section (**Alerts** > **Security Alerts**) consists of three pages:

- ❑ **Alert Overview** – This is a dashboard where you can see alerts that are most relevant to you, analyze risk on IoT devices and on your network, and observe and report alert trends.

- ❑ **All Alerts** – This page displays a table of alerts serially with customizable pagination, columns, and column order. You can filter the information in the table through a dialog box accessed by clicking the Filter icon (≂).

- ❑ **Suppression Rules** – This page is a list of user-defined rules created to suppress the future detection of alerts. For information, see Act on Security Alerts.

Alert Overview

The Alert Overview page is a dashboard with four main sections designed to help you identify top priority alerts, analyze risk, and easily report on alert trends for IoT devices.

Respond to IoT Security Alerts

At the top of the page is an alert summary with information about the alerts matching the filters set for sites, device category, and time range.

- **Active Alerts to Date** – The is the total number of open alerts. An alert can be in one of four states: Detected, Investigating, Remediating, and Resolved. Any alert in one of the first three states—that is, any state except Resolved—is considered open, or active, and is included in this count.

  IoT Security retains security alerts in its database up to one year. If you've been using IoT Security longer than that, keep in mind that this count will not include any alerts discovered more than a year ago.

- **New Alerts in** <time range> – This is the total of all open alerts that were detected within the time range specified in the data filter at the top of the page.

- **Alerts resolved in** <time range> – This is the total of all alerts that were resolved within the time range specified in the data filter at the top of the page.

- **Active Alerts Assigned to Me in** <time range> – This is the total of open alerts that were assigned to the person currently logged in during the time range specified in the data filter at the top of the page.

Alerts of Interest – Define criteria for alerts that matter most to you. IoT Security will then display the top ten alerts in response to your query with the more severe and newer alerts displayed first. For example, if you want to see alerts for a specific vendor or profile that were detected within the last week, click the gear icon (⚙) and configure a query to show the alerts that interest you. IoT Security then displays the ten most recent and most severe alerts that match your terms.

By default, IoT Security uses the predefined "Major Alerts" query to search for critical and high severity alerts detected in the past week for all IoT devices. You can edit this query to define other attributes of interest and then click the bookmark icon (🔖) to save it for reuse.

You can also toggle on **Assigned to me** so that IoT Security displays only alerts within the top ten that were assigned to you. If there are more than ten alerts, **View All** <number> **Alerts** to see the all the alerts that matched your criteria. IoT Security displays these on the All Alerts page. Click an alert name to open the Alert Details page for it.

Alert Distribution – The Sankey chart lets you see the distribution of active alerts across different groupings of devices. Reading the chart from left to right, you start off on the left with all the active alerts that match the site, device category, and time range filters at the top of the page. The chart then relates these alerts to a type of device grouping in the middle and relates these again to another type of grouping on the right. The choices for these groupings are **Severity**, **Profile**, **Device Category**, **Vendor**, **Status**, **Device Type**, and **Alert Type**. Alerts are distributed vertically in the chart by count with those groupings with the most alerts at the top of the chart. When there are more than five groupings, the Sankey chart shows the top five and then gathers everything else in an "Others" group. Hover your cursor over **Others** to see a list of the next ten groupings, and click **View all** to see a pop-up panel with a complete list.

For example, to see the ratio of critical, high, medium, and low alerts among different device categories, choose **Severity** for the middle post and **Device Category** for the right post. The colored bands between the left and middle posts show how many active alerts are critical, high, medium, and low, and the colored bands between the middle and right posts show how many alerts at each severity level were triggered by devices in different device categories. Each band is labeled and shows the total number of active alerts for its severity (on the left) and for that severity per device category (on the right). The width of the bands lets you see at a glance the relative quantities of alerts by their severity. Hovering your cursor over a section of a post shows the percent of alerts for the adjacent bands.

> *Colors only convey meaning to denote alert severity levels: red = critical, orange = high, yellow = medium, and blue = low. For other types of groupings, semi-transparent shades of gray are used solely to distinguish one band from another.*

To download the data from the Sankey chart for your records or reports, click the download icon (↓) in the upper right above the chart. IoT Security saves it as an .xlsx file with alert distribution information on the first sheet and a complete list of active alerts on the second.

Alert Trend – The Alert Trend chart displays a cumulative count of active alerts over the specified time period and a daily noncumulative count of resolved alerts. This visually shows alert trends to help SOC and management teams see if the number of active alerts has been increasing or decreasing over time. It also displays data for resolved alerts, which can help teams gauge their progress in regard to alert resolution. Hover your cursor over different points on the chart to see the number of critical, high, medium, low, and resolved alerts for different dates.

To download data from the Alert Trend chart for reports or records, click the download icon (↓) in the upper right above the chart. IoT Security saves it as an .xlsx file with the active number of alerts to date and resolved alerts over the specified period of time.

All Alerts

The All Alerts page shows all alerts, or *alert instances*, organized by date up to the previous day, which is the last day for which IoT Security has a complete list of alerts. Define filters at the top of the page to control which alerts to display. There are filters for sites, device category, time range, and response status (active alerts, resolved, assigned, unassigned, detected, and all). You can add more filters as well.

The status of an alert begins in the Detected state. You can leave it there or set it to a different state to reflect where it is in the remediation process:

- **Detected**: This is the state of a newly detected alert instance. It makes sense to keep it in this state if no action has been taken to investigate, remediate, or resolve it.

- **Investigating**: Consider setting an alert instance in this state after preliminary work on it has started and it's being verified, researched, and its impact analyzed.

- **Remediating**: Consider setting an alert instance in this state while action is being taken to remediate it but has not yet completed.

- **Resolved**: An alert instance becomes resolved either by mitigating the issue or by ignoring and accepting it.

To change the state of an alert instance, click the entry in the Status column and choose another state. When you resolve it, IoT Security prompts you to provide a reason for its resolution.

To assign an alert instance to someone to work on, select the check box for the instance, and then click **More** > **Assign**. Enter the username or email address of a user and then click **Assign**. The user then receives an email message that states that an alert was assigned to him or her and provides a link to it in the IoT Security portal for investigation.

The person to whom you assign an alert instance must have an IoT Security user account so that it can send a message to the appropriate email address.

IoT Security provides an option for copying the details of an alert instance and creating a work order for use with an asset management system. Select the check box for an instance, and then click **More** > **Copy Alert Information**. Select the sections of the alert description that you want to include in the work order, add additional instructions or relevant information in the Information field, and then click **Copy** to copy the text in those sections.

Paste the copied content into the description field in your asset management console as you manually create a work order there. You can then copy the work order number from the asset management console, paste it back in the Work order field in the Create work order manually dialog box in IoT Security, and then click **Save & Close**.

To add a note about an alert instance or the work being done on it, select the check box for the instance, and then click **More** > **Add notes**. Enter the note and then click **Add**.

To see previously added notes and any previous status changes that were made to an alert instance, click or hover your cursor over the entry in the Last Action column for it. An historical record about the response to the instance appears in a pop-up window.

You can set the number of rows you want to see on each page (from 5 to 200) and navigate among multiple pages.

Security Alert Details Page

Clicking the name of a security alert instance opens the Device Details page.

The Alert Details page is organized into three major sections. At the top is information about the incident itself. The client is always shown on the left, the server on the right, and a rightward pointing arrow between the two—solid if they formed a connection, dashed if a connection was only attempted. The protocol or protocols used in the connection—or attempted connection—are listed below the arrow. The device on which the alert was raised is shown inside a box color

coded to match the severity of the alert. In this way, you can easily see device roles and where the alert occurred.

The client on the left formed a UDP connection with the Avaya IP phone in the server role on the right. The IP phone is the device that raised the alert.

The blue icon next to a device name (arrow pointing out of box) opens a new browser tab showing the Dynamic Topology Viewer with that device in focus (see IoT Security Device Details Page). There you can see how many other devices it communicates with and what they are. This can be extremely useful when investigating a compromised device because it can reveal the location of remote devices participating in the attack and local devices that might be targets of further attacks launched from the victim.

The reference links to a Palo Alto Networks knowledge base article about the Conficker worm.

The Impact section explains how the issue might impact the security of a user, device, or network. (Not all alerts have an Impact section.) The Recommendation section lists options for addressing the issue.

The second major section on the Alert Details page examines the impacted device and summarizes its security status.

You can learn about the identity and activity of the impacted device, its physical location (site), and its logical location on the network. In the Current Behaviors diagram, hover your cursor over any of the five small red circles or the information icon to see more information. The Security section provides security-related information about the device.

The third major section on the Alert Details page shows a snapshot of the network traffic of the impacted device in a Sankey diagram. The diagram includes the IP addresses of other endpoints and the applications used in their communications. The lines indicate various network connections. The ones in red represent the connection involved in the high-severity alert.

If a device has multiple alerts, all relevant lines are colored according to the severity of each one.

Create Alert Rules

IoT Security uses AI and machine-learning algorithms to automatically generate security alerts based on anomalous network behavior and to detect vulnerabilities when device attributes match those in published vulnerability databases such as those at nvd.nist.gov and www.cisa.gov as well as vulnerabilities added to the IoT Security database by its team of security experts. With these automatic detection mechanisms built into the system, IoT Security continuously monitors your network and can notify you of Security threats without any need for you to configure and enable rules or settings for it to do so. However, if you want to detect specific network events (like new device discoveries or a specific device using a specific application), you can define some conditions to identify these events and trigger security alerts and perform actions. To do this, you create custom rules and add them to the set of internal rules that are already in place.

A given rule defined in IoT Security can be triggered based on a single change event such as the discovery of a new device. It can also be triggered by a given traffic pattern such as a specific application command or an accumulation of traffic volume over a period of time. It can even be triggered by a combination of the two. A rule only triggers an action once per day per device to avoid generating excessive noise. To see how many times observed conditions matched a rule, view the Hit Counts column on the **Alerts** > **Custom Alert Rules** page.

The following list shows several types of conditions you might define:

- One device communicates with another device
- A device appears on a specific VLAN or network segment, connects to a specific wireless access point or network switch, or shows up in a specific next-generation firewall zone
- The subnet or IP address of a device changes
- A risky device communicates with the Internet
- The risk level for a device changes
- A device transmits a certain volume of network traffic
- A device uses a particular application or uses something other than a particular application
- A device uses a particular application command or a specific value in a command

If detected, these conditions would trigger IoT Security to take one or more configured actions— generate an alert, notify users, quarantine the device involved.

> *Although the conditions above use the singular form "device" for simplicity, the rule conditions can also apply to multiple individual devices, one or more types of devices (device profiles), or one or more device groups (defined by user tags, Purdue level, or category).*

The rules engine is at **Alerts** > **Custom Alert Rules** and consists of three sections: Basic Information, Rules Details, and Rule Preview.

To help you get started using the rules engine, IoT Security provides a collection of example templates for common rules. Study these preconfigured rules to become familiar with rules engine capabilities, enable and use them as they are, or use them as models for building similar rules of your own.

> *Predefined rules are disabled by default so that they don't trigger unwanted alarms.*

To see the preconfigured example rules, select **Alerts** > **Custom Alert Rules**.

The preconfigured templates differ somewhat based on the vertical theme that's active on your IoT Security portal. Each vertical theme has two or three example rule templates. Here's an example for each theme:

Enterprise IoT Security Plus

- Rule Name: [Example] Suspicious Printer Communication
- Description: Raise a critical alert any time a printer communicates with any other endpoint using applications that aren't on the allow list.
- Rule: WHEN category = "Printer", application =! Dhcp, dns, dns-base, ldap, netbios-ns, ntp-base, smtp-base, snmp-base, snmpv1, ssl, ws-discovery ; DO Publish "Critical" alert
- Action: Raise a critical severity alert

Industrial IoT Security

- Rule Name: [Example] Industrial Device Offline
- Description: Raise a high alert when an industrial controller or remote terminal unit (RTU) is offline during business hours.
- Rule: WHEN category IN ("Industrial Controller", "Industrial RTU"), Offline Device; DO Publish "High" alert
- Action: Raise a high severity alert

Medical IoT Security

- Rule Name: [Example] New Camera Asset Discovered
- Description: Raise a critical alert any time a new IP camera is detected on the network.
- Rule: WHEN: category = "Camera", New Device Discovery; DO Publish Alert
- Action: Raise a critical severity alert

If you want to try a rule, enable it by opening the Rule Engine Editor and toggling the Status from **Disabled** to **Active**. You can edit, clone, and delete the example templates using the options in the Actions column on the **Alerts** > **Custom Alert Rules** page.

STEP 1 | Identify your network concerns and what events you'd want IoT Security to watch for and notify you about.

STEP 2 | Create a rule to address your concern beginning with some basic information.

In the Basic Information section, enter a name and description for the rule and when you want it enforced.

- **Rule Name**: Enter a unique name for the rule.

- **Description**: (Optional) Enter a description of the rule, such as its overall intent, for future reference.

- **Apply rule during**: When days and times of day when you want IoT Security to enforce the rule. By default, a rule is enforced all the time; that is, all day everyday.

- **Status**: If you want IoT Security to monitor the rule conditions, toggle its status to **Active**. If you don't want IoT Security to apply the rule, toggle its status to **Disabled**.

STEP 3 | Define criteria for the rule.

In the Rule Details section, define the criteria necessary to trigger an action that IoT Security will take.

- **All(AND)** or **ANY(OR)**: Choose **All(AND)** when you want all conditions to be met for IoT Security to perform a defined action. Choose **ANY(OR)** when you want any one of the conditions to trigger it.

- **Add Condition**: Choose **Traffic Pattern** to define a condition based on network traffic behavior. Choose **Change Event** to define a condition based on a change to a device, such as a device changing its IP address or going offline, or a new device coming onto the network.

  If you choose Traffic Pattern, IoT Security displays two fields for target devices and extra criteria options for traffic volume and app usage.

  - **Add Target Devices**: In the first target device field, identify the device or devices to which you want to apply the rule. You can do this by choosing up to 10 attributes. These can be the IP addresses and names of devices, the subnets and VLANs to which they belong, previously defined tags and custom attributes, device categories and profiles, the switches and wireless access points through which devices access the network, and traffic destinations (two-character codes for countries and territories defined in the ISO 3166-1 standard). You can specify if the target devices have a specific attribute (=) or not

(**!=**) or if they have any one of a set of attributes (**IN**) or not (**NOT IN**). Building the target device criteria works in a manner similar to the query builder.

- **Target Devices (optional)**: If you want to define a traffic pattern between two specific devices or types of devices, use the second target device field to identify the other end of communications. If you don't enter anything here, it's treated as "any destination" (this can be an internal device or an external web address).

- **Show Extra Criteria**: You can set extra conditions by selecting **Traffic Volume** and **App Usage** and setting parameters (one or both conditions can be selected).

> *The configuration of these settings requires insight into traffic flow volumes and knowledge of application settings and their appropriate values.*

If you select **Traffic Volume**, then enter the volume of traffic and a time period in which it occurs as a condition to trigger a rule. You might want to use this option to watch for unexpected surges in traffic volume, especially to unusual destinations.

