DARKTRACE

DARKTRACE THREAT VISUALIZER 6

USER GUIDE

CONTENTS

THREAT VISUALIZER BASICS 5

THE THREAT VISUALIZER

USING THE WORKSPACE

THE THREAT TRAY

CYBER AI ANALYST 30

AI ANALYST INCIDENTS

AI ANALYST INVESTIGATIONS

INVESTIGATING A DEVICE 40

FOCUS ON A DEVICE

DEVICE EVENT LOGS

OTHER INVESTIGATION METHODS

DARKTRACE MODEL BREACHES 57

MODEL BREACH ALERTS

MODEL BREACH LOGS

OTHER INVESTIGATION TOOLS 69

EXTERNAL SITES SUMMARY

ALTERNATIVE APPROACHES

APPLYING TAGS

EXPLORE

REPORTING

THE SAAS CONSOLE 88

GETTING STARTED

THE DASHBOARD

CYBER AI ANALYST

PROFILES

OTHER FEATURES

ADVANCED SEARCH AND INTEL 108

ADVANCED SEARCH

THREAT INTEL

CONTENTS

MODEL EDITOR 123

USING THE MODEL EDITOR

UNDERSTANDING A MODEL

MAKING A NEW MODEL

ADMIN INTERFACES 146

DEVICE & SUBNET ADMIN

PERMISSIONS ADMIN

SYSTEM STATUS

DARKTRACE RESPOND 156

DARKTRACE RESPOND ACTIONS

DARKTRACE RESPOND MODELS

DARKTRACE RESPOND STATE

DARKTRACE RESPOND/APPS

DARKTRACE RESPOND/NETWORK

THE MOBILE APP 192

GETTING STARTED

CYBER AI ANALYST

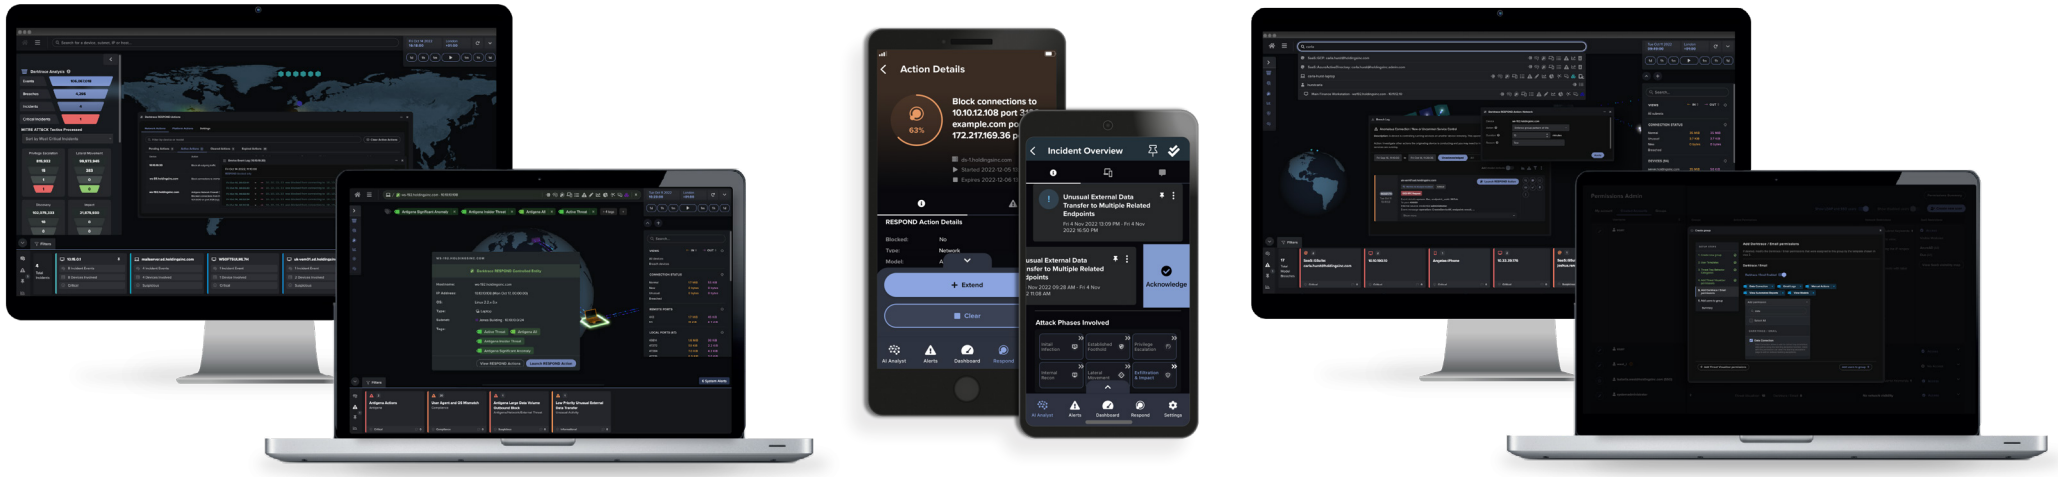
DEVICES AND MODELS

DARKTRACE RESPOND

DARKTRACE/EMAIL

OTHER VIEWS AND SETTINGS

INTRODUCTION



Darktrace Threat Visualizer

Next generation threat detection is about more than simply finding what you can conceive of in advance – it is about implicitly understanding what you, as a security professional, need to know about.

Darktrace's threat detection capability uses a self-learning approach. Instead of trying to pre-define what 'bad' looks like, Darktrace builds an evolving understanding of an organization's 'pattern of life' (or 'self'), spotting very subtle changes in behaviors, as they occur to enable rapid investigation and response to in-progress attacks.

Darktrace's Threat Visualizer is a powerful tool for intuitive visual storytelling alongside rich data that can be used to identify and investigate potential emerging threats as they develop.

This document is intended to help Darktrace users get the best possible results from the Threat Visualizer.

CHAPTER 1 - THREAT VISUALIZER BASICS

THE THREAT VISUALIZER

SUGGESTED WORKFLOWS FOR INVESTIGATING AN ALERT 6
LOGGING INTO THE THREAT VISUALIZER FOR THE FIRST TIME 8
INTRODUCING THE THREAT VISUALIZER WORKSPACE 9
MAIN MENU GUIDE 10

USING THE WORKSPACE

UNDERSTANDING THE SUMMARY PANEL 15
ADVANCED SUMMARY PANEL INFO 17
VIEWING THE NETWORK 18
EXPANDING A SUBNET 19
OTHER SUBNET VIEW FEATURES 21
USING THE OMNISEARCH BAR 22
CHANGING THE TIME 23

THE THREAT TRAY

THE THREAT TRAY 24
FILTERING ALERTS WITH BEHAVIOR CATEGORIES 26
FILTERING ALERTS WITH BASIC FILTERS 27
FILTERING ALERTS WITH ADVANCED FILTERS 28

SUGGESTED WORKFLOWS FOR INVESTIGATING AN ALERT

Where to Start

Exploring behavior can be useful for improving the understanding of what is truly happening in the digital business and how it is all interconnected.

There are five primary ways in which analyst teams can begin seeing and reacting to identified alerts or incidents produced by Darktrace

1. The Cyber AI Analyst automatically analyzes, investigates and triages all model breaches on your network. The incidents it creates give a concise summary and actionable steps to investigate any detected threats further.

2. A “threat tray” is shown at the bottom of the 3D Threat Visualizer interface in most screens and will be displayed on login. The 3D Threat Visualizer enables deep investigation of behaviors.

3. A Dynamic Threat Dashboard triage screen (accessible from the menu at top left of the home screen). This screen is intended for extremely rapid triage with a minimum of interaction. Note that it will scroll through incidents if left unattended which can be useful on SOC TV display screens.

4. A simplified SOC dashboard triage screen is available via the Darktrace mobile app.

5. Automated alerts may be exported into SIEMs or via API to other SOC systems and will include a link back to the incident data in the Threat Visualizer.

Organizations with a Darktrace/Apps, Darktrace/Cloud or Darktrace/Zero Trust module can utilize the SaaS Console for triage and analysis. This specialized interface is focused on the investigation of incidents within SaaS and Cloud environments. For more information, please see Getting Started with the SaaS Console.

Suggested Workflow from a Cyber AI Analyst Incident

Cyber AI Analyst will review and investigate all Model Breaches that occur on the system as a starting point for its analysis process. It will then form incidents - a collection of one or more related events - centered around a device. Incidents involving multiple devices will be classified as ‘cross-network’.

The Cyber AI Analyst, with its global network awareness and machine-speed investigation time, performs the heavy-lifting of the analysis process.

1. Log into the Threat Visualizer and click the Cyber AI Analyst icon in the Threat Tray or open the AI Analyst tab on the Darktrace mobile app. Review the incidents it has created.

2. Select the most severe incident or the most interesting to you based on your knowledge of the business and network setup. Review the summary created by Cyber AI Analyst to quickly understand what each incident may involve.

3. Review the summary of each event within the incident and understand how the events relate chronologically using the activity-over-time visualization. On the mobile app, read the incident overview and swipe between the events.

4. Scan the detailed event information to gauge the involved connections and review the related anomalies. Confirm if AI Analyst is currently processing any further breaches for the device. Optionally check the attack stages that AI Analyst has derived for each event.

5. If the activity of interest relates directly to a model breach, investigate the breach log.

6. Check the Actions section to see if Darktrace RESPOND/Network blocked the associated activity. Follow up the suggestions made by AI Analyst to resolve the incident.

7. Optionally create a PDF report describing the events.

8. Acknowledge each event as the investigation is concluded or acknowledge the entire incident if resolved.

Suggested Workflow from a Model Breach

When investigating an alert, a typical workflow will involve starting with summary information. This is shown by default in the Dynamic Threat Dashboard or can be seen by clicking the eye icon. The analyst is not swamped with too much to deal with all at once – enabling you to triage quickly whether the anomaly is worthy of further review.

You can then visually playback the behavior and event information, drilling down into increasing levels of detail, and broadening the investigation to correlate for related behaviors across the network at each stage of the investigation. As an example:

1. Log into the Threat Visualizer and either access the Dynamic Threat Dashboard, or filter the threat tray to include a manageable number of Model Breaches using the severity slider and adjusting the time period for which to display model breaches i.e. last 24 hours.

2. Investigate the alerts that seem most interesting to you based on your knowledge of the business and network setup. Regardless of how anomalous the activity is, if you are more interested in malware infections than data exfiltration, then looking at a 'beaconing' Model Breach first might be more relevant than looking at an 'unusual data transfer' Model Breach.

3. Identify what type of device this is and how it normally behaves.

4. Look at which events contributed to the breach.

5. Look at whether the device has related anomalies at the time or previously.

6. See exactly how anomalous this activity was or plot against related activities to spot any other anomalies or to display how important Darktrace considers that particular anomaly (overlay additional metrics including importance metrics in the graph).

7. Consider replaying the events to better understand the context of what was happening. If in the dynamic threat dashboard, click on the displayed 3D Threat Visualizer widget. Once in the 3D View, open up the device event log at the top left and press the play button on the timeline at top right.

8. Check whether similar devices are behaving in a similar way (via similar devices in device summary).

9. If appropriate, create and review raw packets (pcap) for interesting connections.

10. Consider using third-party resources for context regarding a suspicious domain, IP address or file.

LOGGING INTO THE THREAT VISUALIZER FOR THE FIRST TIME

The Threat Visualizer interface, powered by Darktrace's self-learning technology, continuously assesses the behavior of devices and users across the wider digital estate. Activity from remote endpoint devices, servers located in local datacenters, serverless architecture in public cloud environments or users working in third-party SaaS services is surfaced side-by-side.

Anomalous interactions and activity are highlighted to the end-user and prioritized for investigation by Darktrace's Cyber AI Analyst. Many operators start with these high-level investigations and then analyze detailed metadata and device or user activity in specialized interfaces.

The Threat Visualizer interface is flexible and accommodates a number of workflows and threat-investigation processes.

The Threat Visualizer interface can be accessed from a browser by navigating to the IP (*physical instance only*) or hostname (e.g., `https://[region]-XXXX-01.cloud.darktrace.com`) of the instance. The login screen will be displayed. The minimum supported browsers to access the Darktrace Threat Visualizer application are Chrome 85, Firefox 81, Safari 14.

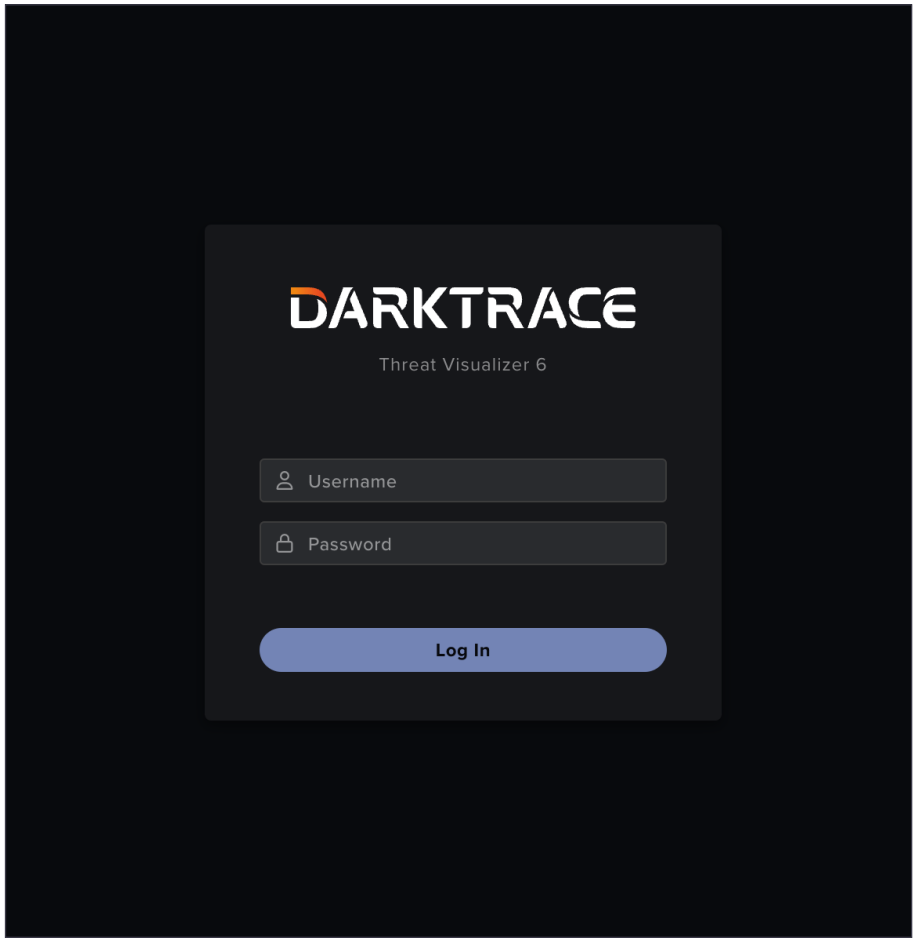
Enter your **username** and **password** to log in. The password can be reset at any time from the **Permissions Admin** page of the interface if required.