If you select **App Usage** and choose **Application: is**, you can then choose a single OT or IoT/IT application and enter whatever commands, parameters, and values must be present in network traffic to trigger an action. If you choose **Application:not**, you can choose a single application that must not be present to trigger an action. If you want

to create a condition that applies to multiple applications, choose **Application:in** or **Application:not in**.

> 📋 *Additional selector fields for commands, parameters, and values are only available after choosing **Application: is**.*

If you choose **Change Event**, IoT Security displays the same Target Devices field it does for **Traffic Pattern** plus an **Event** drop-down list. You can choose the following events to trigger an action:

- IP Change
- New Device Discovery
- New Vulnerability Discovery (includes both confirmed and potential instances)
- Offline Device
- Purdue Level Change (You must also select a Purdue level.)
- Risk Level Change (You must select **Any** or a specific risk level.)
- Subnet Change

Add Condition Set: Adding a condition set lets you create a subgroup of conditions with its own All(AND) or ANY(OR) operator. It's useful for chaining multiple conditions under the main set of conditions.

For example, the following criteria has four conditions with the logic of Condition A **AND** Condition B **AND** { Condition C **OR** Condition D }. To apply an action, conditions A and B and either C or D must be met.

Set the action that IoT Security takes when the defined criteria are met.

To avoid rules generating excessive noise, IoT Security only triggers the specified actions once per device per day. You can configure IoT Security to take up to three of the following actions:

Generate alert + additional actions (**Send to third-party systems** and **Assign to Users**) – When rule conditions are met, IoT Security generates a Security alert and displays them on the Alerts

> Security Alerts page. In addition, IoT Security can automatically push the alert to third-party systems triggering additional actions by the third-party system such as initiating a NAC quarantine or triggering a work order for example. It can also assign an alert to one or more users to investigate for remediation.

Notify users – Set IoT Security to notify multiple users by email or to notify you yourself by SMS text. (To receive text notifications, you must enter your mobile phone number and enable text notifications in *user-name* **> Preferences**.)

Restrict network access – Inform a Palo Alto Networks next-generation firewall to restrict network access to the device whose behavior matches the conditions necessary to trigger the action.

STEP 5 | Check settings in the Rule Preview.

Review the rule displayed in a readable SQL-like format. This is a high-level snapshot of the Criteria and Action sections that lets you check the logical relationships of settings within the rule. Any later changes to settings in these sections will update the rule preview.

Rule Preview {}

WHEN
((
DEVICE GROUP ()
COMMUNICATE WITH
DEVICE GROUP (device.remoteNetwork == '10.0.0.0/8')
)
)
AND ((
DEVICE GROUP (device.localProfile in ['Texas Instruments Device', 'Super Micro Computer'])
)
AND (evt_purdue_level_change == 'Level 0')
)
AND (((
DEVICE GROUP ()
)
AND event.type == 'device_discovered')
OR ((
DEVICE GROUP (device.localCategory == 'Industrial Automation')
)
AND ()
)
)

DO
Publish Alert "High" alerts and retrieve metaData:
byteCount,pktCount,rxBytes,txBytes,appName

Learn about Security Alerts

There are several ways to learn about security alerts. IoT Security can automatically notify you by text and email, depending on the methods you enable in your account preferences. Even if you don't have alert notifications enabled, you might still be notified when another user assigns you an alert for investigation.

You can also learn of alerts in the IoT Security portal itself by checking the Alerts section on the Security Dashboard, hovering over device names on the Devices page, and by viewing the Security Alerts page.

A way to learn about alerts in the IoT Security portal is in the Alerts section on the Security Dashboard. You can organize the alerts on display by severity (low, medium, high, critical), status (detected, investigating, remediating, resolved), device category (for example: audio streaming, IT server, point-of-sale system), or alert type (for example: security risk, unsecure protocol, user policy). When viewing by severity, the numbers in the Alerts column are clickable. Clicking one of them opens the **Alerts** > **Security Alerts** > **All Alerts** page with a filter applied to show only the alerts matching the item you clicked.

When you hover your cursor over a device name on the Devices page, the IoT Security portal displays a pop-up panel with information about the device, including a list of alerts if there are any. Clicking one of the alert names opens the Alert Details page for it.

Click the name of an alert to open the Alert Details page in a new browser window.

Security Alert and System Alert Notification

In addition to viewing security alerts in the IoT Security portal or being notified to investigate an alert, IoT Security also sends email and text notifications automatically when events trigger them. It does this for two types of alerts:

- **Security Alerts** – These alerts pertain to the devices IoT Security is monitoring and are triggered by behavioral changes that indicate a potential attack. Here's an example of a security alert notification:

```
Palo Alto Networks IoT Policy Alert for Super Micro Computer
 device: (Warning) SSH User
        Authentication Brute Force. This event indicates a brute
 force attack through multiple
        login attempts to an SSH server.
```

- **System Alerts** – These alerts pertain to next-generation firewalls. Currently only an outdated application content package triggers a system alert notification.

IoT Security sends these notifications after a user with owner privileges enables them to be sent to all owners (enabled by default) or adds users to a list for notification on **Administration** > **Notification Management**.

The owner can add existing admin users by choosing them from a drop-down list that appears. These users receive notifications by email or text or both depending on their user preferences. The owner can also type in the individual email addresses or distribution lists of users whose email addresses share the same domain of one of the owners. (IoT Security rejects any address with a domain that's not shared by an owner.) These users receive notifications by email. If an owner disables **Send to all the owners**, then only those in the email lists will receive notifications.

Act on Security Alerts

After you learn about a security alert, one of the first steps is to read the details and confirm that the event that triggered it actually occurred, possibly by checking firewall event log entries. After confirming the alert, you must quickly assess its importance and urgency, identify the type of equipment impacted, and then decide how to respond and with whom to engage. The responder might be IT security, clinical engineering, a third-party network security service provider, or perhaps the device vendor or manufacturer. Find the responsible party and contact them about the alert.

Take Action when a Security Alert Occurs

There are numerous ways to respond to a security alert. The action you take depends of the remediation requirements of the situation:

- If a device was infected by malware or a virus, unplug the device immediately. If its continued use is essential, work with IT security to quarantine it from the rest of the network. You might need to modify firewall security policies to permit only traffic absolutely required for the device to function and block everything else while you work on a resolution.

- The resolution might require a software patch, and sometimes you might have to get the equipment vendor involved to patch it. If you must continue using the equipment, enforce a strong zero-trust policy until the patch is available.

- If an alert is generated by a security policy violation, you can send policy recommendations to the firewall so it only permits traffic resulting from normal device behavior.

- To assist in your analysis, IoT Security provides alert log files (in .csv and .log formats), which contain several days' worth of network connections involving the device that triggered an alert. You can also download the network traffic data that IoT Security shows as a Sankey diagram and view it as an .xls spreadsheet.

Assign and Track Security Alerts

From the Alerts and Alert Details pages, you can assign a security alert to one or more people for investigation. When you select an alert on **Alerts** > **Security Alerts** > **All Alerts**, a set of actions appears at the top of the alerts table.

To assign an alert to someone to investigate, click**More** > **Assign** . Enter an email address and comment and then **Assign**.

> *If you assign an alert to an external user—that is, someone who doesn't have a Palo Alto Networks user account and can't log in to the IoT Security Portal—a PDF with alert details will be attached to the email.*

You can also assign an alert occurrence to someone from the Alert Details page (**Alerts** > **Security Alerts** > **All Alerts** > *alert_title*) by clicking **Action** > **Assign**.

You can also add notes to an alert, which is a convenient way for you and your team to track the progress of investigations of high-level alerts. From the Alerts page, select an alert and then click **More** > **Add notes**. From the Alert Details page, click **Action** > **Add Notes**. The notes appear in the Alert Events list on the Alert Details page.

Resolve and Reactivate Security Alerts

You resolve a security alert either by accepting it or by addressing the issue in some way, perhaps by assigning it to a network security administrator to investigate and fix.

The Resolve tool is useful for showing how many alerts got resolved in weekly or monthly reports.

If you consider one or more alerts acceptable, such as one at a low severity level, you can resolve them. It is not necessary to resolve each alert occurrence individually. You can select the check box next to the alert group names and then click **Resolve** at the top of the Alerts list.

After clicking **Resolve**, the Resolve Alert dialog box appears. Select the reason for resolving it, add a comment, and then **Resolve**.

If you later decide to reactivate one or more alerts that were previously marked as resolved, you can do so by setting the filter above the Alerts list to **Resolved**, selecting the alerts, and then clicking **Unresolve** . In the Change Status dialog box, enter a comment and then click **Change**.

Suppress Security Alerts

If IoT Security raises a security alert for an expected event, you can suppress future occurrences of the alert so no further resources need be expended on them. You can suppress future alert detections for just the device on which the alert was triggered or for all devices sharing the same device profile, category, or device type. You can suppress the alert indefinitely or for a limited length of time. In addition to suppressing future alert detections, you can also mark the current alert event as resolved.

To suppress an alert, log in to IoT Security as a user with administrator or owner privileges and select **Alerts** > **Security Alerts** > **All Alerts**. Select the alert that you want to suppress and then click **More** > **Suppress Alerts**.

You can select multiple alert instances if they are the same type of alert (with the same alert name). When different alert types are selected, the Suppress option becomes unavailable.

To suppress all future alert detections for the device or devices on which the alert was triggered, add a comment, leave **Resolve this alert** selected, and then click **Save**.

To suppress future alert detections on additional devices as well as this particular device, expand **Suppression Rule**, choose one or more attributes in one or more of the Tag, Category, Profile, and Device Type fields, set the length of alert suppression, add a comment, and then click **Save**. Cortex XSOAR will suppress future alerts occurring on devices matching any of the chosen attributes for the length of time specified.

After you create a suppression rule, it takes IoT Security approximately 30 minutes to apply it throughout the system to all the devices in your inventory. IoT Security also adds it to the rule table at **Alerts** > **Security Alerts** > **Suppression Rules**.

Clicking a rule name opens the Suppress Alert configuration panel where you can view and edit details. The Status column indicates two states. A rule is "In process" during the initial 30-minute application period after it's been created or modified. After that, the status changes to "Success" indicating that IoT Security has applied the rule to all the targeted devices in its inventory.

After you create a rule, you can always modify it to include additional devices by modifying the rule to encompass a wider range of devices. In fact, IoT Security prompts you to do this whenever you are about to suppress an alert on a device and there's already a suppression rule for this type of alert but it just doesn't apply to this particular device. It displays an information icon, which expands into a pop-up message when you hover your cursor over it.

To add just this device to the existing rule, optionally add a comment and leave **Resolve this alert** selected, and then click **Save**. To apply the suppression rule to this device and others like it, expand **View targeted devices**, modify the original rule to include the profile, category, or device type that would make it apply to this and similar devices, and then click **Save**.

To stop alert suppression, log in to IoT Security as a user with administrator or owner privileges and select **Alerts** > **Security Alerts** > **Suppression Rules**. Select one or more rows in the table and then click **Release Suppression**.

Because vulnerability scanners generate traffic that triggers lots of alerts, you most likely want to suppress alerts for them. If you have an IoT Security Third-party Integrations Add-on license or a full-featured Cortex XSOAR server, you might have integrated IoT Security through Cortex

XSOAR with Qualys, Rapid7, or Tenable vulnerability scanners. If so, IoT Security automatically imports the names and IP addresses of all scan engines, and the names of all sites and vulnerability scan templates from the integrated product and adds them to the list of scanners on **Settings > Scanners**. The Source column indicates that a scanner was automatically imported by displaying the integration product name: **Qualys**, **Rapid7**, or **Tenable**. If you don't want to automatically import this information to the scanners list, disable **Automatically Synchronize Scanners with IoT Security** in one of the following Cortex XSOAR jobs, depending on which integration you're using: PANW IoT Get Qualys Scanners and Profiles, PANW IoT Get Rapid7 Scanners and Profiles, or PANW IoT Get Tenable Scanners and Profiles. Disabling this setting doesn't automatically remove previously imported scanners from the list in the IoT Security portal. You must remove them manually by selecting them in the list, clicking **Remove from Scanner List**, and then clicking **Continue** at the prompt.

If you want to suppress alerts triggered by vulnerability scanners that are on your network but not integrated with IoT Security, create a list of scanner IP addresses and upload it to IoT Security. Click **Settings** > **Scanners**, click **Add Scanners**, and then download a CSV template.

For each scanner, add its IP address and optionally its MAC address and a comment.

Upload the file to IoT Security. If IP addresses in the CSV file match those in the device inventory, IoT Security adds them to the scanner list and begins to suppress alerts for them. (It can take up to an hour after the upload for alert suppression to begin.) The Source column in the Scanners table indicates that a scanner was manually uploaded by displaying **User**. If IP addresses are new to IoT Security, it adds them to the scanner list and it adds them to the inventory as scanners after detecting network traffic for them. If there are duplicate entries, IoT Security skips them during the upload process. Finally, if there's a mismatch between the IP-and-MAC-address pairing for an uploaded scanner and the pairing for a device in its inventory, IoT Security does not upload it.

Routine Security Alert Management

Regularly monitor the notes added to the Alert Events list for the high-level security alerts you're tracking. This is an efficient way for team members to coordinate efforts and check on the status.

Review low-severity alerts on a daily basis. Select the ones that you find acceptable and resolve them all with a few simple clicks as explained in the previous section.

On a weekly or monthly basis, download all the alerts and all the resolved alerts. Use the data there to make a status report to show what your team has done.

In addition to reacting to alerts that already occurred, you can proactively address vulnerabilities before an attack takes place. On **Dashboards** > **Security Dashboard**, check the Active Vulnerabilities to Date entry in the Risk panel.

Click **Active Vulnerabilities to Date** to open the **Vulnerabilities** > **All Vulnerabilities** page.

By default, the IoT Security portal sorts vulnerabilities by severity, displaying the most severe vulnerabilities first. When you click a vulnerability name, the Vulnerability Details page for it opens. There you can see which devices are vulnerable so you can take steps to remove the vulnerability before it's exploited in an attack.

Recommend Security Policies

IoT Security uses machine learning to automatically generate policy rule recommendations based on the normal, acceptable network behaviors of IoT devices in the same device profile.

- Policy Rule Recommendations
- Device Profile Overview
- Device Profile Behaviors
- Device Profile Policy
- Create a Policy Set in IoT Security
- Import a Policy Set into Panorama
- Restrict Network Access

Policy Rule Recommendations

IoT Security uses machine learning to automatically generate Security policy rule recommendations based on the normal, acceptable network behaviors of IoT devices in the same device profile. It then provides these recommendations for next-generation firewalls to control IoT device traffic.