If you are accessing the Threat Visualizer via your organization's SSO system, click the **Log in via SSO** button, and log in to your SSO system as standard. You will then be redirected to the Threat Visualizer after a successful login.

For cloud based instances of the Threat Visualizer, or environments where 2FA has been enabled by an administrator, a QR code will be displayed on first access. Please scan this QR code with your preferred multi-factor authentication app such as Google Authenticator or Duo Security.

After you have logged in, the Threat Visualizer home screen will be displayed. The default landing page can be altered from Account Settings, or your administrator may have configured it for you already. The available options are: Threat Visualizer (default), Dynamic Threat Dashboard, Email Console, SaaS Console, and System Status Page.

Please note, five consecutive failed login attempts will result in the account being locked for five minutes. Subsequent failed attempts will increase the lockout time.



INTRODUCING THE THREAT VISUALIZER WORKSPACE



After logging in, the Threat Visualizer home screen will be displayed. The main workspace is used to visualize devices and connectivity, to open investigation windows and event logs, and to view alerts from Darktrace anomaly detection.

In the top left are two icons:

ICON	DESCRIPTION
Home	Clear any focused devices or windows and return to the main workspace
Main Menu	Access a list of key investigation, configuration and administration interfaces. A full, detailed list of menu items is available in Main Menu Guide.

Other key elements of the workspace are always displayed:

- The omnisearch bar can be used to search across devices, users, subnets, and external endpoints observed by Darktrace. The omnisearch bar is covered in more detail in Using the Omnisearch Bar.
- In the top right is the time selector. The chosen time is used as a starting point for replaying visualizations of connections and event logs will display events backwards from this time. Changing the time is covered in more detail in Changing the Time.
- On the left of the workspace is the summary. The summary provides key statistics about events, bandwidth, devices and AI Analyst investigations. Each element of the summary is described in Understanding the Summary Panel and Advanced Summary Panel Info.
- The main workspace displays an overview of your network. The default shows a visualization of subnets in your network on a world map. This is covered in more detail in Viewing the Network.
- Finally, at the bottom of the workspace, is the threat tray which displays alerts from Darktrace analysis. The threat tray is explained over a series of articles - starting with The Threat Tray is recommended.

MAIN MENU GUIDE

The Darktrace Main Menu offers an instant way to get to the main features within the UI. Clicking on the menu icon in the top left will display all available options. These will change depending on the user and the permissions granted to them.

Advanced Search

Opens a new browser tab to Darktrace advanced search. Refer to Darktrace Advanced Search for more information.

Tags

The Threat Visualizer supports a flexible label system called Tags to allow analysts to be able to rapidly label and refer to groups of devices within the platform. One use case for this feature is to enable monitoring of high-risk users or devices such as departing employees or key assets. Refer to Adding and Reviewing Tags for more information.

Packet Capture

Click to view the list of PCAP files you have created which are stored on the Darktrace appliance. Refer to Creating Packet Captures for more information.

Darktrace RESPOND Actions

Presents currently active and expired actions taken by Darktrace RESPOND components (*excludes Darktrace/Email*) and allows the actions schedule to be configured. Current and historic actions can be exported as a CSV. Refer to Reviewing Darktrace RESPOND Actions for more information.

Requires an active Darktrace RESPOND license.

Darktrace RESPOND/Network Summary

The Darktrace RESPOND/Network Summary page displays overviews of your Darktrace RESPOND/Network eligibility coverage, current state (autonomous/human confirmation) and a breakdown of RESPOND actions that have been triggered in your environment. These summary interfaces are key to understanding interaction with Darktrace RESPOND by your users, the types of action recommended - or taken - by Darktrace RESPOND in your environment, and check in on progress towards a fully autonomous environment.

Requires an active Darktrace RESPOND/Network license.

Email Console

Pivots to the Darktrace/Email interface for investigation and autonomous actions on your email domains. Please see the documentation on Darktrace/Email for more information.

SaaS Console

The SaaS Console is a specialized user interface for investigating anomalous activity and user behavior in and across SaaS and Cloud environments, surfacing the events and analysis from Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules in one centralized location. Please see The SaaS Console for more information.

AI Analyst Investigations

Launch Investigation

Manually trigger a new AI Analyst investigation on a device or SaaS user. See Triggered AI Analyst Investigations for more information.

View Investigations

Opens a list of current and previous AI Analyst triggered investigations.

Models

Model Editor

A Model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event. Refer to The Model Editor for more information.

Model Summary

List of the models with analysis of how many times and how strongly each model has breached since the appliance was installed.

Model Updates

This page allows you to view recent updates to the models. Refer to Model Updates or the Darktrace System Administration Guide for more information.

Reporting

Cyber AI Insights Report

Cyber AI Insights reports provide an overview of your organization's current Darktrace coverage, key statistics on threat trends across deployed components, data on response times by Darktrace RESPOND and analyst engagement. Refer to Cyber AI Insights Reports for more information.

Executive Threat reports are high level, visual overviews of network activity and model breach events. Refer to Executive Threat Reports for more information.

Executive Threat Report

Executive Threat reports are high level, visual overviews of network activity and model breach events. Refer to Executive Threat Reports for more information.

Download Reports

This page allows Threat Intelligence reports and both manually generated and scheduled Executive Threat Reports to be retrieved.

Admin

Device Admin

This page contains a list of basic information regarding every internal device Darktrace has ever observed. The list is exportable to Excel and some fields can be altered (e.g., type of device). Refer to Device Admin for more information.

Subnet Admin

Similar functionality but applies to subnets. The DHCP slider shows whether Darktrace should be seeing DHCP for that specific subnet. Refer to Subnet Admin for more information.

Audit Log

This page shows recent interactions with the platform completed by the various users (e.g., acknowledging breaches, logging in and out).

System Status

System metrics such as ingested traffic quality, master/probe reachability, and protocols observed can be reviewed.

System Config

Provides all configuration settings for the Darktrace Threat Visualizer application including Darktrace RESPOND settings and authentication for DETECT modules and integrations. Alerting options can be configured here.

Permissions Admin

The Permissions Admin interface allows user permissions and data visibility to be controlled at an individual and group level. Group permissions for LDAP and SSO users can also be configured. For detailed information on users, groups, permissions and data visibility, please refer to Permissions Admin.

This interface has replaced the legacy User Admin and Group Admin interfaces.

Permissions

This interface has been superseded by Permissions Admin.

Utilities

Ü Punycode Convertor

Enter text in the top window to convert to Punycode; enter Punycode in the bottom window to convert to text. Punycode is used in DNS to encode Unicode international domain names (IDN) into ASCII. Can be identified as it always starts 'xn—'".

(.*) RegEx Tester

Enter a RegEx in the top bar and an example string to test it in bottom bar. If the string matches the RegEx the bottom box's border turns green; otherwise it remains white/yellow.

64 Base64 Convertor

Enter text to be decoded or encoded using Base64.

js JS Beautifier

Tool for 'beautifying' JavaScript to increase readability

Epoch Converter

Converts epoch time in seconds since 1st Jan 1970 (as seen in advanced search) to normal time.

Intel

Trusted Domains

Trusted domains are endpoints which Darktrace will consider as 0% rare; this feature ensures that models relying on domain rarity will not fire for activities involving common domains - a default, editable list is provided. See Trusted Domains for more details.

Watched Domains

Watched domains are endpoints which trigger automatic model breaches if observed in connectivity. The list is not populated by default but may be added to by TAXII feeds, Darktrace Inoculation, via the Threat Visualizer API or by manual edits. See Watched Domains for more details.

TAXII Config

Permits the ingestion of internal or third-party TAXII feeds and STIX files. The last 10,000 observables ingested can be reviewed, whether derived from stand-alone files, subscribed TAXII collections or Darktrace Inoculation. See TAXII Config for more details.

Dynamic Threat Dashboard

The Discovery View provides an easy left-to-right dashboard to investigate an incident down to the connections that caused the alert to fire. Refer to Dynamic Threat Dashboard for more information.

Explore

Explore Subnets

Explore Mode for subnets gives an overview of all connection events between subnets at a given time interval, filterable by connection size and transfer ratio. Drill-down to a device-to-device level is also possible. A master layout can be defined to reflect existing network topology diagrams.

Explore Tags

Explore Mode for tags gives an overview of all connection events between Tags at a given time interval, filterable by connection size and transfer ratio. Drill-down to a device-to-device level is also possible. A master layout can be defined to reflect existing network topology diagrams.

Help

Ask the Expert

Ask the Expert allows for the creation of incidents which can be submitted to Darktrace analysts for feedback. Refer to Ask the Expert for more information.

View Questions

After a request has been generated, you will be able to view your open and closed questions by clicking View Questions; this window will show the conversation and any relevant information, as well as the ability to reopen or delete an historic request.

API Help

Provides a link to the Threat Visualizer API documentation hosted on the Darktrace Customer Portal.

Account Settings

Change settings for your own account including default color-coding in the event log, enabling accessibility mode, selecting the AI Analyst language, changing the default landing page and threat tray behavior category filters.

Customer Portal

Opens the Darktrace Customer Portal in a new browser window or tab. The Customer Portal provides access to all product documentation, support tickets, analyst insights and Darktrace Education materials.

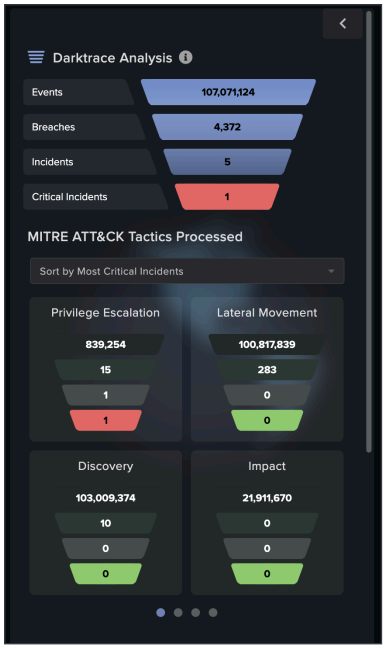
Logout

Log out the current user account.

UNDERSTANDING THE SUMMARY PANEL

Key statistics about your Darktrace environment can be accessed from the summary panel on the left of the Threat Visualizer workspace. By default it is collapsed. Each icon can be hovered for basic statistics or the whole panel can be expanded for detailed information.

Darktrace Analysis



Every second, Darktrace observes thousands of raw datapoints - whether passive network traffic processed by Darktrace Deep Packet Inspection, metadata from Darktrace/Endpoint agents and virtual sensors, streams and signals from telemetry modules and contextual threat intelligence integrations, or user activity data observed Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules.

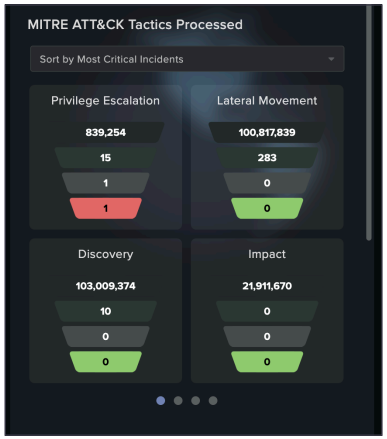
Each pattern is analyzed; is this unusual? Does this fit the pattern of life for the device or its peers? Is this behavior of interest? "Events" represents a subset of all these patterns - the total number of raw events that were passed to Darktrace modeling for further investigation over the last 28 days.

The Darktrace model engine is an evaluation layer; the models framework leverages the underlying 'pattern of life' detection, combining intelligent outputs from Darktrace Self Learning AI with the expertise of the specialized Darktrace analyst team. From the raw events it receives, a small fraction reach the criteria to trigger a "model breach". The total number of these model breaches generated over the last 28 days is displayed.

The Darktrace AI Analyst investigates, analyzes and triages threats seen within your Darktrace environment. By learning from the millions of interactions between Darktrace's expert analysts and Darktrace DETECT components, the AI Analyst combines human expertise with the consistency, speed, and scalability of AI. Every time the conditions for a model are met and a model breach is created, the Darktrace AI Analyst investigates the activity and concludes whether it needs to be surfaced for human analysts to review. "Incidents" represents the total created from this AI-powered investigations into anomalies across your digital environment in the last 28 days.

Finally, behavior categories are high level filters that allow an operator to focus in on specific levels of severity or behavior. There are four categories: Critical, Suspicious, Compliance and Informational. Categories are used across both AI Analyst and Darktrace models. "Critical Incidents" therefore, are incidents created by AI Analyst analysis that have been allocated the highest priority behavior category.

MITRE ATT&CK Tactics Processed



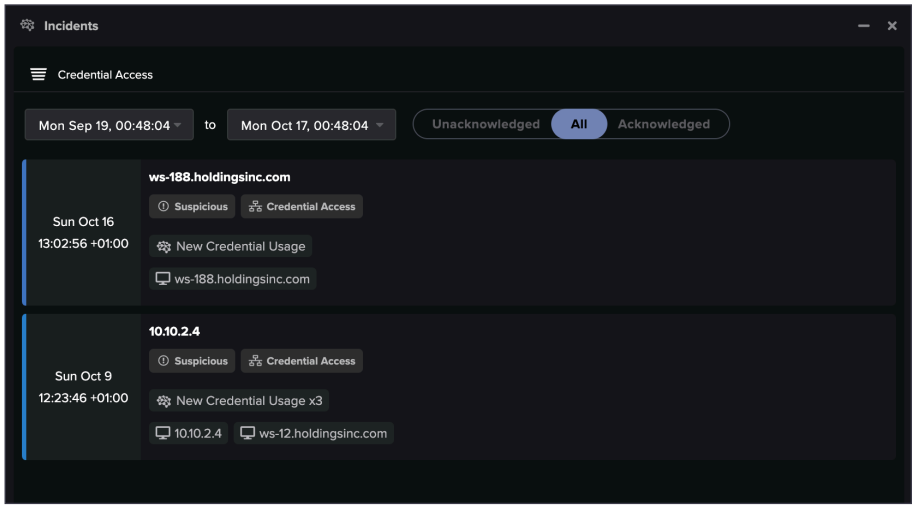
The MITRE ATT&CK Framework is used by SOC analysts, threat intel experts and enterprise security teams to classify the "what", "why" and "how" of cyber threats; the framework is integral to many organizational playbooks on vulnerability mitigation and incident response. MITRE ***tactics*** are used to model the "why" - what a potential malicious actor intends to achieve through their behavior. The Darktrace DETECT platform supports the integration of this framework through the categorization of Darktrace AI Analyst events and Darktrace models.

MITRE ATT&CK Tactics Processed visualizes how raw activity connected to each tactic has passed through Darktrace AI-powered analysis over the last 28 days, distilling thousands of events into just key behavioral alerts.

For each tactic, the number of events that were relevant to the type of activity the tactic encompasses - and therefore were evaluated by the Darktrace model engine - are displayed. Tactics are not mutually exclusive; the same activity may be relevant to, and been evaluated for, multiple tactics.

From these evaluated events, some reach the criteria necessary to trigger a ***model breach***. The Threat Visualizer platform provides a curated set of default models as standard, designed and constantly tuned in parallel with the evolving threat landscape; default models are mapped to relevant MITRE ATT&CK Framework techniques and tactics. For each tactic, the number of model breaches triggered for mapped models is shown. Click to open a breach log filtered on model breaches for the chosen tactic.