IoT Security derives its recommendations from the network behaviors it observes in traffic generated by IoT devices in the same profile across multiple IoT Security tenants. It classifies the applications in the observed behaviors into three groups:

- **Common applications not locally observed** – Applications commonly used by devices in the device profile in multiple IoT Security tenant environments but not observed in yours

- **Common applications locally observed** – Applications commonly used by devices in the device profile in multiple IoT Security tenant environments including yours

- **Unique applications** – Applications that are not typically used by devices in this device profile and were observed in use by devices in your environment only

> *From PAN-OS 11.1, there's a* different process *for recommending Security policy rules to next-generation firewalls from that described here. The following workflow remains applicable to firewalls running PAN-OS versions prior to PAN-OS 11.1.*

IoT Security then formulates a set of policy rule recommendations. These rules allow devices in this device profile to continue network behaviors that are common among multiple tenant environments and those that are unique to yours. The premise is that these behaviors are necessary for devices belonging to this device profile to function. You can accept all these recommendations or disable or modify individual rules to meet the security requirements of your network. When you're satisfied with a policy set, save and activate it. Once activated, it becomes available for firewalls to import—either through Panorama or directly—and then add to their rule set.

When a Panorama or firewall administrator imports a set of Security policy rules from IoT Security, the import operation automatically creates device objects from source and destination profiles in the recommended rules and uses those objects in the Security policy rules it constructs. For the firewall to identify which IoT devices to apply its policy rules to, it uses IP address-to-device mappings that IoT Security provides through Device-ID. The firewall learns the device profile of an IoT device from the mapping and applies rules with matching device objects as the source.

> *The IoT Security app makes policy rule recommendations only for IoT devices that it has identified with a high degree of confidence (a confidence score of 90-100%). It does not consider the network behaviors of low- and medium-confidence IoT devices (0-69% and 70-89% scores). In addition, IoT Security does not provide policy rule recommendations, alert and vulnerability detection, and network behavior analysis for IT devices, which are devices that aren't built for a specific task, such as personal computers, smart phones, and tablets for example. For IT devices, the IoT Security app provides device identification only.*

After allowing sufficient time for IoT Security to collect the full behaviors of IoT devices in a profile, you're ready to create policy rule recommendations for it.

To begin, log in to the IoT Security portal, navigate to **Assets** > **Profiles**, and then click a profile name.

Recommend Security Policies

IoT Security displays three profile details pages:

- **Overview** – View a summary about the high-confidence IoT devices in this profile and their related risk factors for the past day, week, or month. See Device Profile Overview.

- **Behaviors** – View the behaviors of high-confidence IoT devices belonging to this profile in your local network environment and in other IoT Security tenants' environments. Also create Security policy rule sets based on these observed behaviors for next-generation firewalls. See Device Profile Behaviors.

- **Policy** – View previously created Security policy rule sets for next-generation firewalls and ACL rule sets for integration with Cisco ISE. IoT Security generates both types of rule sets from the observed network behaviors of high-confidence IoT devices in this profile in your local network environment and in other IoT Security tenants' environments. See Device Profile Policy.

Device Profile Overview

To access the Overview page of a device profile, select **Assets** > **Profiles >** *profile_name* **>
Overview**.

Recommend Security Policies

The Overview page displays data about the devices in this profile. The data is drawn only from IoT devices with high confidence scores of 90-100%; that is, devices that IoT Security has identified with a high degree of confidence. If the number of high-confidence devices is less than 50%, consider using the recommendations provided on the Data Quality Diagnostics page (**Administration** > **Data Quality**) to increase the number of high-confidence devices in the profile.

Time filter – The time filter controls the data displayed on the Overview page by the number of high-confidence devices in the profile that were active on the network during the past 1 Day (past 24 hours up until now), past 1 Week, or past 1 Month. Clicking the **Reset filter** icon (⟲) sets it to **1 Day**.

The time filter only affects the display of high-confidence devices in the local network, not that of all devices.

Summary bar – The profile summary across the top of the Overview page concisely presents important information about the devices in the profile: the overall number of devices, the number of high-confidence devices, the risk score for this device profile (for risk assessment details, see IoT Risk Assessment), the number of alerts and vulnerabilities of the high-confidence devices, and the number of policy sets configured for this profile.

You can configure multiple policy sets for the same profile but only one of them can be activated at a time.

Below the summary are several sections about key aspects of the device profile and related risk factors. IoT Security produces this information by using machine learning to observe and analyze the network activity of all the high-confidence devices in the profile. It then compares the information about your devices with those in the same device profile in other IoT Security tenant networks to give you a sense of how your device behaviors and risk levels match up with others.

Profile Behavior – This shows the different types of outbound and inbound behavior of the high-confidence devices. Switch between the two behaviors by clicking **Outbound** and **Inbound**.

IoT Security compares the applications that the high-confidence devices in this profile use during the time range set at the top of the page with the applications that devices in the same profile use in other IoT Security tenants. The time filters are 1 Day, 1 Week, or 1 Month. It then shows how many applications were observed in other tenants' environments only (common, not locally observed), in both your and other tenants' environments (common, locally observed), and in your environment only (unique applications).

Most Common Alerts for *profile_name* – This lists up to five of the most common security alerts raised by devices in this device profile across multiple IoT Security tenants and their severity levels. The number of alerts raised by your devices is also shown in the column labeled Your Alerts.

Top Vulnerabilities in for *profile_name* – This lists up to five of the top vulnerabilities affecting devices in this device profile across multiple IoT Security tenants and their severity levels. The number of vulnerability instances in your network environment is also shown in the column labeled Your Vulnerability Instances.

Risk Score – This shows the risk score for the device profile in relation to the overall range and to the average of all IoT Security tenants with the same profile. This helps you see the level of risk for your devices relative to the average level of other IoT Security tenants.

In the following screen capture, the range extends from 10 to 89, which are the lowest and highest risk scores for this device profile among all IoT Security tenants, and the average risk score is 13. With a local risk score of 74, you might consider addressing some threats to reduce risk and lower the score away from the high end of the range.

Device Profile Behaviors

To access the Overview page of a device profile, select **Assets** > **Profiles >** *profile_name* > **Behaviors**.

Recommend Security Policies

The Behaviors page displays the behaviors of high-confidence IoT devices in this profile. These are IoT devices that IoT Security has identified with a high degree of confidence and has calculated a confidence score of 90-100%. The behaviors are those of IoT devices belonging to the same profile in your local network environment and in the network environments of other IoT Security tenants.

A confidence score indicates the level of confidence IoT Security has in its identification of a device. IoT Security has three confidence levels based on calculated confidence scores: high (90-100%), medium (70-89%), and low (0-69%).

Filter the Content Displayed

The behaviors displayed on this page and in the related Sankey chart are controlled by the filters at the top of the page; the option to show either outbound or inbound behaviors; and the option to show common applications, unique applications, or both (the default) under Applications in the Profile Behaviors section.

Recommend Security Policies

The time filter also determines which outbound, or inbound, behaviors are displayed.

> 📋 *You can only create a policy rule set for outbound behaviors; that is, when the source of a behavior is an IoT device in a device profile. IoT Security does not generate policy rule recommendations for inbound behaviors, which is when the IoT device is the destination.*

Time filter – The time filter controls the behaviors displayed on the Behaviors page by when each behavior was observed on the network during the past 1 Day (past 24 hours up until now), past 1 Week, or past 1 Month. Clicking the **Reset filter** icon (🔄) sets the time to 1 Day and removes any additional filters you might have set.

Add Filters – Add filters to show specific types of behaviors. Select one or more of the following:

- **Applications** – Select and choose one or more applications to show on the page. The applications listed were part of high-confidence IoT device behaviors observed on the network during the period of time set in the time filter.

- **Local Observed** – Select and choose **Yes** to show behaviors observed locally in your network or **No** to hide locally observed behaviors.

- **App Usage** – Select and determine what to show based on where behaviors were observed:

  - **Common only** are behaviors that were observed in other IoT Security tenant environments but not in yours.

  - **Common and local** are behaviors observed in other tenant environments and in yours.

  - **Local only** are behaviors observed in your environment but not in the environment of any other tenant.

- **Unexpected behavior** – Select and choose **Yes** to show behaviors that were explicitly not permitted when the policy set was activated but have since appeared on the network. Choose **No** to hide unexpected behaviors.

- **New behavior** – Select and choose **Yes** to show behaviors discovered on the network after the last policy set activation. Choose **No** to hide new behaviors.

> 📋 *None of these filters nor the time filter determines which behaviors to include in any policy sets you might create. They only determine what to show on the Behaviors page. However, once you start the process of creating a policy set, IoT Security presents a similar set of filters to use within the context of policy creation.*

As you add and remove filters, the number in parentheses next to "Profile Behavior" changes accordingly. Refer to this for a quick reference of how the filters affect the number of behaviors that appear on the page while the filters are in place.

Outbound Behaviors and **Inbound Behaviors** – By default outbound behaviors are shown. These are behaviors in which this device profile is the source of network activity.

In the upper screen capture of the two below, there are 66 outbound behaviors:

- 10 outbound behaviors include common applications observed in both your and other tenants' environments. These are indicated with a green fill in the Venn diagram and bar chart.
- 56 include applications that are unique to your environment. These are indicated as gray.

In the lower screen capture, there are 11 inbound behaviors, which are behaviors in which this device profile is the destination of network activity:

- 9 inbound behaviors include common applications observed in other tenants' environments but not in yours. These are indicated with a green outline in the Venn diagram and bar chart.
- 1 includes a common application observed in both your and other tenants' environments. This is indicated with a green fill.
- 1 includes a unique application observed only in your environment. This is indicated as gray.

| Outbound Behaviors | Inbound Behaviors |

Profile Behavior (66) View Sankey Chart

Create Policy

Overview

All common applications are observed locally

All common applications used by devices in this device profile in other IoT Security tenant environments are used by devices in your environment. Your devices also used some unique applications.

Detail
Applications

Common Application   10
Unique Application   56

○ Common applications not locally observed ⓘ 0   ● Common applications locally observed ⓘ 10   ● Unique applications ⓘ 56

| Outbound Behaviors | Inbound Behaviors |

Profile Behavior (11) View Sankey Chart

Overview

Some common applications observed locally

Some common applications used by devices in this device profile in other tenant environments were used by devices in your environment.

Detail
Applications

Common Application   10
Unique Application   1

○ Common applications not locally observed ⓘ 9   ● Common applications locally observed ⓘ 1   ● Unique applications ⓘ 1

The direction you choose—outbound or inbound—controls what's shown in the list at the bottom of the Behaviors page and in the Sankey chart. Your choice also shows or hides the **Create Policy** button, only showing it when **Outbound Behaviors** is active. Clicking the number to the right of the bar charts also controls whether to show common or unique applications on the page and in the Sankey chart. To undo the filter applied by clicking either of these numbers, click the **Reset filter** icon (↻) next to the time filter near the top of the page.

Create a Policy Set

Use IoT Security recommendations to create policy rule sets based on the observed network behaviors of IoT devices in the same device profile. For instructions on creating a policy set, see Create a Policy Set in IoT Security.

📋 *From PAN-OS 11.1, there's a different process for recommending Security policy rules to next-generation firewalls from that described here. The following workflow remains applicable to firewalls running PAN-OS versions prior to PAN-OS 11.1.*

View the Sankey Chart

A Sankey chart is a diagram with lines indicating connections. Click **View Sankey Chart** to open a panel on the right showing the flow of applications from a source (the current device profile in outbound behaviors) to destinations and the destination locations (internal or external). The lines are color coded as explained above and grouped into these three groups:

- Gray for unique local applications
- Green fill for common applications locally observed
- Green outline for common applications not locally observed

A *dash appears for an internal destination when the device profile of a destination is unknown. The number after a destination indicates the number of different IP addresses at the destination profile (for internal destinations) or behind the domain name (for external destinations).*

Because the Sankey chart can become overwhelming when there are lots of lines, you can apply filters to reduce their number. For example, applying a filter that shows only locally observed applications reduces the number of lines in the diagram shown above from 24 to 11 while also increasing line width. See below.

You can also apply an application filter. For example, if there is one application that interests you, you can show only behaviors that include that. You can also filter by multiple applications. The following screen capture shows outbound behaviors just for NTP.

Another feature of the chart is that you can hover your cursor over lines and blue bars to see information pop-ups. In the screen capture above, the cursor is hovering over the destination bar where one of the common behaviors crosses it to show a pop-up identifying its particular destination. This is useful for seeing complete destination profile names and domain names, which are abbreviated in the chart.

View the Behaviors Table

At the bottom of the Behaviors page is a table listing all the behaviors for this profile matching the filters that have been set: the time filter and additional filters near the top of the page, the outbound or inbound behaviors toggle, and the common or unique application numbers under Detail Applications. The data in the table is aggregated with behaviors grouped by application.

Source Pro...	Application	App Risk ⓘ	Locally Obs...	App Usage	Alert Raised	Destination ↓	Location	New Behavior ⓘ	Last Seen
Zoom Video ...	zoom-meeting	2	Yes	Unique	0	zoomsxu112mr	both	Yes	Jan 24, 2022, 1
Zoom Video ...	zoom-base	1	Yes	Common	12 ❗	zoomsxac30zc.:	both	Yes	–
Zoom Video ...	stun	2	Yes	Unique	0	zoomnxo30zc.n	both	Yes	Jan 24, 2022, 1
Zoom Video ...	google-hango...	3	Yes	Unique	1 ❗	www.recaptcha	external	Yes	Jan 23, 2022, 1
Zoom Video ...	quora-base	1	Yes	Unique	0	www.quora.con	both	Yes	Jan 24, 2022, 1
Zoom Video ...	unknown-udp	1	Yes	Unique	1 ❗	www.google-ar	both	Yes	Jan 12, 2022, 1
Zoom Video ...	google-analy...	2	Yes	Unique	0	www-google-ar	external	Yes	Jan 24, 2022, 1
Zoom Video ...	workday-do...	2	Yes	Unique	0	wd5.myworkda	external	Yes	Jan 24, 2022, 1
Zoom Video ...	workday-base	1	Yes	Unique	0	wd5.myworkda	external	Yes	Jan 24, 2022, 1
Zoom Video ...	vimeo-base	4	Yes	Unique	0	vimeo-video.ma	both	Yes	Jan 13, 2022, 1
Zoom Video ...	facebook-base	4	Yes	Unique	0	star.c10r.facebc	both	Yes	Jan 24, 2022, 1
Zoom Video ...	yelp-base	1	Yes	Unique	0	s3-media0.fl.yel	external	Yes	Jan 24, 2022, 1
Zoom Video ...	youtube-base	4	Yes	Unique	1 ⚠	s.youtube.com,ı	both	Yes	Jan 24, 2022, 1
Zoom Video ...	disqus	2	Yes	Unique	0	referrer.disqus.c	both	Yes	Jan 11, 2022, 1
			Yes	Unique				Yes	Jan 12, 2022, 1

The App Risk column contains the risk level for this application as defined in Applipedia. Risk levels are graded from 1 to 5, with numbers approaching 5 carrying increasingly more risk. Hover your cursor over the application name to display a pop-up panel with information about the application retrieved from Applipedia. For explanations about this information, see IoT Device Applications Discovery.

ssl

Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols which provide secure communications on the Internet for such things as web browsing, e-mail, Internet faxing, instant messaging and other data transfers.

Characteristics	Security Information
Category	networking
Subcategory	encrypted-tunnel
Risk Level	4
Standard Ports	tcp/443
Technology	browser-based

The number of Security alert instances and their severity levels are presented in the Alerts Raised column. For outbound behaviors, you can see the number of alert instances that occurred on devices in the source profile for the application in each row.

Source Profile	Applic...	App Risk ⓘ	Locally Obs...	App Usage	Alert Raised ↓	Destination	Locati...	New Behavior	Last Seen
PRTG Network Monitor	ping	2	Yes	Common	15 ❗	clc.stackoverflow.co	both	Yes	—
PRTG Network Monitor	ntp	2	Yes	Common	7 ❗	0.pool.ntp.org,0.clou	both	Yes	—
PRTG Network Monitor	dns	3	Yes	Common	6 ❗	resolver1.opendns.c	both	Yes	Jan 24, 2022, 16:00
PRTG Network Monitor	ssh	4	Yes	Unique	6 ❗	-,Palo Alto Network	internal	Yes	Jan 24, 2022, 16:00
PRTG Network Monitor	kerberos	2	Yes	Common	6 ❗	PRTG Network Mor	internal	Yes	Jan 24, 2022, 16:00
PRTG Network Monitor	ldap	2	Yes	Common	6 ❗	PRTG Network Mor	internal	Yes	Jan 24, 2022, 16:00
PRTG Network Monitor	msrpc-base	2	Yes	Common	6 ❗	PRTG Network Mor	internal	Yes	Jan 24, 2022, 16:00
PRTG Network Monitor	snmp-base	2	Yes	Unique	6 ❗	HPE Networking Sw	internal	Yes	Jan 24, 2022, 16:00
PRTG Network Monitor	snmpv3	1	Yes	Unique	6 ❗	HPE Networking Sw	internal	Yes	Jan 24, 2022, 16:00
		1	Yes	Common	6 ❗	...rnal		Yes	Jan 24, 2022, 16:00

For inbound behaviors, the Alerts Raised column shows the number of alert instances that occurred on devices in the destination profile for an application.

Source	Locati...	Application	App Risk ⓘ	Locally Obs...	App Usage	Alert Raised ↓	Destination Profile	New Behavior	Last Seen
Palo Alto Netw	internal	paloalto-updates	2	Yes	Unique	2 ●	PRTG Network Monitor	Yes	Jan 24, 2022, 16:00
Macintosh,Win	internal	paloalto-device...	1	Yes	Unique	1 ●	PRTG Network Monitor	Yes	Jan 24, 2022, 16:00
Windows Table	internal	dns	3	Yes	Common	0	PRTG Network Monitor	Yes	Jan 23, 2022, 16:00
Arista Network	internal	icmp	4	Yes	Unique	0	PRTG Network Monitor	Yes	Jan 23, 2022, 16:00
PC-Windows,M	internal	lpd	3	Yes	Unique	0	PRTG Network Monitor	Yes	Jan 24, 2022, 16:00
Windows Table	internal	ms-ds-smbv3	3	Yes	Common	0	PRTG Network Monitor	Yes	Jan 24, 2022, 16:00
Windows Table	internal	msrpc-base	2	Yes	Common	0	PRTG Network Monitor	Yes	Jan 24, 2022, 16:00
Macintosh,PC-\	internal	netbios-ns	2	Yes	Common	0	PRTG Network Monitor	Yes	Jan 24, 2022, 16:00
-	internal	paloalto-gp-mf...	1	Yes	Unique	0	PRTG Network Monitor	Yes	Jan 24, 2022, 16:00
		...share...	1	Yes	U...		...k Monitor	Yes	Jan 24, 2022, 16:00

Alert instance totals in the Alerts Raised column are grouped by their severity level: critical, high, medium, and low. The following icons indicate these four levels:

A behavior for a source device profile and application might have numerous destinations. You can drag the destination column to widen it but that still might not be sufficient to see all of them. To open a panel with detailed information, click anywhere in the destination field.

Source Pro...	Application	App Risk ⓘ	Locally Obs...	App Usage	Alert Raised	Destination
Zoom Video ...	zoom-meeting	2	Yes	Unique	0	zoomsxu112mmr.sx.zoom.us,147.124.97.57,zoomsjcgm152mmr.sjc.zoom.us,zoo...
Zoom Video ...	zoom-base	1	Yes	Common	12 ⓘ	zoomsxac30zc.sx.zoom.us,zoomsxab30zc.sx.zoom.us,zoomsxt31zc.sx.zoom.us,zo...
Zoom Video ...	stun	2	Yes	Unique	0	zoomnxo30zc.nx.zoom.us,zoomsjccl213zc.sjc.zoom.us,zoomsxp31zc.sx.zoom.us,z...
Zoom Video ...	google-hango...	3	Yes	Unique	1 ⓘ	www.recaptcha.net,www.gstatic.com
Zoom Video ...	quora-base	1	Yes	Unique	0	www.quora.com,quora.map.fastly.net,-
Zoom Video ...	unknown-udp	1	Yes	Unique	1 ⓘ	www.google-analytics.com,www.google.com,client3.google.com,www.googleapis...
Zoom Video ...	google-analy...	2	Yes	Unique	0	www-google-analytics.l.google.com,ssl.google-analytics.com,ssl-google-analytics.
			Yes	Unique		

The View Destination for *application_name* panel provides its own table with rows for each individual destination to which devices in the source device profile sent a particular application. Hover your cursor over a number in the Destination IP column to see a pop-up with a list of IP addresses.

If you are looking for a specific destination IP address and the list of addresses is too long for the Destination IP pop-up to display them all, click the number in the Destination IP column and a dialog box appears with a search option.

In the Behaviors table, the Location column indicates where the destinations of a behavior are. If all the destinations are in the local network, the location is *internal*. If all the destinations are outside the local network, the destination is *external*. If some destinations are internal and some external, then the location is both. In this case, you can see the location of individual destinations by clicking in the Destination column in the Behaviors table and looking at the Location column in the View Destination for *application_name* panel.

View destination for google-base

Destinations (73)

Destination	Locally Observed	App Usage	Destination...	Location
-	Yes	Unique	9	internal
www.google.com	Yes	Unique	13	external
gcdn.2mdn.net	Yes	Unique	1	external
	Yes	Unique		external

Source Pro...	Application						New Behavior ⓘ	Last Seen
Zoom Video ...	stun						Yes	Jan 25, 202:
Zoom Video ...	zoom-meeting						Yes	Jan 25, 202:
Zoom Video ...	ping						Yes	—
Zoom Video ...	google-play						Yes	Jan 25, 202:
Zoom Video ...	reddit-base						Yes	Jan 25, 202:
Zoom Video ...	gmail-base						Yes	Jan 25, 202:
Zoom Video ...	google-hangouts-b...						Yes	Jan 24, 202:
Zoom Video ...	google-meet	Yes	Unique	2	0	www.googleapi external	Yes	Jan 25, 202:
Zoom Video ...	unknown-udp	Yes	Unique	1	0	www.googleapi external	Yes	Jan 25, 202:
Zoom Video ...	google-base	Yes	Unique	4	1 ⚠ 2 ⓘ	www.google.co both	Yes	Jan 25, 202:
Zoom Video ...	websocket	Yes	Unique	3	0	wd5.myworkda both	Yes	Jan 25, 202:
Zoom Video ...	workday-download...	Yes	Unique	2	0	wd5.myworkda external	Yes	Jan 25, 202:

371

Device Profile Policy

📋 *From PAN-OS 11.1, there's a* different process *for recommending Security policy rules to next-generation firewalls from that described here. The following workflow remains applicable to firewalls running PAN-OS versions prior to PAN-OS 11.1.*

To access the Policy page of a device profile, select **Profiles >** *profile_name* **> Policy**.

This page lists all the policy sets that were created for the device profile, when they were last updated, whether they were activated, and if so, when. When there are no policy sets for a device profile, the Policy page is empty.

If you create a policy set for a device profile and save it without activating it, it's added to the Policy page. In this case, there's a dash in the Last Set as Active column.

After you activate a policy set, it's marked with an Active label and IoT Security adds a timestamp in the Last Set as Active column.

If you later deactivate the policy set, the Active label is removed. However, the timestamp in the Last Set as Active column remains indicating that it once was active and when.

374

New behaviors are behaviors discovered on the network after the active policy set was activated or last updated. Unexpected behaviors are behaviors that were explicitly not permitted when the policy set was activated or last updated but have since appeared on the network, which means the enforcement implemented in a next-generation firewall is missing them. If IoT Security detects new or unexpected behaviors on the network after some time has passed since the policy set was first activated, it lists them on the **Assets** > **Profiles** > *profile_name* > **Policy** page and presents you with an opportunity to modify the active policy set to account for these behaviors.

Recommend Security Policies

Search devices, security alerts, vulner...

OVERVIEW BEHAVIORS POLICY

Filter Query 1 Week ⌄

New Behavior ⓘ

Modify Active Policy

We found 24 new behaviors since the active policy set was created. (Note that new common behaviors might have been observed in other IoT Security tenants' environments but not necessarily in yours.) Review the new behaviors to make sure all critical behaviors are allowed and no critical behaviors are blocked.

Source Profile	Application	App Usage	Destination
Zoom Video Con...	dns	Common	1
Zoom Video Con...	dhcp	Common	1
Zoom Video Con...	traps-manageme...	Unique	1
Zoom Video Con...	apple-push-notifi...	Unique	1
Zoom Video Con...	netflix-base	Unique	1
Zoom Video Con...	ms-office365-base	Unique	1
Zoom Video Con...	okta	Unique	1
Zoom Video Con...	icloud-mail	Unique	1
Zoom Video Con...	directv	Unique	1
Zoom Video Con...	apple-update	Unique	1

Unexpected Behavior ⓘ

Modify Active Policy

We found 42 behaviors blocked in your policy is still actively hitting the traffic. Review your policy and also firewall implementation to make sure only necessary behaviors are allowed.

Source Profile	Application	App Usage	Destination
Zoom Video Con...	ssl	Unique	139
Zoom Video Con...	google-base	Unique	70
Zoom Video Con...	zoom-base	Common	29
Zoom Video Con...	facebook-base	Unique	4
Zoom Video Con...	youtube-base	Unique	62
Zoom Video Con...	apple-maps	Unique	8
Zoom Video Con...	gmail-base	Unique	17
Zoom Video Con...	itunes-base	Unique	8
Zoom Video Con...	ocsp	Unique	15
Zoom Video Con...	ntp	Unique	10

Policy Sets (2)

Create

Policy Set	Sync with	Last Update	Last Set as Active
Zoom Video Conferencing System 164358495... Active	Palo Alto Networks NGFW	Jan 30, 2022, 16:16	Jan 31, 2022, 14:44
444	Palo Alto Networks NGFW	Nov 16, 2021, 17:07	Nov 16, 2021, 17:07

ACL Rule Sets (0)

You don't have an ACL Rule Set yet. Create an ACL Rule Set based on the Network Behavior of this Profile.

Create

When integrating IoT Security with Cisco ISE, you can send ISE automatically generated ACL rule sets for IoT devices. For information about providing ISE with access control lists for IoT devices, see Apply Access Control Lists through Cisco ISE.

Create a Policy Set in IoT Security

IoT Security provides the automatic generation of policy rule recommendations to control IoT device traffic. The recommendations are based on the network behaviors of all the high-confidence IoT devices in the same device profile in your local network environment as well as that of devices in the same profile in other IoT Security tenant environments.

> *High-confidence devices are those whose identity IoT Security is highly confident about and has calculated a confidence score of 90-100%. IoT Security has three confidence levels based on calculated confidence scores: high (90-100%), medium (70-89%), and low (0-69%).*

After allowing sufficient time for IoT Security to collect the full behaviors of IoT devices in a profile, you're ready to create a set of policy rule recommendations for it.

> *From PAN-OS 11.1, there's a different process for recommending Security policy rules to next-generation firewalls from that described here. The following workflow remains applicable to firewalls running PAN-OS versions prior to PAN-OS 11.1.*

STEP 1 | Log in to the IoT Security portal and select **Assets** > **Profiles** > *profile_name* > **Behaviors**.

STEP 2 | Review the data on the Behaviors page, choose **Outbound Behaviors**, and then click **Create Policy**.

For a description of the content on the Behaviors page for a device profile, see Device Profile Behaviors.

> *You can also create a policy set by navigating to the Profiles page, hovering your cursor over a profile name, and then clicking **Create Policy Set** in the information pop-up that appears.*

STEP 3 | Read the introduction to the creation of a Security policy rule set that IoT Security can recommend to next-generation firewalls and then click **Next**.

STEP 4 | Select the recommended policy rules to include in the policy set.

IoT Security automatically generates a list of policy rule recommendations. These are based on common applications that devices in the same device profile in multiple IoT Security tenant environments use and on unique applications that only devices in your environment have used during the last month (note the **1 Month** label to the right of the breadcrumbs at the top of the page). The rule recommendations are organized by application with all rules selected by

default. Clear any you don't want to use based on your organization's policy and practice as well as the information provided.

In addition to the automated policy rules that IoT Security generates based on observed network behaviors of devices in the same profile, you can manually add other rules to the set. In the policy set creation workflow, click **Add Rule** and then set an application and destination. By default, **Any** appears in both the Application and Destination Type fields. To change the application, delete **Any** and start typing the application for which you want to create a rule until autocompletion provides enough letters to select it. To set a destination, first choose a destination type: Destination Profile (for internal destinations), FQDN, IP, or Netmask. Then choose one or more destination profiles from the list or enter one or more FQDNs, IPv4 or IPv6 addresses, or netmasks. When done, **Create** the rule.

The App Risk column contains the risk level for this application as defined in Applipedia. Risk levels are graded from 1 to 5, with numbers approaching 5 carrying increasingly more risk. Hover your cursor over the application name to display a pop-up panel with information about the application retrieved from Applipedia. For explanations about this information, see IoT Device Applications Discovery.

The Alert Raised column provides the number of alert instances involving each application that occurred on devices in the source profile. This information is useful when deciding whether to include recommended behaviors in a set of policy rules. For example, if you notice a behavior is associated with a high number of alerts, you might delay adding a rule that permits this behavior until you investigate how serious the alerts are. If they are all low severity alerts, you

might decide they're acceptable. On the other hand, if they are high or critical severity alerts, you might decide to resolve them first before proceeding.

By default, IoT Security recommends permitting IoT devices in the source profile to use applications with all destinations detected in observed network traffic. This is indicated by Any in the Destination column. If you don't want to allow certain destinations, click **Any**, toggle off **Allow any destination**, clear these destinations from the list, and then close the Select Destination panel.

STEP 5 | Use the automatically generated policy rules configuration or modify it as necessary.

Use the default policy set name or enter your own. Optionally add a description for future reference.

If you want to reduce the number of policy rules that IoT Security generates, enable **Automatically condense policy rules by grouping applications**. When multiple rules have

different applications but everything else is the same—the same destination or set of destinations and, if configured, the same tags, security profiles, source and destination zones, and services—IoT Security gathers them all into a single rule and puts all the applications that had previously been the only differentiating element in the rules into a single list of applications. For example, if this option isn't enabled (its default state) and there's one destination for ten different applications, IoT Security creates ten rules. However, if you enable this option, IoT Security creates just one rule that includes a set of ten applications.