For Darktrace Threat Visualizer 6, all AI Analyst incident event types have also been mapped to one or more MITRE ATT&CK Framework tactics. Each incident is comprised of events which detail a specific anomalous activity - or chain of activities - investigated by AI Analyst. Where an incident comprises one or more events with activity relevant to the tactic, it is counted here. Click to open an AI Analyst incident log filtered on these relevant incidents for the chosen tactic.



As above, these AI-powered incidents can be further distilled down to just those of the highest priority - "Critical Incidents". Click to open an AI Analyst incident log filtered on AI Analyst incidents with events relevant to the chosen tactic.

Refer to Opening the Model Breach Log and Cyber AI Analyst Incident Logs respectively for more information on how to understand the Model Breach logs and AI Analyst incident logs launched from this panel.

Incidents vs Model Breaches

Sometimes there may no relevant model breaches, but AI Analyst incident has still identified activity of interest that is linked to the tactic. Although a model breach may be the trigger for an investigation, that does not mean the activity AI Analyst surfaces is directly related to the original model breach. Like a human, the Cyber AI Analyst uses a model breach as a starting point to investigate the device. The behavioral analysis it performs may discover anomalies or patterns of activity that were not the original trigger point for the model breach but are worthy of investigation.

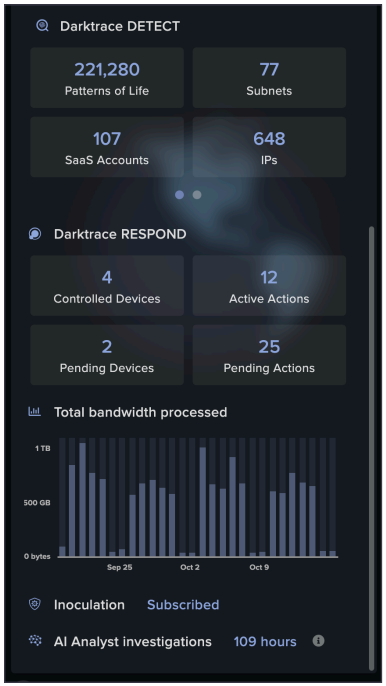
Therefore, AI Analyst may sometimes uncover and report events relevant to a tactic from an investigation triggered from an unrelated model breach, or a model associated with a different tactic.

ADVANCED SUMMARY PANEL INFO

Summary Pane

Additional sections display information about the number of patterns of life and modeled entities in your environment.

Darktrace DETECT



Darktrace creates an overall pattern of life for every entity it sees - whether endpoint devices, users interacting with resources in external platforms, or servers within the internal network. Networks are alive with constant activity and repeated patterns; Darktrace's self learning AI observes these patterns and derives a sense of "normal" behavior for devices, users and peer groups.

Each pattern of life for a device or a peer group is itself built from of thousands of fragmentary patterns of life - a unique connection, event or activity that contributes to the overall understanding of your digital environment. The number of patterns of life is not expected to match the number of modelled entities such as devices or users.

Device counts display the number of entities - users detected in external platforms (SaaS users), devices, credentials - that Darktrace has observed and is actively modelling a pattern of life for. These counts will change as new devices appear and inactive devices are no longer seen. Statistics are calculated over 7 days, 28 days or 12 weeks.

Darktrace RESPOND

If Darktrace RESPOND/Network is enabled, statistics on the number of actions taken and devices actioned will be displayed.

Total Bandwidth Processed

The total bandwidth processed over the last 28 days is displayed for both network and endpoint data; each bar can be hovered for a breakdown between the two input sources. Processed bandwidth is not equivalent to ingested bandwidth - some ingested connections may not be processed due to configuration settings such as exclusion rules.

AI Analyst Investigations

AI Analyst investigates anomalies detected by the Darktrace system at machine speed and surfaces only those that need human attention. The summary displays the equivalent number of hours it would take a human analyst to perform the same investigations.

VIEWING THE NETWORK

The main workspace is used to visualizer connections and devices in 3D in real time. The default view is a simplified subnet view.

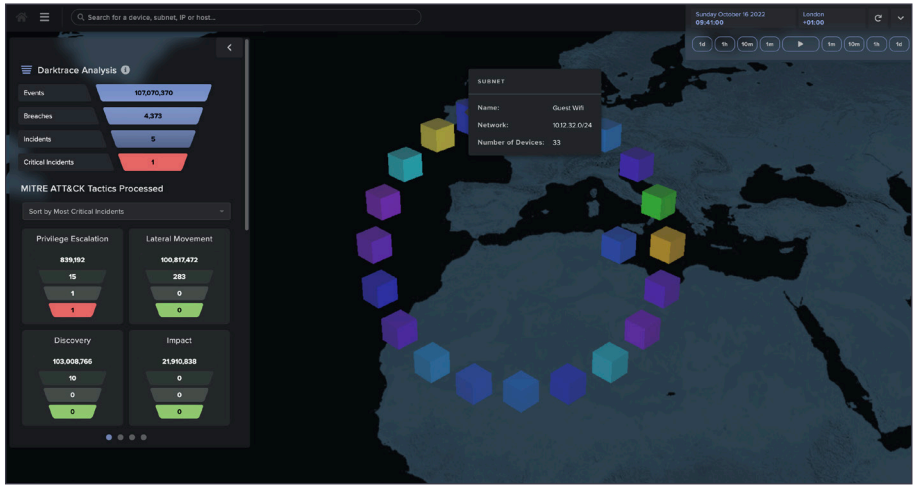


Workspace

The Threat Visualizer homepage displays an overview of your network. Each subnet is displayed as a "cube" icon. Hovering over a subnet icon will show the label or equivalent subnet range, if no label has been provided. Where many subnets are located in close geographic proximity, these are collapsed into a single cube. The list of collapsed ranges is shown on hover.

For devices monitored by Darktrace DETECT & RESPOND/Endpoint agents, subnets are automatically created and grouped by the country that the devices are located in. Subnets created directly from network traffic can be relocated geographically by providing a latitude and longitude on the Subnet Admin page. Across the top of the workspace are special-purpose subnets for internal (unmodeled) traffic, internal multicast and broadcast traffic. These subnets cannot be moved.

Individual Subnets



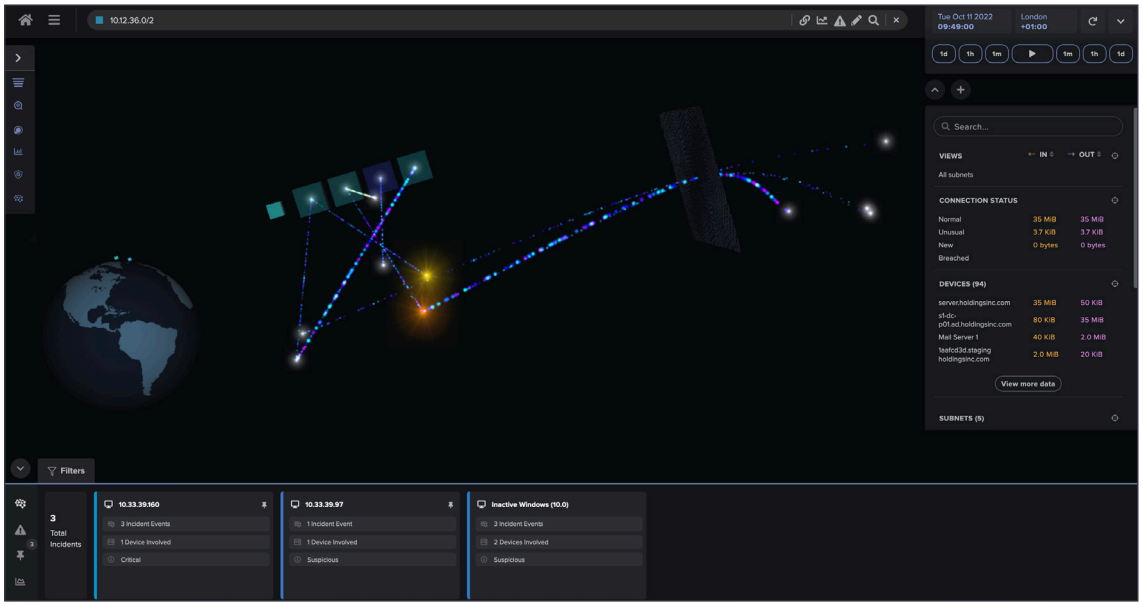
The color of the subnet cube can be used to quickly gauge the level of anomaly:

- A green icon indicates at least one device in the subnet has been subject to a Darktrace RESPOND action.
- An icon from yellow to red indicates at least one device in the subnet has been subject to a model breach, with the color representing severity.

Where many subnets are collapsed into one icon, the icon will show the highest anomaly level of the underlying subnets.

Clicking once on an icon will center the workspace; when centered, hovering over the icon will give more detailed information on device membership. Clicking again will open a detailed view of the subnet.

EXPANDING A SUBNET



Focusing on a Subnet

Click on any cube on the main workspace to focus on a specific subnet. The visualization will switch to subnet mode. To return to the full network visualization, click the small cube at the top left of all connecting cubes, or click the home button.

Visualization

The main workspace now shows a visualization of network devices in the subnet connecting to internal and to external locations in real time. SaaS devices are not visualized by subnet view.

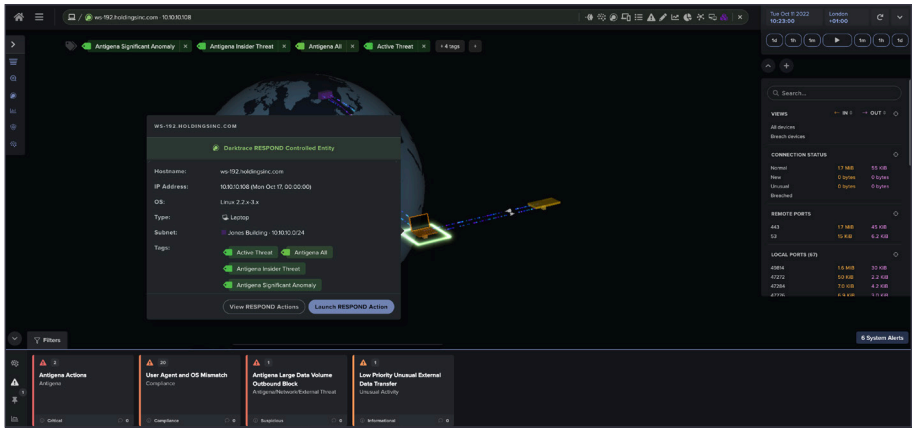
Devices

Glowing nodes are devices with active connections at the time chosen in the top right. When the visualization is unpaused, nodes will appear and disappear as connections start and finish. Nodes which are colored from yellow to red are devices displaying anomalous behavior.

If the device was subject to a Darktrace RESPOND action during the visualization, a green line will appear around it for the duration of the action. If the device currently has pending Darktrace RESPOND actions, an orange line will appear around the visualization.

Hover over a node to display a summary tooltip. The summary includes device details like the device's IP or hostname, device type, current subnet, any applied tags or any recently observed credentials. Hovering over devices will also show whether they are currently controlled or pending control by Darktrace RESPOND actions, and provide the option to trigger a manual Darktrace RESPOND action against the device. See Manual Darktrace RESPOND Actions for more information on triggering these actions.

Click the device to change the focus to the device.



Connectivity

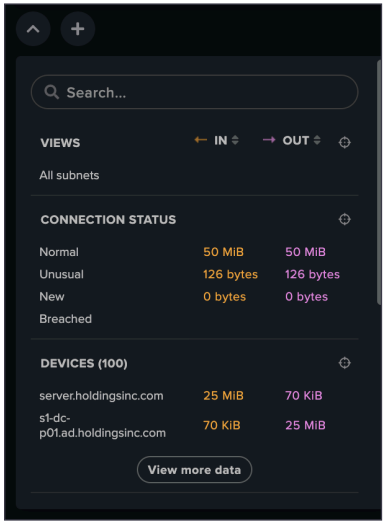
Blue lines represent active connections transferring data. Devices may connect to internal locations, to external locations or connect to one another within the subnet.

Internal subnets are represented by cubes on the left of the visualization. Hover over these cubes to see the subnet they represent or click to change the focus to the alternate subnet.

Pseudo-subnets are created for devices monitored by Darktrace DETECT & RESPOND/ Endpoint agents, grouped by the devices' geographic location. As agents monitor devices in remote networks, inter-subnet connectivity will not be displayed. Nodes for these devices will always be displayed connecting to external locations.

A boundary on the right of the visualization represents the border between internal and external locations. External connections from devices in the subnet are shown passing through this boundary.

OTHER SUBNET VIEW FEATURES



Subnet Statistics

On the right of the workspace is a panel with statistics about the subnets' connectivity over the time window chosen in the selector. The time window - default 5 minutes - is shown below the timezone. The panel includes the most-used protocols and the largest data flows to other internal subnets and the wider internet. Each section of the panel can be collapsed by clicking the heading.

The network statistics panel is useful, for example, to quickly identify devices causing the highest inbound data transfer in the subnet or to check for the presence of prohibited protocols. Filters applied to the panel also impact the main workspace, allowing visualization to be restricted to certain protocols (e.g. UDP) or to certain external endpoints.

Click a key on the left of the panel to apply it as a filter. The free-text search can also be applied to any key, allowing specific protocols or endpoints to be located for filtering.

Multiple filters can be applied and will appear above the panel. Filters applied on this panel will impact other windows and event logs.

Subnet Options

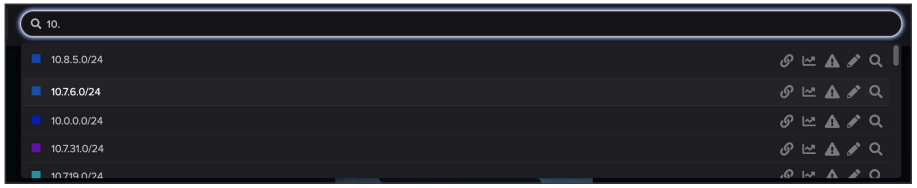
The network range that represents the subnet is populated in the omnisearch bar. Beside the subnet range are a selection of icons.



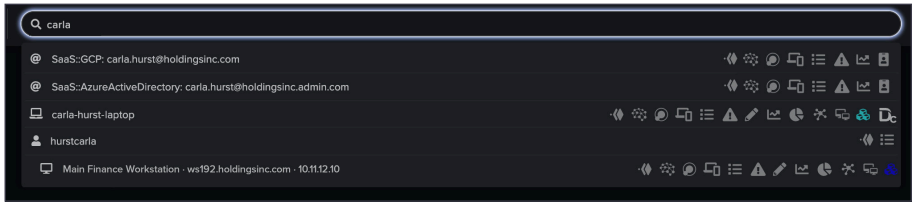
ICON	DESCRIPTION
Explore subnet connectivity	Explore is an alternative interface for investigating subnet-subnet connectivity. This icon opens the Explore interface focused on the chosen subnet. Explore subnets is covered in more detail in Explore Mode.
Open subnet graph	The graph allows metrics of device behavior to be plotted over time. In subnet mode, metrics are plotted for all devices in the subnet. The graph is described in more detail in Other Device Summary Sections in the context of individual devices but is also applicable to subnet mode.
View Breaches	This icon opens the model breach log - a list of model breaches triggered by all devices in the subnet. The default timeframe is 7 days but can be altered. Breach logs in general are covered in more detail in Opening the Model Breach Log.
Edit Subnet Info	Edit provides quick access to change the subnet label, geographic location and tracking information. This information can also be changed in bulk in the Subnet Admin interface.
Devices on this subnet	Clicking this icon will trigger an omnisearch query for all devices in the subnet, not just those currently visualized on the workspace. Scroll to browse or click to change the focus to a specific device. Detailed information about omnisearch queries is also available in Using the Omnisearch Bar.