IoT Security always groups destinations together to reduce the number of recommended policy rules. Unlike the application group option, it doesn't require you to enable it.

Optionally apply tags, a security profile, source and destination zones, and services so that they become part of the policy rules when the Panorama or firewall administrator imports them. This saves the administrator from having to edit imported rules to apply them later. Select the rules you want to apply these to and then click **Tags**, **Security Profile**, **Source Zones**, **Destination Zones**, or **Services** at the top of the page to see your choices. Create or select previously-defined options and then click **Apply** or **Create**. You can apply one or more tags, source zones, destination zones, and services to the same application.

By default, an application uses its standard port and displays application-default in the Services column. When you edit a service, the Edit Services dialog box shows any non-standard ports that IoT Security has observed an application using plus the two options service-http and service-https. Select the service to use in the rule and click **Create**.

STEP 6 | Carefully review the rule set and then, when you're satisfied with it, **Create** the recommended policy rules set.

When reviewing the policy set, notice that IoT Security displays the default service ports for many of the allowed applications. These are the service ports that the selected applications

have been using on the network during the past month. If an application hasn't been observed in over a month, its service ports will no longer show up in the list.

IoT Security learns the service ports for applications by observing network traffic. Be sure to allow it enough time to collect the session data it needs, keeping in mind that IoT Security needs more time for applications that are used less frequently.

After you click Create, IoT Security creates and saves the policy set. You can view all the policy sets that you created for a device profile on the Policy page. IoT Security also prompts you to activate the policy set, which is necessary to make it available for Panorama and individual firewalls to import.

STEP 7 | To activate a policy set to make it available for Panorama and individual firewalls to import, click **Activate Policy Set**.

 A device profile can have only one active policy set at a time.

If you see anything you want to change before activating the policy set, click the **More Actions** icon (**...**) and then click **Edit**. IoT Security returns to the first page (Select Policies) so that you can make the changes.

From the same More Actions menu, you can download the policy set as a spreadsheet and you can delete it.

To save the policy set without activating it, navigate to any other page in the IoT Security portal.

Import a Policy Set into Panorama

STEP 1 | Log in to your Panorama management server and navigate to **Panorama** > **Policy Recommendation** > **IoT**.

When you do, Panorama fetches the latest active recommendations from the IoT Security cloud. If you already have the Policy Recommendations page open when you activate a policy set in IoT Security—or modify or deactivate an existing active policy set—then you must refresh the page to see the changes. Neither Panorama nor the firewalls cache any policy recommendations.

STEP 2 | Click **Import** and import the policy rule recommendations to either the pre-rulebase or post-rulebase and then select the rule to place the imported rule after.

> *Pre-rules are rules written in Panorama that are added before the rules defined locally on a firewall. Post-rules are rules written in Panorama that are added after rules defined on a firewall.*

If you don't select a rule, Panorama places the imported policy recommendations at the top of your rulebase.

> *So that any other Security policy rules for the same devices as those in the recommended rules do not occlude them, position the recommended rules before the others in the rulebase.*

STEP 3 | Click **OK**.

The import operation automatically creates the supporting objects a policy rule requires—device objects, service objects, address objects—and then it creates the policy rule itself.

You can either apply a log forwarding profile to each policy rule manually or—before importing the rule recommendations—create a log forwarding profile and name it "default" to have it applied automatically. See the section about log forwarding profiles in Prepare Your Firewall for IoT Security and also Configure Policies for Log Forwarding.

STEP 4 | Commit the configuration change.

> *For more information about importing a policy set into Panorama (and directly into firewalls), see Configure Device-ID.*

Restrict Network Access

Although policy recommendations enforce trusted behaviors for IoT devices, they only take effect when device behavior changes. However, if IoT Security detects elevated risk on a device, perhaps caused by business-critical devices running obsolete operating systems, and you want to take preventive action before an exploit is launched, you need to take a different approach from behavior-based policy rules.

IoT Security provides another option that lets you restrict network access to a specific IoT device or group of IoT devices that have the same issue, such as those susceptible to or suspected of compromise.

To accomplish this, first create a Security policy rule in which Source Device is any device whose category is "Restricted" and the action in the rule is Deny. Position this rule above all other device-based rules in the rules list. Otherwise, there's a chance that a rule based on the profile attribute, or on some other attribute, will occlude it. Similarly, make sure the "Restricted" rule is above *any* rule that might occlude it, even those not using Device-ID.

Then, in the IoT Security portal, enable the network traffic restriction feature but don't use it to restrict access yet. Notice that firewalls won't apply the new rule because none of the IP address-to-device mappings have a category attribute that matches "Restricted".

When you restrict network access for one or more devices, IoT Security immediately changes the category attribute for them from their real device categories to "Restricted" and sends firewalls new IP address-to-device mappings for them. When traffic reaches a firewall from a device with the "Restricted" category attribute, it applies the security rule you created, denying it access to the network.

📄 *Although the accompanying illustrations show how a firewall enforces a "category=Restricted" rule instead of another device-based Security policy rule, it's not necessary for the other rule to be device based. You can also restrict network access for an IoT device even when a firewall permits its access based on source IP address, service, application, or any other factor or combination of factors.*

Later, after the security issue is resolved, you derestrict devices, which returns the IP address-to-device mapping for them to their previous categories. As a result, their category attributes no longer match the "Restricted" rule and the devices will be permitted to access the network as determined by other rules.

Notes:

- To support Device-ID and IP address-to-device mappings, firewalls must be running PAN-OS 10.0 or later. To support the traffic restriction feature, firewalls must have device dictionary file 16-253 or later. Both the PAN-OS software version and device dictionary version appear in the General Information section on the PAN-OS web interface Dashboard.

- Traffic restriction is only applicable for devices with a high identity confidence score of 90 or above.

    A confidence score indicates the level of confidence IoT Security has in its identification of a device. IoT Security has three confidence levels based on calculated confidence scores: high (90-100%), medium (70-89%), and low (0-69%).

- This feature restricts network access but doesn't completely quarantine a device. Depending on network design, a restricted device can still access those parts of the network it can reach without traversing a firewall.

- Only an IoT Security user with owner privileges can enable and disable the feature.

STEP 1 | Configure a Security policy rule that denies traffic from any device whose Device-ID attribute for Category is "Restricted".

> 📋 *These instructions explain how to configure a security policy rule in the PAN-OS web UI. You can also configure it through Panorama.*

Log in to the web UI on your firewall, click **Policies** > **Security**, and then click **Add** to create a new Security policy rule. On the General tab, enter a name for the rule such as `Restrict IoT network access`.

On the Source tab, click **Add** in the Source Device section and then click **Device**. In the Device Object dialog box that appears, enter a name, choose **Restricted** for Category, and then click **OK**.

Select the device object you just created as the source device and select **Any** for the source zone and address.

On the Destination tab, select **Any** for the destination zone, address, and device.

On the Actions tab, choose **Deny** as the action. If the firewall forwards logs to Strata Logging Service, Panorama, or some other external log server, choose a log forwarding profile. Even for a rule that denies traffic, logs provide visibility into what the restricted device was attempting

to connect with and are useful during remediation. Click **OK** to save the Security policy rule configuration.

Move the rule above other policy rules.

STEP 2 | Enable traffic restriction in the IoT Security portal.

Log in to the IoT Security portal with owner privileges, click **Policy Sets** > **Settings**, and then toggle **Restrict device traffic via firewall policy**.

The following user roles *have IoT Security owner privileges: account administrator, app administrator, instance administrator, and owner.*

Policy Sets > Settings

Settings

Restrict device traffic via firewall policy

A pop-up panel appears. Read how traffic restriction works and then click **Next**.

Enable Traffic Restriction

Step 1/3: Understand How it works

When you restrict device traffic, the value of its Device-ID attribute for Category changes to "Restricted". A PAN-OS device-based policy rule (created in the next step) then restricts traffic from this device based on its new attribute value.

Cancel   Next

Select **I have created the policy** and then click **Next**.

Enable Traffic Restriction

Step 2/3: Create a PAN-OS policy rule to restrict traffic

Follow the instructions to create a PAN-OS device-based policy rule on your firewall that restricts traffic on devices whose Device-ID attribute for Category is "Restricted".
☑ I have created the policy.

Cancel   Next

Read where to restrict traffic in the IoT Security portal and then click **Enable**.

Enable Traffic Restriction

Step 3/3: Start Restricting Traffic

Follow instructions to restrict traffic for one or more devices that have a vulnerability (Risks> Vulnerabilities> vulnerability-name), or for a device with a security alert (Alerts> Security Alerts> alert-name), or for a single device on its Device Details page.

Cancel   Enable

STEP 3 | Restrict IoT devices.

As stated in step 3/3 of the Enable Traffic Restriction panel, there are three places in the IoT Security portal where you can restrict network traffic: vulnerability instances on a Vulnerability

Details page, a Security Alert Details page, and a Device Details page. Each place, or point of restriction, is described below.

> *Although only an owner can enable and disable the ability to restrict network traffic, either an owner or an administrator can use the feature to impose a restriction on a device or release one from restriction. For more information about user roles, see* Create IoT Security Users.

Vulnerability Instance as the Point of Restriction

To restrict one or more IoT devices on the Vulnerability Details page, click **Risks** > **Vulnerabilities** and then click a vulnerability name.

If the Confidence Level column is hidden, click the Columns icon (▥) and select it. Select one or more vulnerability instances with a high confidence score of 90 or above and then click **More** > **Restrict Traffic**.

Review the list of vulnerable or potentially vulnerable devices whose traffic will be restricted, optionally add a note for future reference, and then click **Confirm**.

The entry for this device in the Restricted Traffic column changes from No to Yes, indicating that its traffic is being restricted. If you don't see the Restricted Traffic column, click the Columns icon (▥) and select **Restricted Traffic**. A new entry appears in the Vulnerability Responses column. Hover your cursor over the entry to see a history of actions taken.

The Device Details page for the traffic-restricted device adds a *Restricted Device* label next to the device name. If you hover your cursor over the label, a pop-up appears with the time and point of restriction and a link to a vulnerability, security alert, or device details page. In this

case, it would be a link to a Vulnerability Details page. The pop-up also includes any notes you made.

Security Alert as the Point of Restriction

To restrict an IoT device with a specific security alert, click **Alerts** > **Security Alerts** and then click an alert name. On the Alert Details page, click **Action** > **Restrict Traffic**.

If the confidence score of the impacted device is below 90, the following message appears. The confidence score appears in the Impacted Device section on the Alert Details page.

If the confidence score is 90 or above, the Restrict Traffic dialog box appears.

Review the device whose traffic will be restricted, optionally add a note for future reference, and then click **Confirm**.

A new label appears at the top of the Alert Details page stating `Traffic Restricted Yes` and a new entry appears in the Alert Events column.

The Device Details page for the traffic-restricted device adds a *Restricted Device* label next to the device name. When you hover your cursor over the label, a pop-up appears with the time you started restricting traffic; a link to the point of restriction, which in this case would be to a Security Alert Details page; and any notes you made.

IoT Device Details as the Point of Restriction

To restrict a single IoT device on the Device Details page, click **Devices** and then the name of one of the devices in the inventory table. In the Identity section at the top of the Device Details page, click the **Action** icon (three vertical dots) > **Restrict Traffic**.

Check that the device whose traffic will be restricted is correct, optionally add a note for future reference, and then click **Confirm**.

The IoT Security portal adds a *Restricted Device* label next to the device name on the Device Details page. When you hover your cursor over the label, a pop-up appears with the time you

started restricting traffic; a link to the point of restriction, which in this case would be to the same Device Details page you're already on; and any notes you made.

On the Devices page, the entry for this device in the Restricted Traffic column changes from No to Yes, indicating that its traffic is being restricted. If you don't see the Restricted Traffic column, click the Columns icon (▥) and select **Restricted Traffic** in the Traffic section.

	Status	Risk	↓	Confidence Score	Device Name	Profile	IP Address	MAC Address	Restricted Traffic
☐	●	⌁ 49		◯ 97	Polycom_64167f75a4ea	Polycom Video Conferencing Device	10.193.2.143	64:16:7f:75:a4:ea	Yes
☐	◀▶	⌁ 45		◯ 97	0c:c4:7a:6f:1d:d4	PRTG Network Monitor	10.0.2.43	0c:c4:7a:6f:1d:d4	No
					raspberrypi			:27:eb:58:6c:70	No

Inventory (21)

STEP 4 | View all restricted devices.

On the Policy Sets page, click the number of restricted devices displayed in the Overview panel.

The Devices page opens with a filter applied to show only restricted devices in the inventory table.

STEP 5 | After investigating and remediating a traffic-restricted device, derestrict traffic for it.

To derestrict traffic for a device, repeat the same process as you did to restrict traffic but click **Derestrict Traffic**.

You can derestrict multiple vulnerability instances in bulk. Select one or more instances on the Vulnerability Details page and then click **More** > **Derestrict Traffic**.

For other traffic-restricted devices, view the inventory on the Devices page with the Restricted Traffic filter applied. Then click device names one by one to open the Device Details page for each one and click the **Action** icon (three vertical dots) > **Derestrict Traffic**.

> *To disable the feature completely, click **Policy Sets**, toggle off **Restrict device traffic via firewall policy**, and then **Confirm** the action. When you do, IoT Security cancels all existing device traffic restrictions. It also changes the entries in the Vulnerability Response column (Risks > Vulnerabilities > vulnerability_name) and Last Action column (Alerts > Security Alerts) for these devices to* `Device was derestricted`*.*

Medical IoT

IoT Security provides a means to monitor how its supported categories of medical IoT imaging devices and infusion systems are utilized and if vendors have issued recalls on medical equipment in your network. It also supports the upload of MDS2 files, which it uses to discover vulnerabilities and raise security alerts for medical IoT devices.

The IoT Security portal only displays the Utilization and Biomed dashboards and the MDS2 page when the portal theme is Medical IoT Security. It only shows the Recalls page when devices in your inventory have been recalled.

- Biomed Dashboard
- Utilization Dashboard
- Utilization Dashboard Filters
- Utilization Information Panels
- MDS2
- MDS2 Community
- Recalls

Biomed Dashboard

IoT Security gathers statistics about medical IoT devices it's monitoring, assesses their risk, and displays its findings on the Biomed dashboard. You can leverage this data to track medical device inventory and utilization as well as evaluate and address the risk of medical IoT devices.

To view the Biomed dashboard, make sure **Medical IoT Security** is the activated vertical theme for your portal and then select **Dashboard** and choose **Biomed** from the **Manage Dashboards** drop-down list.

Medical IoT

The dashboard is organized into three broad sections. At the top is a set of filters for sites and time ranges. Directly below that is the Medical Assets section, which has a high-level summary of medical device information and two panels showing top medical device categories and medical device utilization. At the bottom of the dashboard is the Compliance Risk section, which has several panels showing potentially risky types of medical devices.

Medical Assets

At the top of the Medical Assets section is a list of totals for all medical IoT devices, new medical IoT devices, their vendors, and those medical IoT devices with MDS2 forms. To provide context for these numbers, the totals for all devices, subnets, and sites in the network are also provided.