USING THE OMNISEARCH BAR

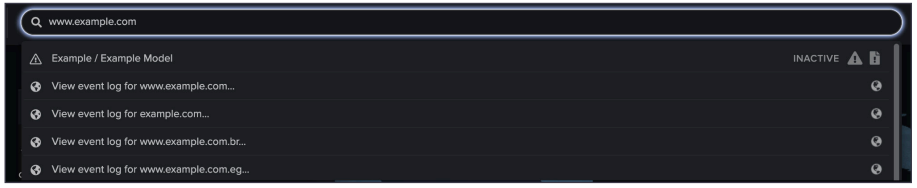
The omnisearch bar is another way to focus on subnets, devices, models and specific connections or events. The search autocompletes and suggests the most relevant search results.



For devices, it is possible to search by label, mac address, IP and hostname. When searching for devices, a list of suggested devices will be displayed along with their derived device type and simple information such as IP, hostname, MAC address (where available). If network devices are subject to a Darktrace RESPOND action at the current time, the omnisearch entry will appear shaded green.



Users - credentials Darktrace has observed being used by a device - can also be returned. The devices currently associated with the user will appear as an indented list underneath their entry.



To focus on a device, subnet or connection - click the entry. The options for each result can be clicked without changing the workspace focus. This means event logs and summaries can be opened within focusing the view completely on the device.

There are many types of information that can be returned by an omnisearch query:

EXAMPLE RESULT	TYPE	ACTIONS
ws183.holdingsinc.com · 10.15.15.2	Device	For device results, the device options described in Device Investigation Options are also accessible from the omnisearch bar.
10.15.15.0/24	Subnet	Change the focus to a specific subnet. The same subnet options described in Other Subnet View Features are also accessible from the omnisearch bar.
view event log for google.com...	External location	For external IPs and hostnames, the external sites summary can be opened from the omnisearch. The summary covers interaction between internal devices and the external host over time. The summary is covered in more detail in External Sites Summary.
Compliance / Telnet	Model	Models can also be searched for from the omnisearch. If the model is inactive (disabled), this is indicated. Click the triangle icon to open the model editor for the model, or click the page icon to open a breach log for the model.
S3a9BRg6DoFtrbpu2e	UUID	Unique uuids are assigned to every event, connection or group of connections observed by Darktrace. Click the lines icon to open an event log for the UUID. The UUID will also be populated in omnisearch if a pivot from Advanced Search was performed.

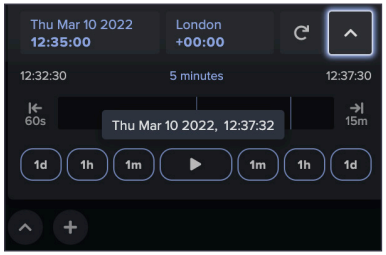
Advanced searches can also be performed with prefixes. For example, using `subnet:10.0.0.0/24` will return devices within this subnet, `tag:Test` will return devices tagged with the “Test” tag.

CHANGING THE TIME

The time selector in the top right controls the connection visualization and the starting point for event logs in the main workspace. Focusing on an alert will change the time in the selector to the time of the alert so events can be replayed.

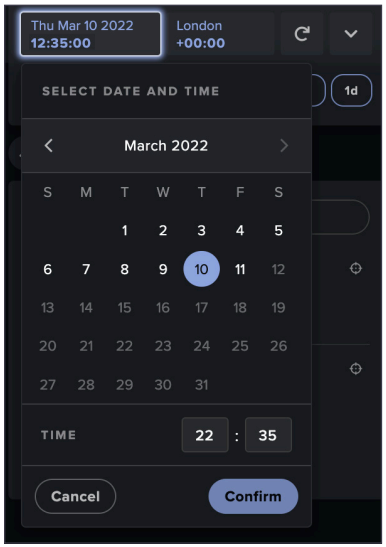
Some elements, such as the threat tray, show alerts over a large time window. These elements have their own time selector and this will be clearly indicated.

Changing the Time



- Click the time in the top right to select a specific time and date.

  To move relatively, use the 1 day, 1 hour, 1 minute buttons at the bottom. Those on the left of the ▶ play icon move the time backwards, those on the right move forwards.



- Click the timezone to change the current location-based timezone. By default, the browser timezone is used. Any open event logs will not show the new timezone until closed and reopened.
- Click the ↻ return icon to reset the time to the current time.
- Click the ⌄ dropdown arrow to change the time window. The window is discussed in further detail below.
- Click the ▶ play button to play the 3D visualization currently on the workspace real time from the time chosen.

Time Window

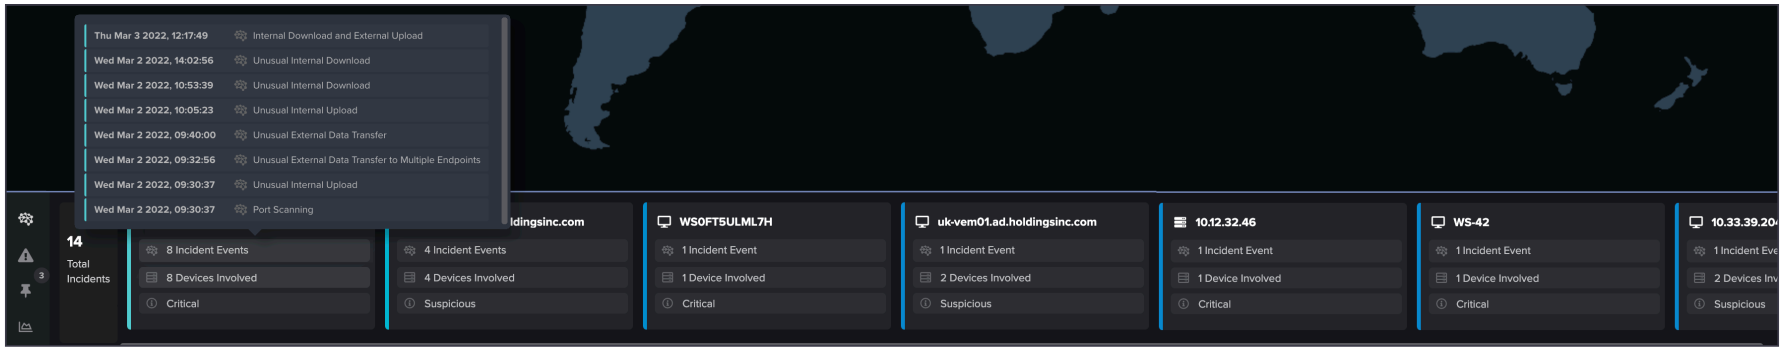
Some elements use a window of time to display information over. When focusing on a device or subnet, a panel with statistics appears on the right of the workspace. This element uses the time window to show statistics on data transfer volumes and connections - changing the window will alter this panel.

The size of the time window is shown above the bar. The start and end times are displayed on either side.

- Values on the left of the window ⇤ make the window smaller when clicked.
- Values on the right of the window ⇥ make the window larger when clicked.

Hover within the window at a specific time and click to set the main time selector to that time.

THE THREAT TRAY



Threat Tray

The panel at the bottom of the main threat visualizer workspace contains alerts from Darktrace anomaly detection. For any investigation workflow, the tray is the first place to start. By default, it will show incidents created by AI Analyst on first access. Afterwards, it will remember the last mode used.

The tray has four options: AI Analyst, Model Breaches, Pinned Alerts, Model Breach Cluster View.

AI Analyst

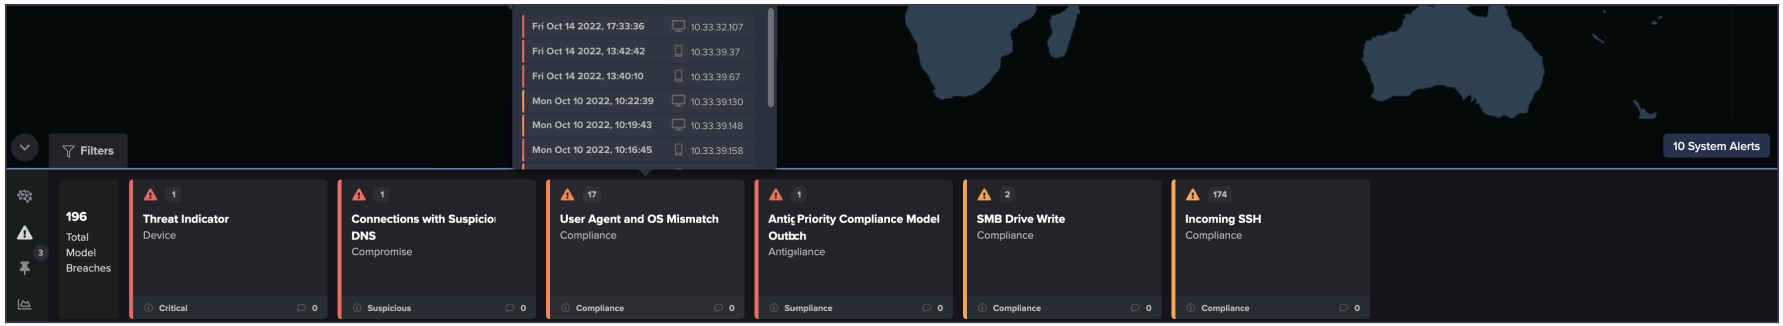
Incidents are created from AI-powered investigations into anomalies across your digital environment. Every time the conditions for a model are met and a model breach is created, AI Analyst investigates the activity and concludes whether it needs to be surfaced for human analysts to review. The results of this analysis is then grouped in incidents.

AI Analyst only reports upon the most important or interesting activity - it performs the heavy lifting and should be the starting point for any Darktrace operator.

Model Breaches

Darktrace models are a logic framework that allows operators to interact with the output from 'pattern of life' and anomaly classifiers. Models are primarily used to identify and alert on anomalous and potentially malicious behavior. The Threat Visualizer platform provides a curated set of default models as standard, designed and constantly tuned by the specialized Darktrace analyst team.

When conditions for a model are met, this creates a model breach and can trigger an alert or action. AI Analyst investigates every model breach that occurs on the system and reports only the most interesting activity as incidents.



Pinned Alerts

Model breaches can be pinned for quick access. Pinning is specific to each user and persists between logins (***requires local storage***). A list of pinned model breaches is retrievable by clicking this icon.

Clicking on a model breach from the pin list will open the model breach log. Model breach logs are covered in further detail in Opening the Model Breach Log.

Globally pinned AI Analyst incidents will also be returned here for reference.

Model Breach Cluster View

In cluster view, model breach events are plotted onto a graph of severity score against time, allowing patterns and clusters of model breaches to be quickly identified. A key of what each point represents is shown on the right. Model breaches are aggregated into single dots by the model breached, by the device, or by the credential that triggered the model breach. Aggregation is dependent on the “Sort By” chosen.

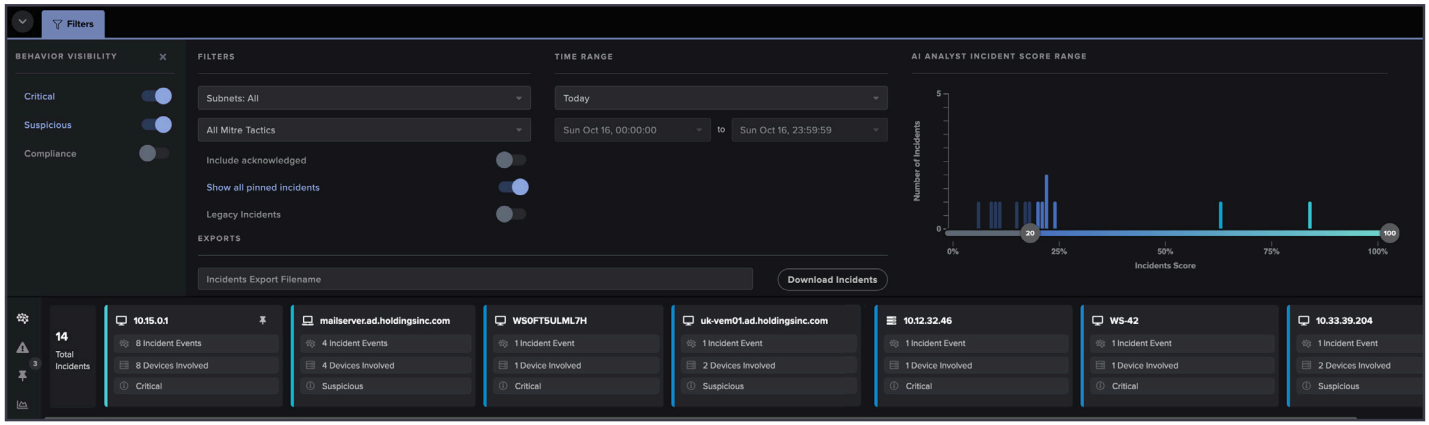
Clicking on a point on the graph will open the model breach log. Model breach logs are covered in further detail in Opening the Model Breach Log.

Filtering the Log

For AI Analyst and Model Breaches, the total number of alerts present in the threat tray after filters are applied will appear on the left.

Filtering the tray to just alerts relevant to your workflow will now be covered in Filtering Alerts with Behavior Categories, Filtering Alerts with Basic Filters and Filtering Alerts with Advanced Filters.

FILTERING ALERTS WITH BEHAVIOR CATEGORIES



The alerts in the threat tray can be filtered in many different ways to support different priorities and investigation workflows. The most important filter to be aware of is behavior categories.

Behavior categories are high level filters that allow an operator to focus in on specific levels of severity or behavior. There are four categories: Critical, Suspicious, Compliance and Informational. Categories are used across both AI Analyst and Darktrace models.

- **Critical** alerts represent events with the highest anomaly and are where an analyst should focus their attention in the first instance.
- **Suspicious** alerts are those which suggest moderate levels of anomaly but do not warrant prioritization over critical.
- **Compliance** alerts are raised for behaviors that may be counter to organizational policies and best operating practices.
- **Informational** alerts cover low level indicators and notable events which do not indicator malicious activity in isolation. The informational category is not applicable to AI Analyst alerts as only events worth analyst attention are surfaced.

Categories can be applied to AI Analyst and Model Breaches separately. To see the currently applied categories, open the Filters tab and click the ··· three dots icon.

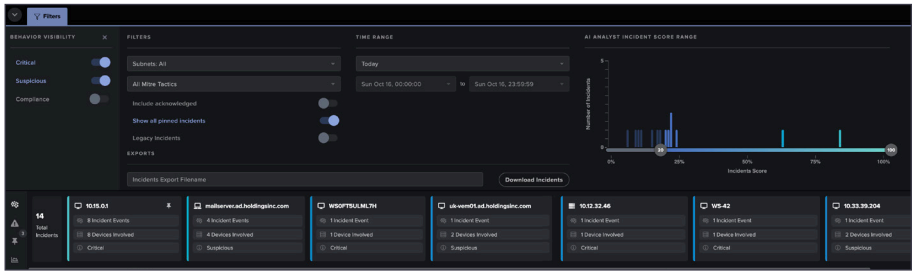
When new users are created, the behavior categories they see as default when they first log in can be customized. For existing users, the default behavior categories they see can be customized in two ways:

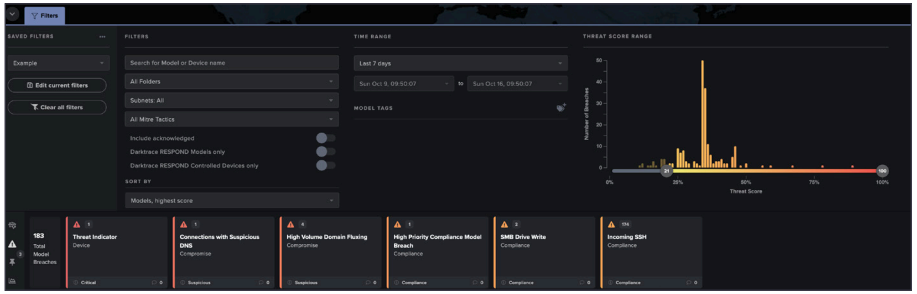
- In **Account Settings**, by altering the settings “Default Threat Tray AI Analyst Behavior Visibility” and “Default Threat Tray Model Behavior Visibility”.
- By creating a saved filter with the desired behavior categories and setting it as default. This saved filter takes priority over account settings, if set.

FILTERING ALERTS WITH BASIC FILTERS

Filters

For many operators, there are two more key filters they will use when investigating alerts: time and by score. The time range and score range over which alerts are returned can be customized.





More granular filters are also available for both AI Analyst incidents and model breaches so that specific types of activity or parts of the network can be focused on. These filters are entirely optional and are in addition to the key filters already discussed.

Filter sets can be saved for AI Analyst incidents and model breaches and set as the default view.

Time

AI Analyst incidents and model breaches can be filtered by a time range to control the quantity of alerts. The range is controlled separately for the two alert types. If a custom range is chosen, the two date fields will become interactable so a range can be chosen.

AI Analyst incidents are made up of multiple alerts that may occur over a longer time frame. An incident will be returned if any of its underlying events happened within the range, ensuring that active, developing incidents are not excluded.

This time range may also impact other investigation elements - this will be noted where relevant.

Score

Both AI Analyst incidents and model breaches can be filtered by a severity score using a slider. By default, the minimum is set at 20% for AI Analyst incidents and 60% for model breach alerts. This can be altered manually or overwritten by a saved filter.

AI Analyst incidents are given an overall score from 0-100%, represented by a gradient from dark (lowest severity) to light blue (highest severity). The score is calculated from the number of devices, the type of activity and the level of anomaly across all events that make up the incident. High scores indicate incidents that AI Analyst believes to be high priority for human attention.

When the criteria for a Darktrace model are met, a model breach is created with a score from 0-100%, represented by a gradient from yellow (lowest severity) to red (highest severity). The score is calculated from the type of activity, the priority of the devices involved and the rarity of the model breach for the device.

Score can be a useful metric to use when prioritizing within a category - for example, starting with category “critical” and addressing alerts in order of score. The two scores should not be seen as comparable. AI Analyst only surfaces the most important and interesting activity for users, whereas model breaches cover the full range of activity from low level, informational activity to high priority alerts.

FILTERING ALERTS WITH ADVANCED FILTERS

Filters

For many operators, there are only two more key filters they will use when investigating alerts: time and by score. The time range and score range over which alerts are returned can be customized.

More granular filters are also available for both AI Analyst incidents and model breaches so that specific types of activity or parts of the network can be focused on. These filters are entirely optional and are in addition to the key filters already discussed.

Changing the Default View

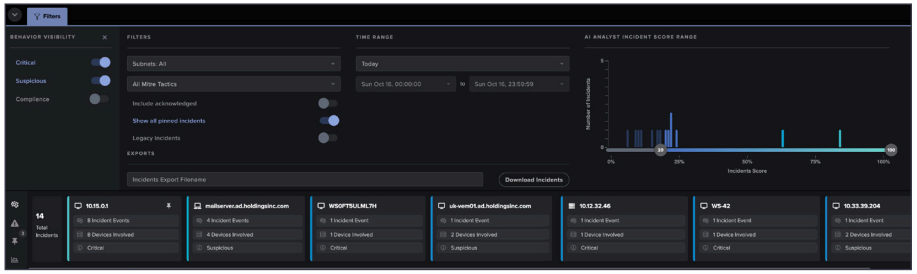
The current set of filters can be saved for AI Analyst Incidents and Model Breaches and applied as the default for subsequent sessions.

Saved filters for both alert types include the behavior category visibility, the score range, the time range and subnet filters. Model breach saved filters also include model folder filters, model tag filters, the sort type and whether results should include acknowledged or exclusively Darktrace RESPOND model breaches.

Turning on "Default Filter" when a filter is saved or edited (and saved) will mean the filter set is applied in subsequent sessions.

Saved filters take priority over account settings, if set. Custom time ranges - those specifying exact dates - are not valid for saved filtersets.

AI Analyst Threat Tray Options and Filters



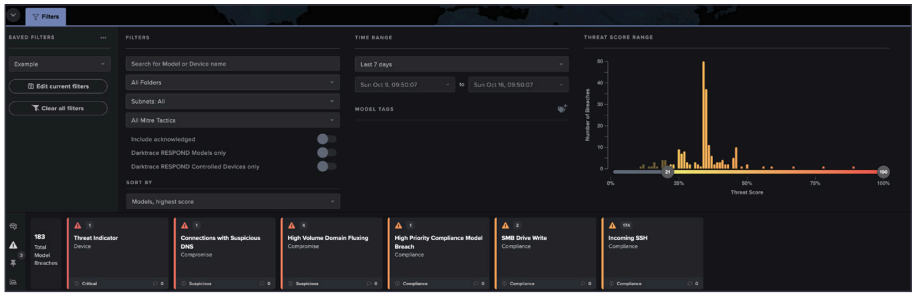
AI Analyst incidents can be further filtered by network subnet and acknowledgement status. There are no sort or grouping configuration options for AI Analyst alerts as incidents are always sorted by severity and grouped by high-level analysis.

The current alerts in the tray can be exported as a PDF. To do so, enter a filename under the "Export" header and click **Download Incidents**.

FILTER	DESCRIPTION
Subnets	This filter allows devices from specific subnets to be removed from the returned results. Subnets can be selected from a list or a custom range (such as 10.0.0.0/8) can be defined. When at least one subnet is selected, the filter can be set to "Include Subnets" (only show devices from these subnets) or "Exclude Subnets" (remove all devices from these subnets). Subnet filters are not applicable to geographic subnets created for Darktrace DETECT & RESPOND/Endpoint agent devices.
MITRE Tactics	As of Threat Visualizer 6, all AI Analyst incident event types have been mapped to one or more MITRE ATT&CK Framework tactics. This filter allows the returned incidents to be filtered down to only those that contain an event mapped to the specific tactic. Events can be mapped to multiple tactics and incidents may contain multiple events, each with distinct tactics assigned.
Include Acknowledged	Model breach and AI Analyst alerts can be acknowledged as part of an investigation workflow. By default, acknowledged alerts are removed from the returned results. Turn this setting on to include these results.
Show all pinned incidents	AI Analyst incidents can be globally pinned - this means they will be returned as part of the results, regardless of whether they are within the timeframe. Turn this setting off to only include pinned incidents that are within the chosen timeframe.
Legacy Incidents	Darktrace Threat Visualizer 5.2 introduces a significant development in the way AI Analyst incidents are grouped. Events are now linked persistently through meta-analysis, creating incidents with a wider scope that previously possible. Turn this setting on to see incidents created using (pre-Darktrace Threat Visualizer 5.2), or compatible with, the legacy method. This setting is not compatible with all filter options - some filters may be ignored when enabled.

Model Breach Threat Tray Options and Filters

Model breach alerts have a large number of additional, optional filters. These filters allow an operator to focus in on very specific activity and return to frequent queries using saved filter sets.



FILTER	DESCRIPTION
Search	A free text filter which will filter results by device or model name. The value of this field is not saved in saved filters and does not persist if the page is refreshed.
Model Folders	Darktrace models are sorted into folders that categorize the type of behavior they are designed to detect. When a folder is unticked, model breaches of models within that folder will disappear from the tray.
Subnets	This filter allows devices from specific subnets to be removed from the returned results. Subnets can be selected from a list or a custom range (such as 10.0.0.0/8) can be defined. When at least one subnet is selected, the filter can be set to “Include Subnets” (only show devices from these subnets) or “Exclude Subnets” (remove all devices from these subnets).
MITRE Tactics	Models curated and maintained by the Darktrace analyst team are mapped to the MITRE ATT&CK Framework. This filter allows the returned model alerts to be filtered down to only those models mapped to the specific tactic. Models can be mapped to multiple tactics but will not appear multiple times.

FILTER	DESCRIPTION
Include Acknowledged	Model breach and AI Analyst alerts can be acknowledged as part of an investigation workflow. By default, acknowledged alerts are removed from the returned results. Turn this setting on to include these results.
Darktrace RESPOND Only	If turned on, only model breaches for Darktrace RESPOND (Antigena) models will be included.
Darktrace RESPOND Controlled Devices Only	Restricts the results returned to just model breaches for devices that experienced a Darktrace RESPOND Action within the timeframe selected.
Sort By	The sort controls how model breaches are grouped into tiles. There are three modes: device, model, and user. The impact of the sort is covered in more detail in Model Breach Alerts in the Threat Tray.
Tags	In addition to being sorted into folders, models can also be tagged to group them together. Common examples are the “AP” tags which map models to attack phases. If a tag filter is added, model breaches will only be returned if the model has at least one of the chosen tags applied.

CHAPTER 2 - CYBER AI ANALYST

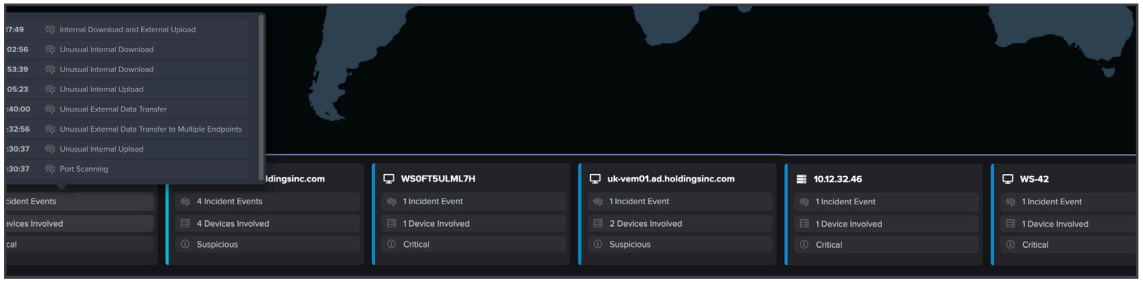
AI ANALYST INCIDENTS

CYBER AI ANALYST 31
CYBER AI ANALYST INCIDENTS 32
CYBER AI ANALYST INCIDENT LOGS 33

AI ANALYST INVESTIGATIONS

INVESTIGATING AN AI ANALYST INCIDENT 34
DETAILS OF CYBER AI ANALYST INCIDENT EVENT 35
UNDERSTANDING HOW AI ANALYST INVESTIGATED 36
LINKING AI ANALYST EVENTS TOGETHER 37
AI ANALYST INCIDENT OPTIONS AND ACTIONS 38

CYBER AI ANALYST



AI Analyst

AI Analyst incidents are an excellent place to start when investigating alerts for the first time in the Threat Visualizer. Incidents are only created for the most interesting and high priority activity observed by Darktrace.

Incidents are created from AI-powered investigations into anomalies across your digital environment. Every time the conditions for a model are met and a model breach is created, AI Analyst investigates the activity and concludes whether it needs to be surfaced for human analysts to review. The results of this analysis is then grouped in incidents.

What is AI Analyst?

The Darktrace AI Analyst investigates, analyzes and triages threats seen within your Darktrace environment. By learning from the millions of interactions between Darktrace's expert analysts and Darktrace DETECT components, the AI Analyst combines human expertise with the consistency, speed, and scalability of AI.

Darktrace models - a logic framework that allows operators to interact with the output from 'pattern of life' and anomaly classifiers - are used as a trigger to invoke AI Analyst. When the conditions for a model are met, a model breach is created; AI Analyst reviews and investigate all model breaches that occur on the system as a starting point for its analysis process.

The output from this analysis process are AI Analyst Incidents - a collection of one or more related events of anomalous activity. Incidents (Darktrace Threat Visualizer 5.2+) are formed through a meta-analysis of activity type, devices, and endpoints involved in each event. Each incident can encompass multiple stages of activity as it develops.

The layer of abstraction means only high priority activity is surfaced.

Incident to Model Breach Relationship

Although a model breach may be the trigger for an investigation, that does not mean the activity AI Analyst surfaces is directly related to the original model breach. Like a human, the Cyber AI Analyst uses a model breach as a starting point to investigate the device. The behavioral analysis it performs may discover anomalies or patterns of activity that were not the original trigger point for the model breach but are worthy of investigation.

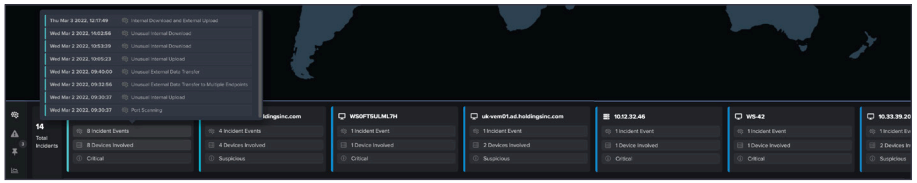
Similarly, not all model breaches that are investigated will result in an incident - only activity the AI Analyst considers high priority. A model breach will indicate if an AI Analyst incident was created as a result. Users who wish to review all alerts raised by Darktrace should also consult the model breach view.

Please note, the Cyber AI Analyst will only investigate model breaches of default Darktrace models. Custom models are not currently supported. From Darktrace Threat Visualizer 5 and above, AI Analyst investigations can be triggered by third-party alerts using a dedicated meta-model.

CYBER AI ANALYST INCIDENTS

Investigating an AI Analyst Incident

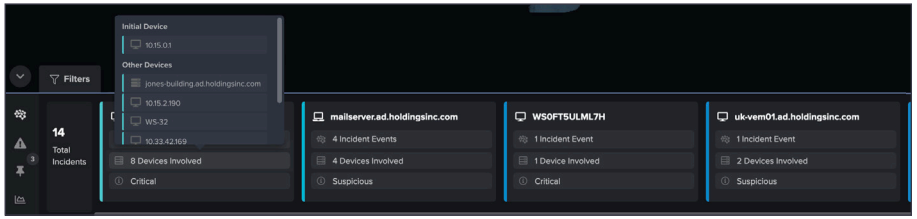
AI Analyst is accessible from the AI Analyst icon in the bottom left of the threat tray. For new users, it will be the default view when they first access the Threat Visualizer. The number of AI Analyst alerts in the tray (after filters) is displayed on the left.