In more detail, the high-level summary contains the following device statistics:

- **Total Medical Devices**: This is the total number of medical IoT devices whose traffic was detected on the network at the sites and during the time range set on the Biomed dashboard. Clicking the total opens the **Assets** > **Devices** page to show entries for all medical devices detected within the defined site and time-range filters.

- **New Medical Devices**: This is the number of medical devices that IoT Security discovered at the specified sites and within—but not before—the specified time range. Clicking the total opens the **Assets** > **Devices** page to show just entries for the medical devices discovered within the defined site and time-range filters.

- **Total Vendors**: This is the number of vendors for medical devices referenced in Total Medical Devices.

- **Devices with MDS2**: This is the number of medical devices for which IoT Security has an MDS2 form.

- **Total Devices**: This shows the total number of devices on the network as determined by the sites and time range filters set on the Biomed dashboard and the global filter for device type set on another page such as Devices. Clicking the number opens the **Assets** > **Devices** page to show device entries matching defined filters for site, device type, and time range.

- **Medical Subnets**: This is the total number of subnets containing medical devices as determined by the site and time range filters set on the Biomed dashboard and the global filter for device type set on another page such as Devices. Clicking the number opens the **Networks** > **Networks and Sites** > **Networks** page.

- **Total Sites**: This is the absolute total number of sites for the tenant regardless of the current site and time range filters set on the Biomed dashboard and the global filter for device type set on another page. Clicking the number opens the **Networks** > **Networks and Sites** > **Sites** page.

- **Improve Device Visibility Coverage**: This button opens the Data Quality Diagnostics page (**Administration** > **Data Quality**). There you can see the quality of data that IoT Security is receiving. In particular, the page focuses on IP endpoints and low-confidence devices, how they can lower data quality, and ways to reduce their numbers through improved network coverage.

The two panels in the Medical Assets section contain information about the main categories of medical IoT devices and their utilization:

- **Top Medical Device Categories**: This panel lists the medical device categories and ranks them by device count, with those that contain the largest number of medical devices at the top. Clicking a category name opens a new browser window displaying the **Assets** > **Devices** page

filtered to show just entries matching this category. Clicking **View All** in the lower right opens the **Assets** > **Devices** page filtered to show all medical devices.

- **Medical Device Utilization**: This panel shows all medical IoT categories and how the devices in each one are being utilized. A bar chart shows the percentages of time the devices in a category are detected in use, online but not in active use, and offline. Hover your cursor over a bar to see a pop-up with numbers for each kind of utilization. Click a medical device category to open the **Assets** > **Devices** page in a new browser tab or window. The page is filtered to show devices in the selected category.

Compliance Risk

This section of the dashboard shows information about medical IoT devices that affect their risk exposure.

- **Top End-of-Life Operating Systems**: These are devices running an OS version that the vendor no longer supports with patch updates, making them more vulnerable to attack. The table shows how many device profiles have devices with an end-of-life OS and the percent of affected medical devices relative to all medical devices. Clicking a number in the Devices column opens the **Assets** > **Devices** page filtered to show only devices running this operating system and version.

- Devices with various risk factors are listed. For each one the total number of devices with this risk and their percent relative to all medical IoT devices are shown. Clicking **View All** for the first three opens the **Assets** > **Devices** page with a filter to show just these devices. Clicking **View All** for Devices with FDA recalls opens the **Vulnerabilities** > **Recalls** page.

Devices with outdated endpoint protection: These are devices that have endpoint protection, such as anti-virus protection, but they haven't communicated with their vendor and haven't been updated in over a month. This makes them vulnerable to new types of attacks released since their last update.

Devices without endpoint protection: These are devices that do not have any endpoint protection installed on them.

Devices with PHI: These devices contain personal health information (PHI).

FDA Recall Instances: This shows the total number of devices that have been issued a recall order by the Food and Drug Administration (FDA) because of a product flaw that affects safety and requires it to be fixed or replaced.

Utilization Dashboard

When the IoT Security portal theme is Medical IoT Security, you can see the Utilization dashboard. IoT Security gathers utilization statistics and metrics about the medical IoT devices it is monitoring and displays them on this dashboard. You can then leverage this data to minimize device downtime, reduce total cost of ownership (TCO), and increase revenue through better capital planning. In addition to minimizing downtime and maintenance, you can also use the gathered data to identify unused assets (possibly broken or misplaced) and ensure safe and secure device disposal at the end of the IoT device life cycle.

Make sure the application content version on your firewalls is 8367-6513 or later; that is, the major version, which is identified by the first four digits, is 8367 or above (8368, 8369, 8370, and so on), starting from 8367-6513. These versions include healthcare-specific applications that allow IoT Security to discover medical equipment and provide utilization data. They also allow firewall security policy rules to include healthcare-specific applications.

To view the Utilization dashboard, select **Dashboards** and then choose **Utilization** from the **Manage Dashboards** drop-down list.

The dashboard is organized into two broad sections. At the top is a set of filters for sites, medical IoT device categories, and time ranges that control what appears on the page. Below the filters are information panels that display various types of information about how medical IoT devices are being utilized.

In addition to viewing the dashboard in the IoT Security portal, you can download its data as an Excel spreadsheet. Click the **Download** icon to the right of the top filters, set the filters to include the data you want to save, and then click **Download**.

Download

Choose from the options below and download details for your devices

Site
All Sites | X-Ray Machine

Time Range
3 Months (Nov 16 - Today)

Cancel | **Download**

IoT Security creates an Excel file with the details you specified on multiple tabs and makes it available for download.

Medical IoT

site name	mac address	ip address	hostname	category	profile	last activity	currently in use
Healthcare	90:1b:0e:26:ff:fa	10.163.23.10	XP1028101323459	X-Ray Machine	Siemens Fluoroscopy Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	e4:02:9b:80:48:96	10.23.176.126	WISP-	X-Ray Machine	GE C-Arm	2024-02-11T16:33:06.889Z	yes
Healthcare	b0:b9:8a:6d:06:5b	10.44.178.26	LOGIQe	UltraSound Machine	GE UltraSound Machine	2024-02-14T16:34:18.743Z	
Healthcare	04:92:26:8a:77:22	10.44.178.145	DESKTOP-7KGV839	X-Ray Machine	DICOM-Imager-X-Ray	2024-02-11T16:33:06.889Z	
Healthcare	04:f0:21:0c:28:f6	10.23.176.105	SJHSXRM2333	X-Ray Machine	GE X-Ray Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	04:f0:21:4b:76:f1	10.23.182.64	SJS_GEOPTIMA_82	X-Ray Machine	GE X-Ray Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	00:25:90:f3:44:7c	10.163.23.192	TOMO01_SJHS	X-Ray Machine	Hologic Mammography Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	90:1b:0e:17:5f:d7	10.163.23.11	10502200-60404	X-Ray Machine	Siemens Fluoroscopy Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	00:90:fb:65:90:f0	10.163.23.65	WISP-	X-Ray Machine	GE C-Arm	2024-02-11T16:33:06.889Z	yes
Healthcare	04:f0:21:85:5f:f3	10.23.176.107	SJS_GEOPTIMA_84	X-Ray Machine	GE X-Ray Machine	2024-02-14T16:34:18.743Z	
Healthcare	c8:d3:ff:ba:4c:00	10.160.214.200	minint-3f2eelg	MRI Machine	Siemens MRI Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	88:b1:11:d6:06:14	10.23.176.125	WISP-	X-Ray Machine	GE C-Arm	2024-02-14T16:34:18.743Z	yes
Healthcare	3c:52:82:6f:c2:25	10.163.23.24	minint-43mgoop	MRI Machine	Siemens MRI Machine	2024-02-14T16:34:18.743Z	yes
Default Site	d4:85:64:b9:9b:00	10.10.37.224	X-Ray DR6000	X-Ray Machine	GE X-Ray Machine	2024-02-14T16:34:18.743Z	
Healthcare	c8:d3:ff:bc:2d:ff	10.163.23.189	Admin-PC	X-Ray Machine	Hologic Bone Densitometer	2024-02-14T16:34:18.743Z	yes
Healthcare	00:0e:8e:63:23:9d	10.23.180.62	SURGERY-2LTAIII	X-Ray Machine	GE C-Arm	2024-02-14T16:34:18.743Z	yes
Healthcare	18:60:24:ae:81:22	10.163.23.81	SJS_XR646_81	X-Ray Machine	GE X-Ray Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	f0:03:8c:99:bd:fc	10.23.178.100	SJHSUSM2353	UltraSound Machine	GE UltraSound Machine	2024-02-14T16:34:18.743Z	
Healthcare	00:0b:ab:c2:f6:42	10.163.23.53	SJHSUSM2353	UltraSound Machine	GE UltraSound Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	a0:42:3f:29:e9:2a	10.163.23.136	minint-g34bt9i	X-Ray Machine	Siemens Angiography Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	48:0f:cf:4b:fb:9f	10.163.23.74	SJSD63074	Nuclear-Medicine Imager	GE Nuclear-Medicine Imager	2024-02-14T16:34:18.743Z	yes
Healthcare	ac:1f:6b:1e:47:57	10.163.23.70	SJS_VOL_03	UltraSound Machine	Volcano UltraSound Machine	2024-02-14T16:34:18.743Z	
Healthcare	00:19:99:ec:d9:46	10.163.23.102	CTAWP66611	CT Scanner	Siemens CT Scanner	2024-02-14T16:34:18.743Z	
Healthcare	3c:52:82:5f:3f:7a	10.163.23.101	minint-g20usvj	MRI Machine	Siemens MRI Machine	2024-02-14T16:34:18.743Z	yes
Default Site	d4:85:64:b9:9b:81	10.10.37.224	X-Ray DR6000	X-Ray Machine	GE X-Ray Machine	2024-02-14T16:34:18.743Z	yes
Healthcare	90:1b:0e:ea:12:c2	10.163.23.28	SJHSCT28	CT Scanner	Siemens CT Scanner	2024-02-14T16:34:18.743Z	yes
Healthcare	00:13:95:25:9a:fd	10.163.23.107	LS817091575A	UltraSound Machine	GE UltraSound Machine	2024-02-14T16:34:18.743Z	yes

Utilization Dashboard Filters

At the top of the dashboard are filters for sites, medical IoT device categories, a time range (1 week, 1 month, or 3 months), and device profiles. The filters determine the data that appears throughout the dashboard.

These four filters control what data appears in the Utilization dashboard:
· Site
· Device category
· Time period
· Device profile

Sites: The choice of site filters includes **All Sites** and one or more individual sites. IoT Security site filters provide great flexibility by letting you combine multiple selections.

Medical IoT device categories: The contents of this list are dynamically determined by the devices discovered on your network and are listed in alphabetical order. When you initially navigate to the Utilization dashboard, it uses the filter for whatever device category comes first alphabetically. If you change the category filter, navigate away, and then return to the Utilization dashboard, it remembers your previously chosen filter and continues to use it.

The following are the supported medical IoT device categories that can appear as filters based on whether such devices are found on your network:

- CT scanner
- Infusion System
- MRI Machine
- Nuclear-Medicine Imager
- PET Scanner
- UltraSound Machine
- X-Ray Machine

Time range: The time range filters for the Utilization dashboard consist of 1 Week, 1 Month, and 3 Months, referring to the last seven days, last 30 days, or last 90 days. When you initially navigate to the Utilization dashboard, it inherits the time filter set on another page or dashboard. If the time filter is not 1 Week, 1 Month, or 3 Months, the inherited filter is still displayed but the contents on the dashboard are set for 1 month.

Together with the filters for sites and medical IoT device categories, the time filter determines the scope of data in the information panels. However, the device total shown in the device profile filter is always for the past year regardless of the time range filter.

Device profiles: Below the page title and top filters bar is a panel with the total number of medical IoT devices in the selected device category during the past year and the device profiles within that category to which devices belong. These profiles are additional filters that allow you to zoom in on utilization details from the broader device category level to individual device profiles.

By default, the Utilization dashboard displays device profiles in order from those with the most devices to those with the least. To list them alphabetically, click **Device Count** > **Profile Name**.

In addition, the Utilization dashboard displays data for all device profiles within the category encompassing them by default. To filter it further, select a device profile.

417

Utilization Information Panels

The Utilization dashboard contains various information panels. The types of panels differ depending on the medical IoT device category filter you select.

Data does not immediately appear in the Utilization dashboard. It requires a minimum of 24 hours to collect enough data to populate the information panels with meaningful data.

Trend – The Trend information panel displays graphs that help you spot trends in the way your devices are being used.

For medical imaging devices, the Trend panel shows two graphs. The line graph shows the number of images taken at intervals throughout the period set as the time filter. The bar graph shows the total case studies created during the same period. At a glance, you can see patterns of activity and any periods of lulls and spikes. Hovering your cursor over a data point shows the number of images and case studies at that point. If you only want to see one chart or the other, click **Images** or **Case Studies** in the upper right corner of the Trend panel to show or hide them. If you want to focus attention on one graph but not completely hide the other, hover your cursor over **Images** or **Case Studies** to cause the other one to fade.

For infusion systems, the Trend panel displays a line graph to track the number of systems in use and the number connected to the network (online) but not necessarily in use. If all connected infusion systems were both connected and in use, the two lines appear to be a single line because all their data points align. However, if you hover your cursor over **Devices Used** or **Devices Online Only** in the upper right corner of the Trend panel, you can show or hide one line or the other. Hover your cursor over a data point to see the number of devices that were used and those that were connected at that point.

The number that appears to the left of the chart is not an overall total. It shows how many infusion systems were used today, using the most recent time for which a total can be calculated.

Imaging Scan Analysis – This panel summarizes the sections of the human body that were scanned by imaging devices. It is shown for all DICOM devices except ultrasound equipment, which doesn't identify scanned body parts in its traffic.

The panel is divided into two sections. On the left is a human figure consisting of four major anatomical regions:

- Head and Neck
- Upper Extremity
- Trunk
- Lower Extremity

There is a fifth grouping called "Other". This is for scans that are unidentifiable.

The coloring of each body region represents the volume of scans performed on it. The darker a section is the more scans were done there.

On the right are five bar charts, one for each of the four major anatomical regions and a fifth for the Other grouping. Chunking the data in this way makes it easier to see how your imaging equipment is being used. The bars in each chart represent the number of scans for more specific body areas (for example, a bar for the more specific Mouth and Throat is within the bar chart for Head and Neck). The charts include bars for areas where the most scans were performed. To see a full list, hover your cursor over one of the main sections and a popup appears.

Trunk Analysis (5 Profiles)	Images
Skin & Breast	213
General Anatomical Regions	25
Axial Skeleton, Except Skull	24
Non-Axial Upper Bones	2
Non-Axial Lower Bones	2

Devices – This information panel shows the total number of devices on the network in the past year and how many devices were connected to the network and used, connected but not used (online only), and disconnected (offline) during the filtered time range.

- **Used** – These are devices that matched all the filters and were not only on the network but sent traffic with a protocol indicating that they were being used.

- **Online Only** – These are devices that matched all the filters and were detected on the network but did not send traffic that indicates they were being used.

- **Offline** – These are devices that IoT Security detected on the network within the past year but not during the time range set in the time filter.