Incidents have a behavior category and an overall score from 0-100%. Behavior categories are high level filters that allow an operator to focus in on specific levels of severity or behavior. There are three categories for AI Analyst incidents: Critical, Suspicious, Compliance. Categories are assigned to incident events and the incident is then given the most severe of all underlying categories.

The score can be used to filter on a sliding scale; incidents are categorized from dark to light blue, where light blue indicates a high score and darker blue a lower score. The score can be a useful way to prioritize within a behavior category - for example, starting with “critical” alerts and investigating from highest to lowest score before moving on to “suspicious”.

For more information about the filter options available, please see Filtering Alerts with Behavior Categories, Filtering Alerts with Basic Filters and Filtering Alerts with Advanced Filters.



Each alert in the tray lists the initial device (the device that triggered the first activity AI Analyst detected), the number of events the incident comprises, the number of devices involved in the incident, and the behavior category. If the initial device or user is subject to an active or pending Darktrace RESPOND action at the current time, a Darktrace RESPOND icon - - will appear beside the device name on the tile. The device name is highlighted green if the action is currently active or orange if currently pending.

Hovering over a tile will show more details of the devices and events.

Users can trigger AI Analyst investigations manually if they wish to look into a device’s behavior in greater detail. If an event or incident was created from an investigation triggered by a manual investigation or a third-party alert, this will be indicated on the tile.

Cyber AI Analyst Incidents

An incident is composed of one or more events; events are a group of anomalies or activity investigated by Cyber AI Analyst that it considers high priority enough to raise to a human operator for attention. For new users, start with the first incident in the tray with “critical” severity.

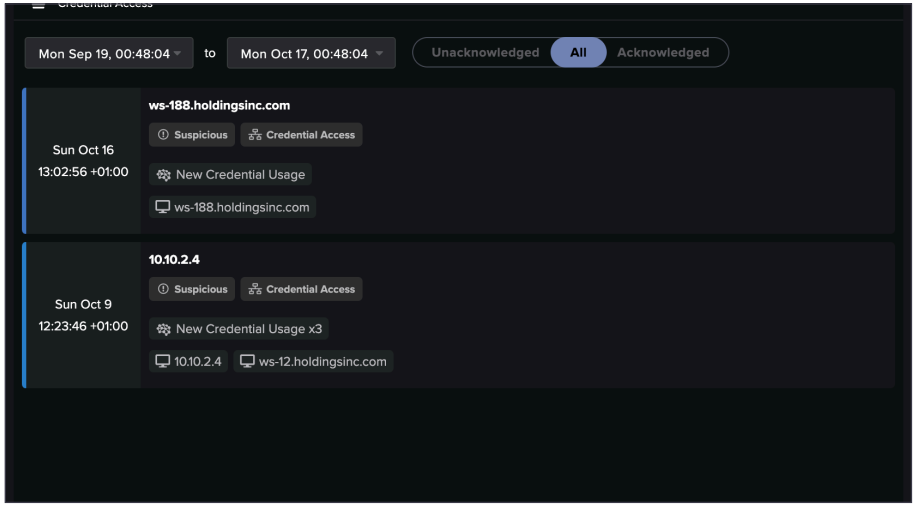
CYBER AI ANALYST INCIDENT LOGS

In addition to the Threat Tray, AI Analyst incidents can also be viewed in an AI Analyst incident log; the incident log is opened by interacting with the "MITRE ATT&CK Tactics Processed" statistics on the Threat Visualizer homepage summary.

The panel displays a count of events and alerts that are relevant to "tactics" from the MITRE ATT&CK Framework (***refer to Understanding the Summary Panel for more information***). For Darktrace Threat Visualizer 6, all types of activity AI Analyst investigates (event types) have been mapped to at least one MITRE ATT&CK Framework tactic. As AI Analyst incidents can comprise multiple events or types of activity, created incidents may therefore be associated with multiple tactics.

To open the AI Analyst incident log, click the "Incidents" or "Critical Incidents" count for a MITRE tactic.

AI Analyst Incident Log



The AI Analyst incident log opens a new window. The window title will reflect whether the log was opened from "Critical Incidents" or "Incidents". The MITRE tactic filtered upon is displayed in the top left of the window.

Underneath is the time range for AI Analyst incidents to be included in the log. The default range is taken from the 28 day summary panel range but can be altered using the time selectors. The log can also be filtered to show only unacknowledged AI Analyst incidents, only acknowledged AI Analyst incidents, or both states of AI Analyst incidents. By default, acknowledged alerts are removed from the returned results.

Incident Entries

Each incident appears as a separate entry and lists the initial device - the device that triggered the first activity AI Analyst detected - at the top of the entry. The time that the incident was first created appears on the left.

The entry also contains important information about the incident priority, involved devices, types of activity detected and the associated MITRE tactics.

EXAMPLE INCIDENT INFORMATION	MEANING
Critical	The incident behavior category - in this case, "Critical".
Exfiltration	A MITRE tactic associated with the incident.
Unusual External Data Transfer	The title of an event that is part of the AI Analyst incident.
ws192	A device that was involved in the incident.

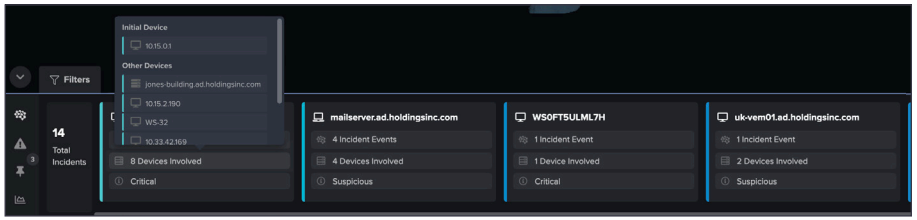
Click anywhere on the log entry to open the incident information, or click on the title of a specific event to open the window focused on that event. Incidents opened from the incident log are no different than those opened from the Threat Tray and can be investigated as described in Investigating an AI Analyst Incident and Details of Cyber AI Analyst Incident Event onward.

INVESTIGATING AN AI ANALYST INCIDENT

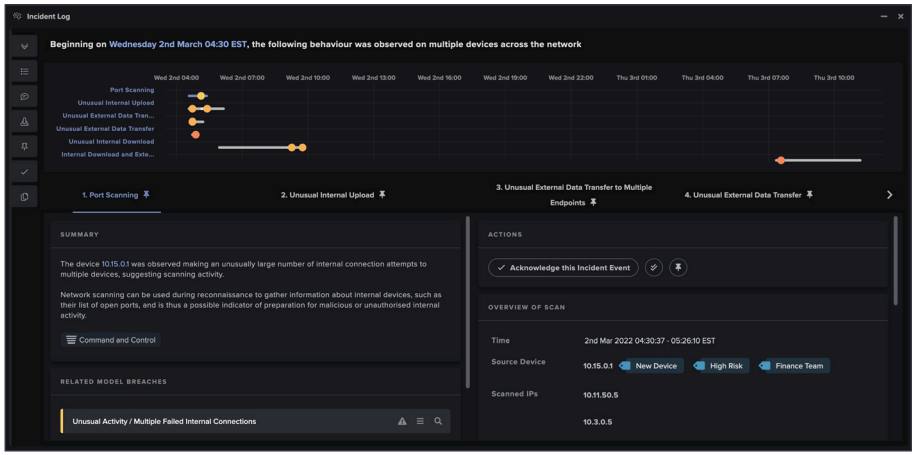
An incident is composed of one or more events; events are a group of anomalies or activity investigated by Cyber AI Analyst that it considers high priority enough to raise to a human operator for attention. For new users, start with the first incident in the tray with "critical" severity.

Investigating an AI Analyst Incident

Returning to our previous example, the first incident in the tray is a critical incident with 8 underlying events and 8 devices involved. The incident has been initially triggered by the behavior of a desktop device ("10.15.0.1"). Hovering over the tile, we can see that the incident is mostly unusual data transfer after potential port scanning activity. The devices involved are a mixture of desktop devices and servers.



For users investigating from an AI Analyst incident log, the same information is instead displayed on the log entry itself as described in Cyber AI Analyst Incident Logs.



Once clicked, the incident launches a Cyber AI Analyst pane with the results of the analysis and triage. It will open on the first event to occur as part of that incident. In our example, the first event is port scanning activity.

At the top of the panel is a representation of the incident time period. We can see that the first three events AI Analyst has highlighted occurred at around the exact same time. The device "10.15.0.1" was observed performing potential port scanning, making unusual large uploads internally and externally at 4:30am. An alternative view is also available from the ☰ incident timeline option on the left.



Model breaches which may be associated with each event will appear as dots, where color indicates severity. The activity associated with the event selected right now will be highlighted in blue; long-lived events such as large data transfers may cover a large chronological period.

Like a human, the Cyber AI Analyst uses a model breach as a starting point to investigate the activity. The behavioral analysis it performs may discover anomalies or patterns of activity that were not the original trigger point for the model breach but are worthy of investigation. Consequently, the event period may not correspond with the model breach time.

Additionally, some model breaches require sustained behaviors such as repeated connections, so the final model breach trigger may be later than the connection of interest. For this port scanning activity, the "Unusual Activity / Multiple Failed Internal Connections" model was triggered after repeated activity. The activity AI Analyst has identified therefore stretches back before the model breach.

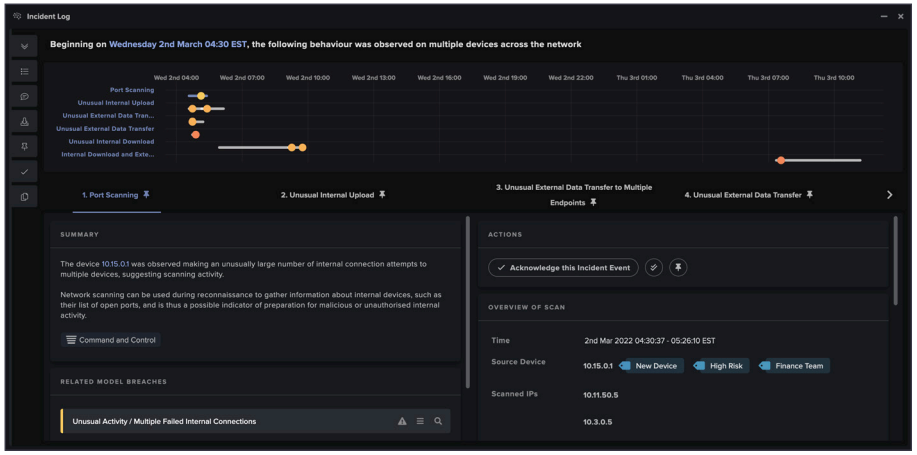
Events

The next step is to understand the activity AI Analyst has surfaced and understand why it has been highlighted to a human operator. Each incident is comprised of events which detail a specific anomalous activity - or chain of activities - investigated by AI Analyst. Each event will appear as a tab and can be investigated and acknowledged separately.

DETAILS OF CYBER AI ANALYST INCIDENT EVENT

Incident Events

Each incident is comprised of events which detail a specific anomalous activity - or chain of activities - investigated by AI Analyst. When starting to investigate an AI Analyst incident, one approach is to work through each incident event chronologically. Each event will appear as a tab and can be investigated and acknowledged separately. Other approaches which prioritize investigating the most high priority activity first may be more suitable, depending on your organizational priorities and workflow.



Continuing the example incident started in Investigating an AI Analyst Incident, the first event is potential port scanning activity.

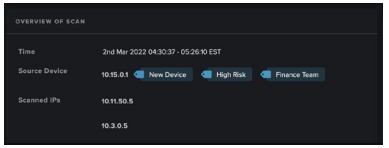
Event Summary

The first panel gives a summary of the event outline. The summary describes the type of activity AI Analyst detected, the device involved, any endpoints connected to and an explanation of why the activity is anomalous. The summary is the best way to quickly understand the activity AI Analyst has flagged for human review. Here, the summary explains why port scanning activity should be of concern to the security team.

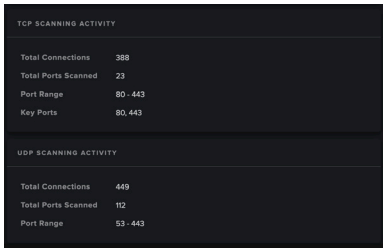
All AI Analyst event types are mapped to one or more MITRE ATT&CK Framework tactics - the relevant tactic appears below the sumamry text.

The summary can be localized into any of the current 12 options: English (US), English (UK), Chinese (Simplified), Chinese (Traditional), French, German, Italian, Japanese, Korean, Portuguese (BR), Spanish (ES) and Spanish (LATAM). This setting can be changed from **Account Settings**.

Event Details



The key details of the activity AI Analyst detected are provided in a series of panels on the right. The type of information is specific to each event type and can differ significantly.



For our example port scanning activity, the right panels include details of the scanned IP addresses, port ranges scanned via TCP and port ranges scanned by UDP. In other examples, the panels may include detailed information on the location of external endpoints or rare user agents detected performing SaaS administration.

AI Analyst will also utilize intelligence from Darktrace Attack Surface Management (ASM) through the Darktrace PREVENT/ASM Integration; if a malicious asset detected by ASM such as a domain is observed as part of an AI Analyst event, a pivot to that asset in the ASM interface will appear in the event details.

Where devices observed by Darktrace are involved in the event, detailed information about the device hostname, IP address and any tags will be included. These details will vary between device. If the device or user is subject to an active Darktrace RESPOND action at the current time, the device name will appear highlighted green. If clicked, this section will center the Threat Visualizer workspace on the device.

It can now be helpful to understand how AI Analyst detected the activity and how it links together with the other devices detected behaving anomalously.

UNDERSTANDING HOW AI ANALYST INVESTIGATED

How did AI Analyst reach the conclusions it did?

Incidents are created from AI-powered investigations into anomalies across your digital environment. Every time the conditions for a model are met and a model breach is created, AI Analyst investigates the activity and concludes whether it needs to be surfaced for human analysts to review.

For each event, AI Analyst explains the activity that initially triggered the investigation and the investigation process it followed to come to the conclusions surfaced in the Threat Visualizer.

Related Model Breaches

Underneath the event summary for an AI Analyst incident event are the model breaches that triggered the initial investigation. AI Analyst is not reliant on model breaches, they primarily provide a starting point for its deeper analysis of the device activity around the anomaly time.

A 💬 comment icon indicates if any related model breaches have been commented on by a user. These can be reviewed in the incident discussion.

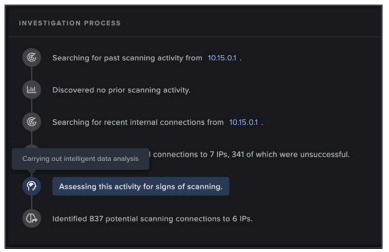
Further investigation prioritize the anomalous activity signposted by Cyber AI Analyst first and the model breach conditions second if the two do not directly align. For each model breach, a series of actions can be performed.

ICON	DESCRIPTION
⚠	This icon opens the model breach log for the related model breach. The Model breach log is covered in great detail in a series of articles, starting with Opening the Model Breach Log.
≣	Opens the model breach event log. Model breach logs are covered in more detail in Exploring the Model Breach Event Log.
🔍	Center the Threat Visualizer on the device that triggered the model breach at the time the model breach occurred.

Investigation Process

The investigation process panel is located below the related model breaches and explains how AI Analyst went from the initial trigger to the final activity presented to the user. Reviewing the process can be a useful way to understand the steps performed.

The Cyber AI Analyst, when triggered by a model breach, by telemetry data or by a human operator, initiates an investigation just like a human analyst. It retrieves and analyzes data to identify related behaviors and activity from the device or user under investigation, and from the wider network environment - just as a cyber analyst would seek to gain context and identify similar patterns of anomalous behavior elsewhere.



However, unlike a human analyst, the Cyber AI Analyst can analyze vast quantities of data at machine speed and investigate potential hypotheses concurrently. It can perform complex data interrogation on a scale and a complexity not possible without AI. The investigation process panel displays a summary of this intelligent analysis process followed by AI Analyst to reach the conclusions displayed in the event.

As the investigation process proceeds, AI Analyst narrows down the activity from a wider analysis to a targeted anomaly. Each investigation blends AI, supervised machine learning and intensive data analysis. These smart steps are identified in blue - "🧠 Carrying out intelligent data analysis".

As noted before, the behavioral analysis AI Analyst performs may discover anomalies or patterns of activity that were not the original trigger point for the model breach but are worthy of investigation.

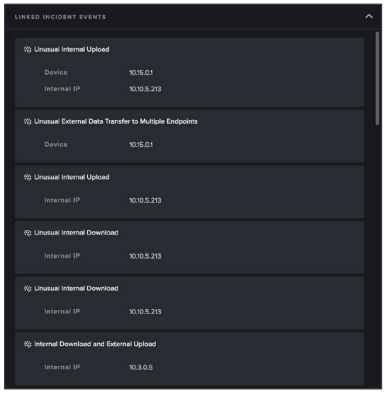
LINKING AI ANALYST EVENTS TOGETHER

The final step is to understand how the activity detected by AI Analyst in the event fits in with the wider incident.

Darktrace Threat Visualizer 5.2 introduces a significant development in the way AI Analyst incidents are constructed. New, 'open' incidents are created by a meta-analysis of the devices, endpoints and activity involved in each event. Events with linking factors are then joined together persistently to create incidents encompassing a much wider scope of time and activity.

Linked Incident events

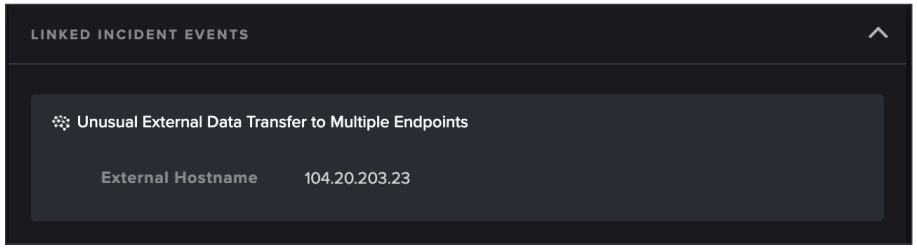
The linked incident events panel displays the reasons why AI Analyst chose to link certain behaviors together.



- The port scanning behavior has been connected to an "Unusual Internal Upload" and "Unusual External Data Transfer to Multiple Endpoints" because they were all performed by the same device ("10.15.0.1").
- This "Unusual Internal Upload" was also to an IP address seen in the scanning activity ("10.10.5.213"), which was then observed receiving an "Unusual Internal Upload" and sending data which triggered two instances of "Unusual Internal Download".
- Finally, another IP observed in the scanning behavior was seen performing an unusual "Internal Download and External Upload"

Seven events have therefore been directly linked to the activity in this event. Some events may not have a direct link to others and are instead linked through a shared event that draws the behavior together.

Switching to the final event - "Unusual External Data Transfer" - we can see from the linked incident events panel that it has been connected to the "Unusual External Data Transfer to Multiple Endpoints" due to a shared external destination. "Unusual External Data Transfer to Multiple Endpoints" was performed by "10.15.0.1", the same device that was observed performing the port scanning. Therefore, the activity has been linked together indirectly through this shared endpoint.



Returning to the example incident from Investigating an AI Analyst Incident and Details of Cyber AI Analyst Incident Event, eight events have been connected together.

Incidents within Incidents

Incidents can also be subsumed into other incidents - what might appear as disparate behavior initially can be linked by subsequent unusual activity. In this case, the later behavior is subsumed into the initial incident. Links to the later incident will redirect to the initial incident.

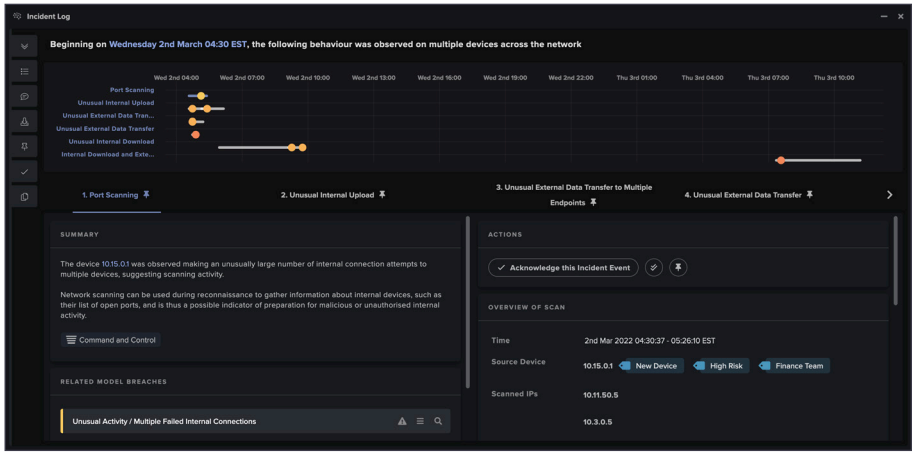
In a simplified example, device A is seen beaconing to two external rare endpoints. Device B is separately seen performing a chain of administrative SSH connections that end at a key server. These activities trigger separate investigations and result in two separate incidents. Subsequently, the credential `admin12` is observed on device A and triggers an investigation for unusual admin credential use. At the same time, the credential `admin12` is also seen on the server connected to by device B, triggering an investigation for unusual admin credential use. The event for device A and device B are now linked by a shared admin credential - `admin12` - and their previous activity is drawn into one overall incident. Where the beaconing and administrative connections were disparate events at first, the linking credential allows AI Analyst to present the full picture: device A and B are compromised by the same actor, and device Bs connectivity is lateral movement associated with that compromise.

AI ANALYST INCIDENT OPTIONS AND ACTIONS

To investigate an AI Analyst event, the next step is to look more closely at the activity and devices involved. Before focusing on a device, it is necessary to understand the actions and investigation steps that can be taken in the AI Analyst incident pane itself.

Incident Actions

On the left of the incident panel are a series of actions and views. During initial investigation, the incident timeline can be used to understand the order of events and filter down by specific devices. The attack phases also visualizes the events alongside the stages of a cyber attack they may represent.

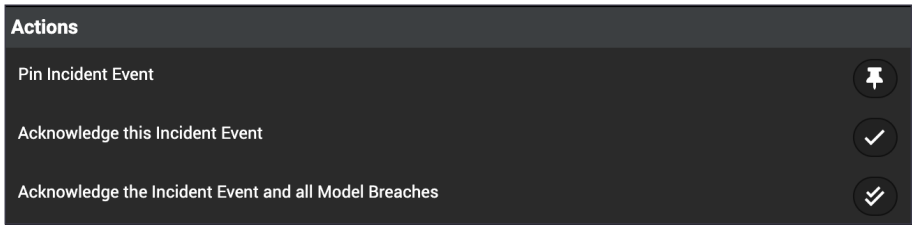


Pinning the incident, commenting or downloading a PDF report can be useful for collaboration with other analysts on event investigation.

ICON	DESCRIPTION
Breaches Processing	A flashing Cyber AI Analyst icon indicates that it is currently analyzing other model breaches that may be relevant. If this analysis deems the breaches related to the incident, Cyber AI Analyst will create additional events. Click the icon to expand the list of processing breaches.
Incident Timeline	This option opens a timeline of all events and their trigger device - it is useful to quickly understand the chronology of the incident and find events that occurred around the same time. A device filter can be applied in this pop-out which also applies to events returned in the main panel.
Attack Stages	Open a visualization of the events alongside the stages of a cyber attack they may represent.
View Incident Discussion	Opens the comment pane where users can discuss the incident. Comments on related model breaches will also be displayed in the discussion. Users may comment from the Threat Visualizer interface as well as the mobile app.
Create a Report	Opens a dialogue to create a downloadable PDF report from the incident. Any text entered into the “Report Name” field will be used as the filename and appear as a title in the generated PDF.
Pin the Incident	Keeps the Incident at the left-hand-side of the threat tray regardless of the time range chosen. When the incident is pinned, the tab will appear lighter. Individual events may also be pinned - if one or more events are pinned, the entire incident will be pinned. Pinned incidents do not interact with the pinned model breaches functionality.
Acknowledge Incident	This icon will acknowledge all underlying events for the incident at the time of click. If all events are acknowledged, the incident will be shown/hidden from the incident tray depending on the filter settings. Events can also be acknowledged individually as part of an incident investigation workflow - some individually acknowledged events do not prevent the overall incident being returned.
Copy to Clipboard	Copies a link to the incident to the clipboard.

Actions for an Event

On each event are a series of actions that can be taken. These actions only affect the chosen event or related alerts.



ICON	DESCRIPTION
✔ Acknowledge this event	Events can be acknowledged individually, rather than as part of the whole incident - this is a common workflow as a threat analyst investigates each underlying event in turn.
✔✔ Acknowledge this event and related model breaches	In addition to acknowledging this event, this icon will also acknowledge the model breach alerts listed under the “Related Model Breaches” heading. A comment is automatically added to model breaches acknowledged this way in bulk.
📌 Pin/Unpin incident event	Click this icon to pin or unpin just this event. Incidents with at least one pinned underlying event will be returned regardless of whether they fall within the time range chosen.

Incidents and events can be acknowledged when they have been investigated. This removes them from the threat tray by default. A common workflow is to investigate each event and acknowledge individually.

Investigating activity

Once the outline of the event and incident as a whole is clear, focus on a device within the incident to better understand its normal activity patterns and gather context about its role in the digital environment. Investigating a device is introduced in Focusing on a Device.

CHAPTER 3 - INVESTIGATING A DEVICE

FOCUS ON A DEVICE

FOCUSING ON A DEVICE 41
DEVICE INVESTIGATION OPTIONS 42

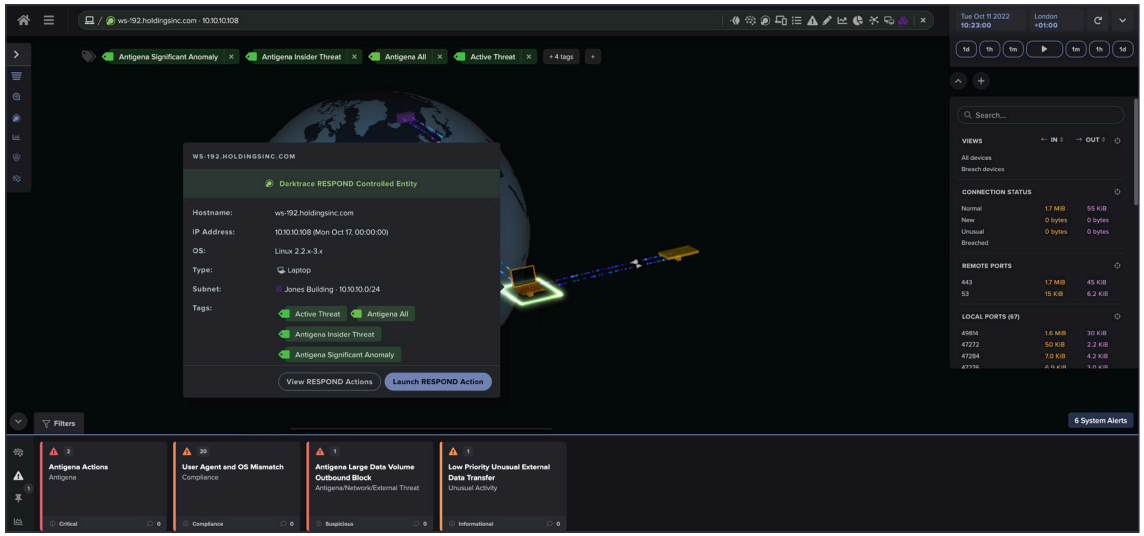
DEVICE EVENT LOGS

DEVICE SUMMARY 44
OTHER DEVICE SUMMARY SECTIONS 46
USING THE DEVICE EVENT LOG 48
TYPES OF EVENT IN THE DEVICE EVENT LOG 49
INVESTIGATING A CONNECTION IN THE DEVICE EVENT LOG 50
INVESTIGATING SAAS ACTIVITY IN THE DEVICE EVENT LOG 52
FILTERING THE DEVICE EVENT LOG 54

OTHER INVESTIGATION METHODS

TRIGGERED AI ANALYST INVESTIGATIONS 55
OTHER DEVICE INVESTIGATION WINDOWS 56

FOCUSING ON A DEVICE



Investigating a Device

Focusing on a device is the next stage in any investigation workflow. To understand whether an anomalous activity might be malicious, it is important to understand the device and to acquire context about its historic behavior.

There are many ways to focus the Threat Visualizer workspace on a device. Three common methods are:

- From an AI Analyst incident event, click the device name in blue above the graph. Click the time in the same subheading to also set the Threat Visualizer to that exact time.
- From a model breach log, click the 🔍 magnifying glass icon. This will center the Threat Visualizer on the device that triggered the model breach at the time of the model breach.
- From the subnet view, click any glowing point that represents a device.

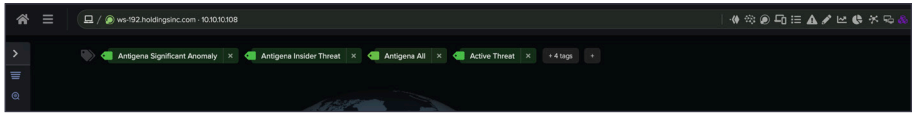
Workspace

The workspace will change. In the center is a visualization of the chosen device and any other devices it is connected to at the time chosen. The time in the selector at the top right controls the visualization. If the ▶ play button is clicked on the time selector, time will progress and the re-enactment of connections will resume. Clicking on any other device will change the focus to it.

Hover over the device to see a summary of information including device type, current subnet, any applied tags, any recently observed credentials and to trigger a manual Darktrace RESPOND action agains the device. See Manual Darktrace RESPOND Actions for more information on triggering these actions.