Devices

Used (19)	Online Only (3)	Offline (0)				
Device Name ↑	Images	Total Scan Time(min)	Case Studies	Site	Location
d8:50:e6:57:d6:ba	2	--	1	Healthcare	
GE X-Ray Machine	1.1K	--	6	Default Site	1100
surgery30nw	7	--	2	Healthcare	
X-Ray Machine	19	--	5	Healthcare	

View all 19 devices   16 – 19 of 19 ❮ ❯

22

● 19 Device Used
● 3 Device Online Only ⓘ
● 0 Device Offline

Click a section of the donut graph or a tab heading to switch lists. Click **View all** <number> **devices** to open the Devices page with a filter set to show only the devices in the active list. Click a specific device name to view device details for it.

When you click an entry in the Device Name column, the Device Details page for it opens. While viewing the Device Details page, click **Utilization** to see more detailed information.

The Utilization section on the Device Details page for any imaging device such as an ultrasound machine, X-ray machine, or CT scanner, shows when that device is and is not being used, how it's being used, and when it communicated with the device vendor. For example, with the time filter set to 1 Month at the top of the Device Details page, the Device Usage by Hours information panel shows the times when the device was in use during the past month.

The legend in the upper right explains how the colors indicate how much the device was used at each hour. The darker the green, the more it was used. You can also hover your cursor over a square to see a tooltip explaining how often the device was used at that time. For example, in the image shown above, the device was used three times on Wednesday during the 2:00 PM hour.

> 📋 *The legend is dynamic based on which time range you selected: for 1 Week, it's 0–1, for 1 Month it's 1–5, and for 1 Year it's 0–40+.*

In Device Usage by Number, you can see the number of images the device took–or, by clicking **Case**, the number of cases for which the device took images. In Vendor Activities, you can see when the device communicated with its vendor, which for many devices is an automated means for obtaining software and security updates. You can see the traffic that the vendor initiated and total vendor traffic, which is a superset of all communication between the device and the vendor regardless of which one initiated it.

When viewing the Utilization section on the Device Details page for an infusion system, the IoT Security portal displays the following information.

The blue bars in the upper graph, show how long the device was in use. For 1 week, 1 day, or 2 hours, each blue bar indicates the number of minutes per hour (60 minutes maximum) that the device was in use over the past 168, 24, or 2 hours. For 1 month or 1 year, it shows the number of minutes per day (1440 minutes maximum) that the device was in use every day for the past 30 or 365 days.

The lower graph shows how this device compares with others in the same device category. You can see how long it was in use and the percent of time it was actively used in relation to the time range set. In the example above, it was used for 9.72 days divided by 30 days or 32.4% of the time. The line graph shows that out of 42 devices in the same category, this device was used more than 72% of the others (indicated to the left in green) and less than the remaining 28% (indicated to the right in white).

The following site-specific information panels appear only when you're using IoT Security for multiple sites.

Used Devices by Site – These bar charts show how many medical IoT devices were used at each site within the filter parameters. The number of devices that were in use during the filtered time range is shown in relation to the total number of devices at each site over the past year.

If you filter the dashboard to show data for a single site, this panel shows not only the devices that were in use at this site but also other sites that had active devices to provide a reference for comparison.

Scanning Data by Site – When the medical IoT device category is for imaging devices, this information panel shows the number of scanned images and cases for each site with active devices. Click **Cases** or **Images** to toggle between them.

If you filter the dashboard to show data for a single site, this panel shows not only scanning data for this site but also other sites that have scanning data to provide a reference for comparison.

Average Usage by Site – When you set **Infusion System** as the medical IoT device category filter, this information panel shows the average device usage at each site. The average is calculated by dividing the total number of hours that devices at a site were in use during the time filter (1 week, 1 month, or 3 months) by the total number of devices in use at that site during the same time. Any devices that were not in use during that time are excluded from the calculation. In short, hours in use / devices in use = average usage.

If you filter the dashboard to show data for a single site, this panel shows not only the average usage of infusion systems for this site but also that of other sites with active infusion systems to provide a reference for comparison.

MDS2

> Note: The **Vulnerabilities** > **MDS2** page only appears when the portal theme is Medical IoT Security.

Medical device vendors often list the security-related features of their products in Manufacturer Disclosure Statement for Medical Device Safety (MDS2) forms, which they share with their customers. Vendors issue these MDS2 documents for each version of a medical device and include valuable information such as whether a device processes PHI (personal health information); if it stores PHI and, if so, if it's encrypted; and if antivirus software is installed on the device.

Over time, healthcare providers can collect thousands of MDS2 documents for thousands of medical devices. When used as intended, MDS2 documents can greatly enhance your security posture and incident response (IR). However, absorbing the details from these documents for the specific version of the software running on their connected devices is a daunting task. As a result, MDS2 files often go unused.

IoT Security simplifies the management and use of the MDS2 files you have. If you upload an MDS2 file for a device to IoT Security, it then includes this data along with other environmental factors when assessing the risk to the device. For example, if the software version of a device specified in an MDS2 file has a known vulnerability, IoT Security more precisely identifies it as a vulnerability instead of just a potential vulnerability. IoT Security supports MDS2 files in 2004, 2008, 2013, and 2019 formats.

You can upload MDS2 files to IoT Security and use files shared by other IoT Security users through the MDS2 community. To join, select **Vulnerabilities** > **MDS2**, click **Learn More**, read about how the MDS2 community works, and then click **Join Now**. After that, IoT Security scans the community and shows previously uploaded MDS2 files from other community members that match your devices. At the same time, Palo Alto Networks security engineers review any MDS2 files that you've already uploaded. If they are approved, IoT Security then shares your files with other community members. In this spirit of cooperation, everyone benefits from the files shared with each other.

If members upload duplicate MDS2 files (that is, more than one file applies to the same vendor, profile, and model), IoT Security uses the following logic in order from the top to prioritize one over another and automatically apply it to your devices:

- If an MDS2 file is excluded, do not use it.
- Use a manually selected MDS2 file over an automatically selected file.
- Use an MDS2 file that you uploaded.
- Use an MDS2 file that's shared in the community.
- Use a version of an MDS2 file released later than another version.
- Use a later format version of an MDS2 file over an earlier format; for example, use a 2017 MDS2 file instead of a 2013 format version.

When you select **Vulnerabilities** > **MDS2** after joining the MDS2 community, IoT Security displays the MDS2 Files Matched page. This lists the MDS2 files that match medical IoT devices in the IoT Security inventory. You can navigate from here to a page with MDS2 files that you previously

uploaded, a page with files uploaded by other IoT Security customers, and a page listing medical IoT devices that match the MDS2 files here.

On **Vulnerabilities** > **MDS2**, you can view files that match medical IoT devices in the inventory, download them, and exclude them if you don't want IoT Security to apply them to your medical IoT devices. You can also download a complete list of all uploaded MDS files or a list of one or more selected files.

To upload a file, click the **Upload** icon (⬆), navigate to an MDS2 file in PDF format, and then select and upload it.

IoT Security matches the uploaded MDS2 file with devices that share the same model, vendor, and profile as those specified in the file. Although you can upload an MDS2 file on the Device Details page, IoT Security only applies the MDS2 file to that individual device. On the other hand, if you upload an MDS file on the MDS2 page, IoT Security searches its inventory for all devices with the same model, vendor, and profile attributes and applies the MDS2 file to all matched devices. Furthermore, if new devices are added to the inventory later, IoT Security will apply the MDS2 file to those devices as well.

Clicking a number in the Matched Device column opens the Devices page with a filter applied to show just those devices that match the MDS2 file.

The number in the Matched Device column on the MDS2 page is the total for all sites. If you have administrative access to device data for a subset of sites, the number of matched devices on the Devices page might be smaller than the number on the MDS2 page.

To view some details about an MDS2 file, click the entry in the File Name column. An information panel slides open on the right side of the main window listing the three attributes that IoT Security uses to map the MDS2 file to devices. Below this, it lists several key points about the device, the document, and security.

When you upload a MDS2 file, check if there are any inaccuracies among the device mapping rule values. It's possible that text alignment issues in the PDF cause characters to be parsed incorrectly. If that happens, IoT Security won't be able to match the MDS2 file with devices. In such cases, click **Edit** to the right of Device Mapping Rule, modify the text as necessary, and then click **Update**.

In addition to the values in the Device Mapping Rule, you can edit other attributes in the MDS2 file if they were parsed incorrectly as well. Whenever you click **Update**—either for changes to Device Mapping Rule or Data from MDS2 File—IoT Security immediately removes any previous matches for the MDS2 document and runs the matching process again.

To view an entire MDS2 file in PDF format, click **Show PDF** in the information panel.

To download the PDF, click the **Download** icon ( ) at the top of the PDF viewer.

To close the information panel (and PDF viewer if it's also open), either click the X in the upper right corner or click the file name again.

To download a list of all uploaded MDS2 files in a .csv file, click the **Download** icon (⬇) above the MDS2 table. To download a list of one or more MDS2 files in a .csv file, select check boxes of the ones you want to download and then click **Download**.

To delete one or more previously uploaded MDS2 files, select the check boxes of the files to delete and then click **Remove**.

MDS2 Community

If you want, you can override the automatic MDS2 file selection and apply a different file of your choice. See the last step below for instructions.

STEP 1 | View your uploaded MDS2 files and shared community files.

1. On the **Vulnerabilities** > **MDS2** page, click <number> **Files Available in Community**.

   The **Vulnerabilities** > **MDS2** > **MDS2 Files** page opens with two tabs: one for the files you uploaded (**Uploaded Files**) and one for files shared in the community (**Community Files**).

	File Name	Mapped Vendor	Mapped Profile	Mapped Model	Matched Device	Software Ver.	Format Ver.	Upload Date
☐	CX-50_4.0_MDS2_form_HN_1-2013.pdf	Philips	Philips UltraSound Machi...	CX50	1	4.0	2013	Aug 24, 2021, 17:32
☐	Carestream-All-DR-DRX-Soft... parsing				0			Aug 24, 2021, 17:30
☐	SecurView_MDS2_2016.pdf	Hologic	Hologic SecurView Work...	SecurView DX Man...	0		2013	Aug 24, 2021, 15:25
☐	Trident_MDS2_Sheet__Rev_2.pdf	Hologic	Hologic Specimen X-Ray ...	Trident HD	0		2013	Aug 24, 2021, 15:24
☐	Insight_FD_MDS2_sheet.pdf	Hologic Inc	Hologic X-Ray Machine	Insight FD	0		2013	Aug 24, 2021, 15:23
☐	Insight_MDS2_Sheet_Rev_004.pdf	Hologic Inc	Hologic X-Ray Machine	Insight FD	0		2013	Aug 24, 2021, 15:23
☐	Manufacturer_Disclosure_Statement_IMP...	Agfa	Agfa PACS Server	IMPAX	0			Aug 24, 2021, 15:23
☐								Aug 24, 2021, 15:23

   The table provides key details about the MDS2 file: the vendor, profile, model, and software version it applies to; how many devices in your inventory it matches; the format version of the MDS2 file; and when it was uploaded to IoT Security.

The table on the Community Files page provides almost all the same information as that on the Uploaded Files page except for the date the file was uploaded.

434

STEP 2 | If there are shared MDS2 files that you don't want IoT Security to apply to your devices, exclude them.

You might not want IoT Security to apply one or more community files to your medical IoT devices, perhaps because you prefer a different version for some reason. In this case, you can exclude them from the automatic selection process.

1. Select **Vulnerabilities** > **MDS2** > **MDS2 Files** > **Community Files**, select the check box of an MDS2 file you want to exclude from the automatic selection process.

   Exclude and **Download** buttons appear above the table.

2. Click **Exclude**.

   📋 *If you change your mind, repeat these steps but click **Include** instead of **Exclude**. IoT Security displays an **Include** button for previously excluded MDS2 files.*

STEP 3 | If there are duplicate MDS2 files and you want to use something different from what IoT Security automatically selected, manually override the selection.

When more than one MDS2 file applies to the same vendor, profile, and model, you might want IoT Security to apply a different MDS2 file other than the one it automatically selected. For example, you might choose an earlier format version if your medical IoT devices are

running earlier software versions and this MDS2 file applies to them more accurately than a later one would.

1. Click an MDS2 file name on the **Vulnerabilities** > **MDS2** page.

   The information panel that opens on the right side of the page lists duplicate MDS files in order of priority.
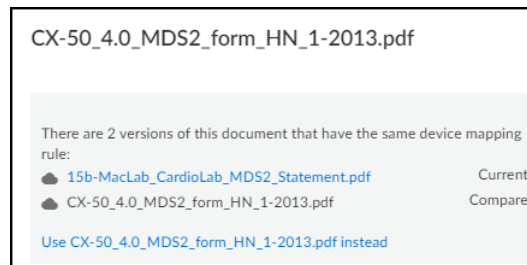

> 📋 *The information panel for a file you uploaded includes an **Edit** option, but the panel for a community-shared file does not.*

436

2. To compare two duplicate files, either click or hover your cursor over **Compare**. IoT Security compares it with the currently applied file. In the example shown here, you can see that the first MDS2 file is prioritized because it has a software version whereas the second file does not.

3. To use a different file, click its name and then click **Use** <file-name> **instead**.

The name of your selected file now appears in the File Name column and is shown as *Current* in the information panel.

Recalls

The **Vulnerabilities** > **Recalls** page lists the Food and Drug Administration (FDA) recalls for devices on your network.

> *The IoT Security portal only displays the Recalls page when the* portal theme *is Medical IoT Security and it detects that at least one of the medical IoT devices on the network has been recalled. The* application content version *on your firewalls must be 8367-6513 or later to detect healthcare-specific applications and include them in security policy rules.*

The **Recalled Devices** column shows the number of devices on your network that have been recalled, and if the status is open, the manufacturer is still accepting device returns for them.

Recall	Status	Affected	Recalled Devices	Recalled Profiles
Z-2671-2017	Terminated on 10/26/18	Alaris PC Unit, Model 8015	293	1
Z-1606-2016	Terminated on 6/29/18	Alaris PC unit, Model 8015The Alaris PC unit is ...	293	1
Z-1380-2016	Terminated on 12/13/17	EPIQ DIAGNOSTIC ULTRASOUND SYSTEM, M...	1	1
Z-0817-2015	Terminated on 8/15/16	EPIQ 7 Ultrasound System, EPIQ 7 systems wit...	1	1
Z-1632-2015	Terminated on 8/08/16	EPIQ 7 Ultrasound System versions 1.3.2 or low...	1	1
Z-1579-2015	Terminated on 2/29/16	EPIQ 7 Ultrasound System with Pediatric Cardio...	1	1
Z-0725-2014	Terminated on 5/06/14	GE Optima XR220amx and Optima XR200amx ...	1	1
Z-1407-2015	Terminated on 4/29/15	GE Healthcare Automatic Mobile X-Ray (AMX) ...	1	1
Z-1399-2012	Terminated on 1/16/14	GE Healthcare Automatic Mobile X- Ray (AMX) ...	1	1
Z-1993-2012	Terminated on 1/10/13	GE Healthcare Optima Mobile X-ray System. Th...	1	1
Z-0768-2016	Terminated on 9/07/16	GE Healthcare, Optima XR220amx, Mobile Digit...	1	1
	13/17	SOMATOM P...		1

Each recall identifier links to a URL on the U.S. Food & Drug Administration website with information about the recall. For example, clicking **Z-1380-2016** in the Recall column shown above opens the following webpage.

Whether or not you receive instructions from a manufacturer when a product is recalled, you can open the recall URL to get identifying information such as lot numbers, serial numbers, or version numbers for the recalled devices and learn the manufacturer's contact information. You can then use these numbers to locate the devices and call or write the manufacturer.

Manage IoT Security Users

Manage IoT Security users with role-based access control (RBAC).

- [Create IoT Security Users](#)
- [User Roles for IoT Security](#)

Create IoT Security Users

When users log in to the IoT Security portal using single sign-on (SSO), they go through a two-step process. In step 1, an SSO identity provider (IdP) authenticates users by verifying their credentials. In step 2, users are authorized and provided with a role to access IoT Security.

When users log in to the IoT Security portal using Palo Alto Networks SSO, their credentials are verified against user accounts in the Customer Service Portal (CSP). Then their user role is assigned according to the Identity & Access section of the hub. User roles determine what they can see and do in the portal. These user roles are referred to as "externally managed user roles" in contrast to "internally managed user roles", which are assigned in the IoT Security portal and are described in a later section.

In addition, IoT Security also provides an option to verify users against an Active Directory (AD) authentication system through SSO. In this case, user accounts are in Active Directory, which verifies user credentials on behalf of IoT Security. You can manage the role of a given user in two different ways, similar to the Palo Alto Networks SSO: (1) managed internally by IoT Security or (2) managed externally by Active Directory.

> *External roles are managed in the AD instead of the hub as done in the Palo Alto Networks SSO option.*

Because the user role can be managed in two different places, when users log in through an SSO, IoT Security might find their external roles are different from their internal roles. In such cases, whichever role is higher takes precedence.

Authenticate Users with the Palo Alto Networks SSO and Manage User Roles in the Hub

IoT Security supports role-based access control (RBAC) through App Administrator, Instance Administrator, Owner, Administrator, and Read-only roles. Creating users for the IoT Security application involves three steps:

- Create a user account in the Customer Support Portal
- Assign a user role in the hub
- (For Administrator and Read-only users) Allow access to all sites or a subset of sites

STEP 1 | Log in to the Customer Support Portal with superuser permissions, which allow you to create new user accounts.

STEP 2 | Click **Members** > **Create New User**, enter the required information, and then **Submit**.

A new user account is created and added to the account as a member. An email notification is sent to the new user with login credentials.

STEP 3 | Log in to the hub.

STEP 4 | Click the gear icon in the upper right of the hub landing page and then **Access Management**.

STEP 5 | Expand the IoT Security section in the left panel, select the IoT Security instance to which you want to assign the user, select the check box for the user account you just created, and then **Assign Roles**.

STEP 6 | Select **IoT Security** in the left panel to display the IoT Security role assignment window in the main panel.

STEP 7 | Choose one of the following roles from the Role drop-down list:

App Administrator

Instance Administrator

Owner

Administrator

Read only

STEP 8 | For information about these user roles, click **Role Definitions**.

To learn more about the App Administrator and Instance Administrator roles, which are common roles for all Palo Alto Networks apps and provide the same privileges in IoT Security as Owner, see Available Roles. To learn more about the Owner, Administrator, and Read only roles, which are specific to IoT Security, see User Roles for IoT Security.

443

Authenticate Users with an Active Directory SSO and Manage User Roles in Active Directory

STEP 1 | Prepare the authentication system.

Before you configure IoT Security, prepare your Active Directory to communicate with it and export the identity provider (IdP) metadata file that IoT Security will need to communicate with the IdP.

1. Configure your IdP with the following URLs, replacing the `tenant-id` variable with your own tenant ID, which is the first part of your IoT Security portal URL:

   `https://tenant-id.iot.paloaltonetworks.com/login`

   Depending on how you configure your IdP, either point it to the IoT Security metadata URL to retrieve all the necessary data or enter the information separately.

   - **Assertion Consumer Service (ACS)** – This is the destination to which the IdP sends authentication assertions in response to user authentication requests.

     `https://tenant-id.iot.paloaltonetworks.com/v0.3/zauth/saml2_sso/acs`

   - **Entity ID** – This is the URL that uniquely identifies the Zingbox SP.

     `https://tenant-id.iot.paloaltonetworks.com/v0.3/zauth/saml2_sso/metadata`

   - **Palo Alto Networks Metadata** – This file includes the ACS URL and entity ID plus other parameters such as its public Security Assertion Markup Language (SAML) 2.0 encryption key.

     `https://tenant-id.iot.paloaltonetworks.com/v0.3/zauth/saml2_sso/metadata`

   📋 *To see the URLs with your specific tenant ID, follow steps 1-2 in the next section and then copy the URLs in the Service Provider (SP) Configuration Details section.*

2. Either copy and save the URL where IoT Security can import the IdP metadata file from your SSO authentication system or download the file and save it in XML format. You will later import it to the IoT Security portal.

STEP 2 | Prepare IoT Security to use an externally managed SSO.

1. Log in to the IoT Security portal as an owner, navigate to **Administration** > **User Accounts**, and then **Manage SSO**.

   Palo Alto Networks is the default SSO identity provider (IdP) that authenticates users accessing the IoT Security portal and assigns user roles to them.

2. To add a user-configured SSO, **Add New SSO**, and then enter the following in the Single Sign-on Configuration dialog box that appears:

   Name: Enter a name for the SSO. It can be up to 16 characters. This name will appear on the login page as shown in the Preview below.

   Logo (Optional): Upload an image file to display next to the SSO name on the login page as shown in the Preview. The image file can be up to 2 MB and must be in .bmp, .jpg, or .png format.

   IdP Metadata: Either enter the URL of the IdP metadata file you copied and saved earlier or click **Choose file**, navigate to the XML file you exported from your authentication system, and select it.

3. **Validate the IdP metadata URL or uploaded file.**

   Validating the IdP metadata URL activates the **Save** and **Test** buttons.

4. Configure the following settings to identify AD user groups whose users you want Active Directory to authorize. If you leave them empty, IoT Security authorizes them locally.

   Attribute to get AD Groups: Enter the attribute in the SAML 2.0 response that identifies user groups from Active Directory.

   AD Group Format: Select whether the attribute is formatted as **Plain Text** or **Regular Expression**. These are how IoT Security maps AD user groups to IoT Security user roles.

   Plain Text identifies the user group with the exact value specified in **Attribute to get AD Groups**. For example, if **AD Group Format** is **Plain Text** and the **AD Group** is `Hospital Administrator`, then IoT Security maps only users in the AD group named `Hospital Administrator` to the specified IoT Security role.

   Regular Expression identifies any user group that contains the value specified in **Attribute to get AD Groups**. For example, if **AD Group Format** is **Regular Expression** and the **AD Group** is `OUI=Hospital*`, then IoT Security maps users in any AD group whose organizational unit identifier (OUI) includes `Hospital`—such as `OUI=Hospital Administrator` and `OUI=Hospital NetSec`—to one or more specified IoT Security roles.

   AD Group and **User Role**: Enter an Active Directory group name and then choose the IoT Security user role to map it to: **Owner**, **Administrator**, or **Read Only**. Click **+** to add more AD group-to-user role mappings. You can create up to 50 mappings. A single AD group cannot

map to multiple IoT Security user roles, but multiple AD groups can map to the same IoT Security user role.

> For information about the IoT Security user roles, see User Roles for IoT Security.

5. **Save** the SSO configuration.
6. **Test** the SSO configuration.

   IoT Security opens a small window to log in using the authentication system.
7. When done with the test, click **Confirm**.
8. **Enable** the SSO configuration.
9. After enabling the configuration, the **Enable** button changes to **Disable and Edit**.

Authenticate Users with any SSO and Manage User Roles in the IoT Security Portal
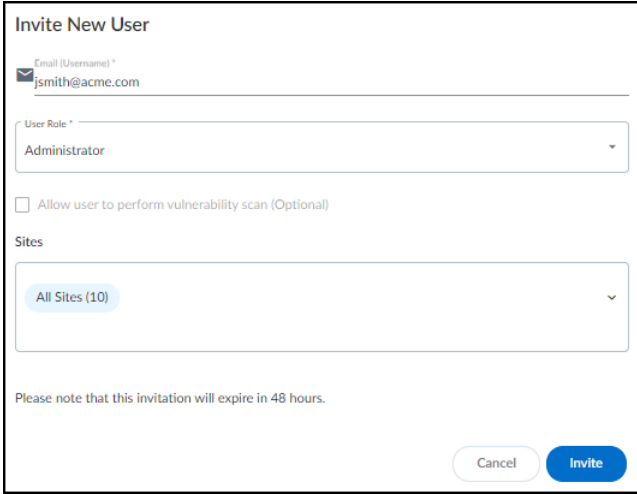
User roles are set for user accounts in external SSO authentication systems—the Palo Alto Network SSO and customer-managed SSOs—but you can also log in to the IoT Security portal with owner privileges and set other roles for administrators and read-only users. If the externally and internally managed roles are different, IoT Security assigns the higher of the two. Therefore, only set user roles internally on IoT Security that are higher than those set externally; otherwise, an internal role will never be assigned. The ranking of roles from highest to lowest is owner, administrator, read-only user.

If user accounts in an external SSO don't have any externally managed roles defined, these users won't be able to log in to IoT Security until a local user with owner privileges sets internally managed roles for them and invites them to log in to IoT Security.

STEP 1 | Invite users who have an account on an external SSO but no externally managed role to access IoT Security.

> *Skip this step if users have an externally managed role that maps to a role in IoT Security.*

1. Log in to IoT Security as a user with owner privileges, select **Administration** > **User Accounts** and then click the **Invite New User** icon (**+**) above the User Accounts table.

2. Enter an email address, choose a role (**Owner**, **Administrator**, or **Read only**), specify which sites the user can access, and then **Invite**.

IoT Security automatically generates an email with a login link and sends it to the user.

> *The invitation is valid for 48 hours after it's sent.*

When the email recipient clicks a link in the email, he or she is directed to the login page. The user clicks the **Log in with <sso-name>** button to log in through SSO. After the user logs in, IoT Security grants him or her access with the local role you specified.

3. If you want to invite more users, repeat the previous steps for each one.

STEP 2 | View users, their externally managed roles, role providers, and internally managed roles and which sites they can access.

You can see a list of users and their roles on the Access Management page in the hub and, if you're logged in with owner privileges, on the User Accounts page (**Administration** > **User Accounts**) in the IoT Security portal.

Externally Managed Role and **Role Provider**: If IoT Security applies the user role that's set on the external SSO authentication system, the role appears in the Externally Managed Role column and the SSO name appears in the Role Provider column. If IoT Security has an internally managed role for a user that's the same as or higher than his or her externally managed role, it applies the internally managed role. In this case, these two columns are empty.
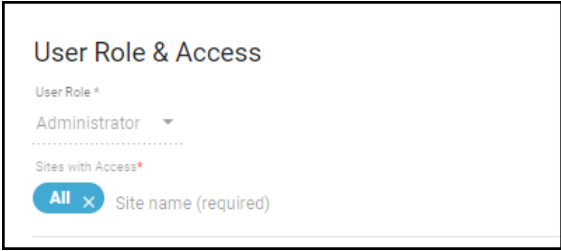
Internally Managed Role: This column lists user roles defined in IoT Security. It's only empty if there isn't a role defined internally.

> *After you create a user account in the Customer Support Portal and hub, the account won't appear on the **Administration** > **User Accounts** page in the IoT Security portal until the user logs in to the IoT Security portal.*

STEP 3 | Assign a user with an internally managed role.

1. When logged in to the IoT Security portal with owner privileges, click **Administration** > **User Accounts** and then click an entry for an administrator or read-only user in the Email (Username) column.

   The User Role & Access dialog box opens.

2. Choose a different role from the User Role drop-down list. When there are different externally and internally managed roles for the same user, IoT Security applies the role with higher privileges. Therefore, when setting an internal role, choose one that is higher than the one assigned by an external SSO authentication system.

STEP 4 | Determine which sites an administrator or read-only user can access.

By default, all users have access to all sites. To give the user access to a subset of sites, click the **x** in the All label and then select the names of the sites or site groups to which you want to permit access.

📋 *For information about site groups and how to use them to control what data users can access, see* Sites and Site Groups.

STEP 5 | When done, **Save** the configuration change.

The next time the user logs in, he or she will only have the privileges of the internally managed role and access to devices and data for the selected sites.

User Roles for IoT Security

Role-based access control (RBAC) enables you to assign privileges and access rights to administrative users through role assignment. You create user accounts in the Customer Support Portal (CSP), assign them roles in the hub, and limit the data they can access by site in the IoT Security portal. For step-by-step instructions about creating users for IoT Security, see Create IoT Security Users.

IoT Security supports the following user roles:

- App Administrator
- Instance Administrator
- Owner
- Administrator
- Read only

The App Administrator and Instance Administrator are common roles that are available to every Palo Alto Networks product application. For IoT Security, they provide the same privileges as Owner. To learn more about them, see Available Roles.

The three user roles specifically for the IoT Security portal are Owner, Administrator, and Read only.

User Role	Role Definition	Access Control
Owner (Also App Administrator and Instance Administrator)	Access to all functions in the IoT Security portal	All read/write privileges as administrators plus: • Set a global idle timeout • Change the device-to-site assignment method from one based on firewall locations to one based on IP addresses • View audit logs for all users • Set scanning permissions per administrator account • Control which sites users with administrator and read-only privileges can access • Control who receives notifications of security alerts and system alerts

User Role	Role Definition	Access Control
Administrator	Access to most functions in the IoT Security portal	Create, edit, and delete IoT Security configurations and manage their own account preferences: • See their own user role and list of sites they can access • Create, download, and delete API access keys • Update contact info • Modify their login preference if accessing multiple deployments • Shorten the idle timeout • Enable and disable alert sounds • Enable and disable alert notifications via SMS and email • Manage their own user account preferences • See the audit log for their own activities
Read only	Can only view data in the IoT Security portal	• View IoT Security data for the sites they can access • Manage their own user account preferences • See the audit log for their own activities

For Panorama-managed Prisma Access tenants with an IoT Security add-on license, add the following types of users to give them access privileges to both Prisma Access and IoT Security:

Prisma SASE Platform User Roles	IoT Security User Roles
Superuser, MSP Superuser	Owner
N.A.	Administrator*
View Only Administrator	Read-only

* There is no user role in Prisma SASE that maps to the Administrator role in IoT Security.

For new Panorama-managed Prisma Access customers as of August 2022, or an existing Panorama-managed Prisma Access customer whose Prisma Access instance has been transitioned to the Prisma SASE platform, use Common Services: Identity & Access for managing user access, roles, and service accounts.

For existing Panorama-managed Prisma Access customers whose Prisma Access instance has not yet been transitioned to the Prisma SASE Platform, you can continue using the existing process to create administrative users until the transition completes.

454