The color of the device is decided by the highest severity of model breaches that it has triggered in or around the time selected in the top right. If the device was subject to a Darktrace RESPOND action during the visualization, a green line will appear around it for the duration of the action. If the device currently has pending Darktrace RESPOND actions, an orange line will appear around the visualization.

The icon for the device is decided by its device type. Device type is derived by Darktrace from observation of the device behavior - it can be overwritten in Device Admin (see Device Admin) or from the icons in the omnisearch bar (see below).



Below the device options are any tags that have been applied to the device. Tags offer a simple labelling system for network devices, credentials and users detected in external platforms (SaaS users) which can be used to identify important resources, control eligibility for Darktrace RESPOND responses and alter model logic . More information can be found in Adding and Reviewing Tags.

On the right of the workspace is a panel with statistics about the device's connections over the time window chosen in the selector. The time window is shown below the timezone. By default, the time window is 5 minutes.

DEVICE INVESTIGATION OPTIONS

Device Options

At the top of the workspace, in the omnisearch bar, the device focused upon is shown. The bar displays the device label, hostname and IP address, where available. For SaaS user devices, the device name is displayed in the format `[SaaS::][service] [username]`. For devices monitored by Darktrace DETECT & RESPOND/Endpoint agents, only the hostname and label are displayed as the device may have many concurrent IPs.



On the left of the device name is an icon representing the device type. The device type is automatically derived and can be manually overwritten.

To the right of the device name are a series of icons; these icons and options will differ depending on the type of device.

Darktrace RESPOND

If the device or user is subject to a Darktrace RESPOND action at the current time, the omnisearch will appear shaded green. If an action is currently pending confirmation for the device, the omnisearch will appear shaded orange. A Darktrace RESPOND icon - - will also appear beside the device type.

Device Options

ICON	APPLICABLE TO	DESCRIPTION
Open Device in ASM	Network	Where the Darktrace PREVENT/Attack Surface Management integration is present, matched devices will display this icon. Click to pivot to the matched asset in ASM.
Open Device in E2E	All	Where Darktrace PREVENT/End-to-End is configured, click to pivot to the device in the E2E interface.

ICON	APPLICABLE TO	DESCRIPTION
AI Analyst Investigation	All	AI Analyst investigates anomalous activity in your digital environment, surfacing only the most high priority findings. Users can trigger their own AI Analyst investigations into devices and users detected in external platforms (SaaS users). This icon opens a dialog to trigger an AI Analyst investigation into the focused device.
Darktrace RESPOND Actions	All	Darktrace RESPOND takes autonomous actions against devices to control and contain anomalous behavior. This option opens the Darktrace RESPOND Actions window, focused on the device. The default timeframe is 7 days but can be altered. This window can also be used to trigger manual Darktrace RESPOND actions against the user account for eligible modules. This option will not appear for users of modules which do not have Darktrace RESPOND available.
Device Summary	All	This icon opens the device summary, a key window for understanding statistics about a device and its historic behavior. Device summary is covered in more detail in Device Summary and Other Device Summary Sections.
Device Event Log	All	The device event log is a detailed log of device connectivity, activity, changes, and anomaly notifications from Darktrace analysis. it is a key investigation tool. The device event log is covered separately in more detail, starting with Using the Device Event Log.
Model Breach Log	All	This icon opens the device breach log for the device - a list of model breaches triggered by the device. The default timeframe is 7 days but can be altered. Breach logs in general are covered in more detail in Opening the Model Breach Log.
Edit Info	Network, Endpoint	For network and endpoint devices, this option allows a label (or nickname) to be set for the device. The device type and priority can also be altered. The priority impacts the score for models the device triggers. These settings can also be changed for multiple devices in Device Admin.

ICON	APPLICABLE TO	DESCRIPTION
Open Graph	All	The graph allows metrics of device behavior to be plotted over time. The default metric for non-SaaS devices is connections and activity for SaaS devices. The graph is described in more detail in Other Device Investigation Windows.
Open Profile in SaaS Console	SaaS	The Darktrace SaaS Console is a purpose built interface for investigating SaaS and cloud user behavior. This icon opens the equivalent user profile in the SaaS console to continue investigation there.
View Connections Data	Network, Endpoint	The connections data window visualizes the connectivity of a chosen device against those Darktrace has derived as peers. This is described in more detail in Other Device Investigation Windows.
View Similar Device Map	Network, Endpoint	For every device, Darktrace identifies peer devices - those it derives as having a similar pattern of life. The similar device map is a 3D visualization of these similar devices, clustered by how similar they are. The list of peer devices is also available in device summary and by clicking the “view similar devices” option.
View Similar Devices	Network, Endpoint	As above, for every device, Darktrace identifies peer devices - those it derives as having a similar pattern of life. This option opens a clickable list of similar devices.
Go to device subnet	Network, Endpoint	Returns the focus to the subnet view for the subnet the device is in.
Monitored by Darktrace/Endpoint	Endpoint	This icon is not interactive - it indicates that the device is monitored by a Darktrace DETECT & RESPOND/Endpoint cSensor agent.

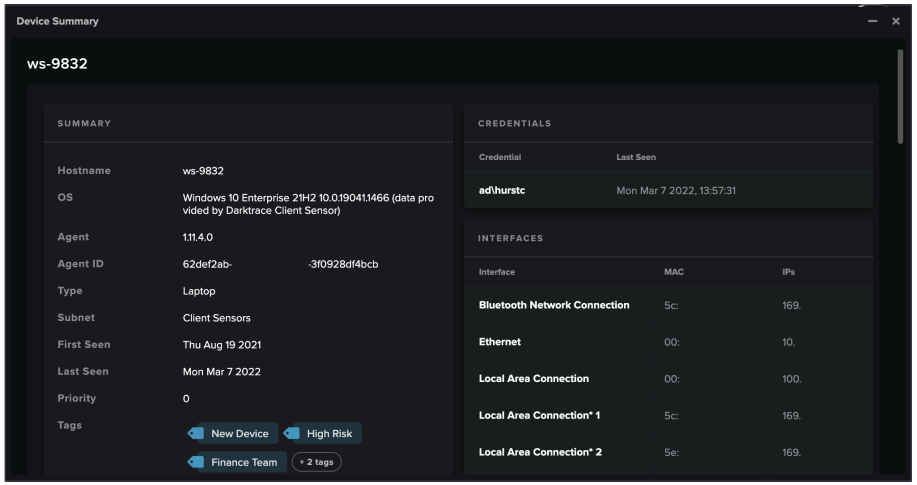
DEVICE SUMMARY

When investigating a device, it is important to understand the context of how it behaves and where it fits within the network. Device summary is a key tool for gathering this information quickly. It displays key information about devices, their historic behavior, any acknowledged or unacknowledged alerts and similar peer devices.

The timeframe for connectivity data is 1 month and 7 days for credential history. For summary data and interface data, the information is the most up to date at the time the summary is opened. The timeframe for other elements can be altered on the section itself.

Summary

The summary section appears for all device types. It contains key information such as IP address, hostname, the time the device or user was first seen, the device priority and any tags applied to it. The information that appears in the summary differs between network, endpoint and SaaS devices. For example, devices monitored by Darktrace DETECT & RESPOND/ Endpoint cSensor agents will include information about the software version of the agent itself.



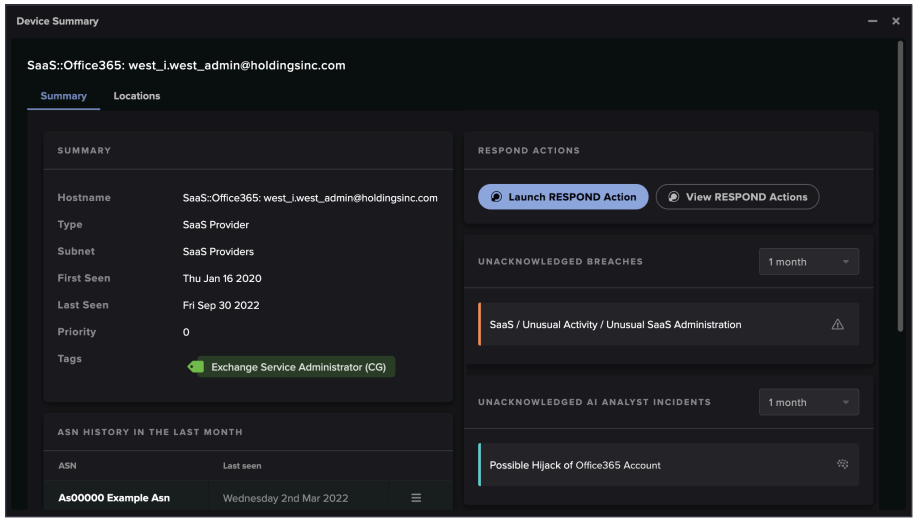
Darktrace will attempt to derive operating system information for a device from passive observation of its behavior and display it in the summary. For devices monitored by Darktrace DETECT & RESPOND/Endpoint cSensor agents, this information will be populated directly by the agent. The Darktrace Defender Advanced Hunting integration also populates OS information for devices observed by both Darktrace and Microsoft Defender in the summary.

RESPOND Actions

If Darktrace RESPOND is licensed for the type of device selected, the RESPOND Actions section will appear.

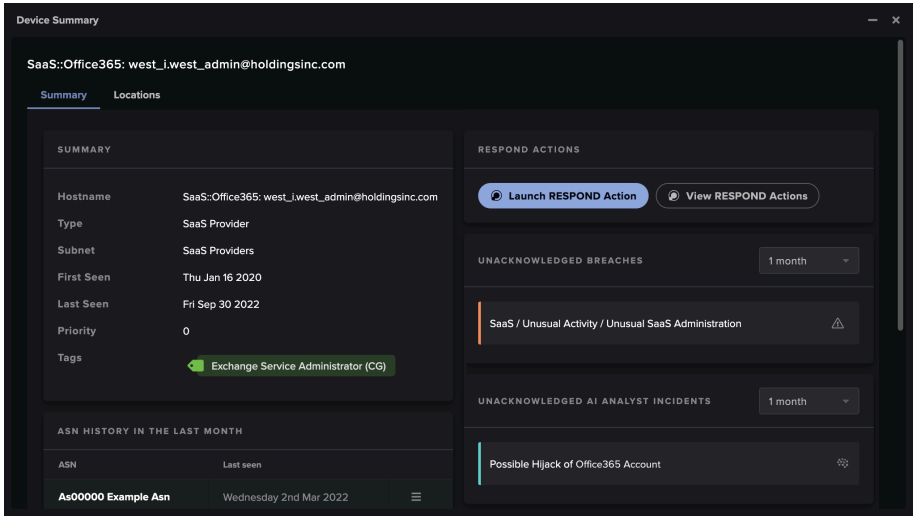
- Click ""Launch RESPOND Action" to trigger a manual Darktrace RESPOND action on the device. See Manual Darktrace RESPOND Actions for more information on triggering these actions.

- Click ""View RESPOND Actions" to open the Darktrace RESPOND Actions window filtered on the chosen device. Refer to Reviewing Darktrace RESPOND Actions for more information on how to understand the entries in this window.

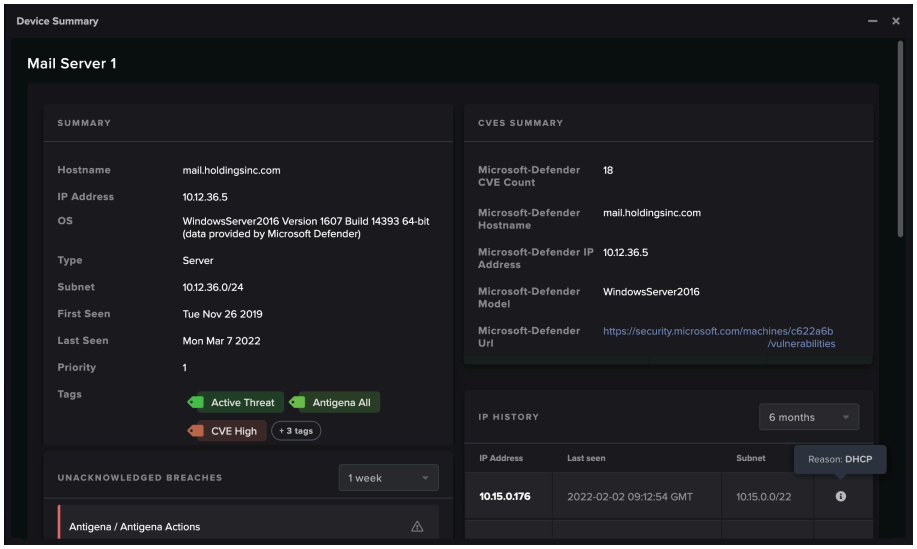


IPs, Credentials and ASNs

Device summary populates information about recent IP addresses, credentials, ASNs and interfaces associated with the device. Credentials associated with network or endpoint devices will be displayed first, if present. The historic information will differ between network, SaaS and endpoint devices and will only be returned where possible. For example, if no credentials have been associated with the device, this subsection will not appear.



Historic IP information is displayed for network devices - the timeframe can be extended up to 6 months. For each IP change, the reason why a new IP was associated with the device is provided. Devices monitored by Darktrace DETECT & RESPOND/Endpoint cSensor agents can have multiple concurrent IPs due to multiple network adapters - instead, details of all network interfaces observed by the cSensor agent (both IP and MAC Address) are listed.

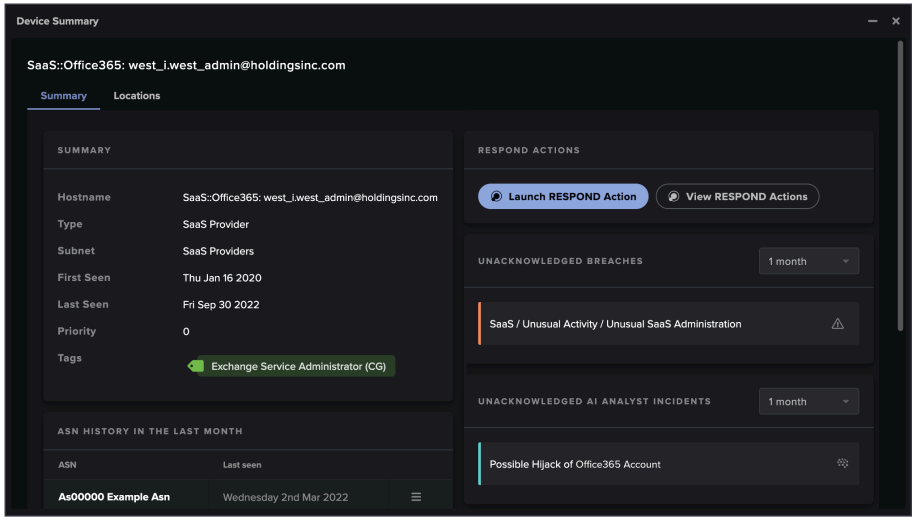


For SaaS user devices, ASN history - the ASNs for IP Addresses seen connecting to the SaaS account - is provided. An ASN is a group of IP addresses controlled by one Internet Service Provider or entity. As SaaS activities are often performed by users on mobile devices or tethering to mobile devices, their IP may change regularly but the ASN will stay consistent. Greater variation in IP address is therefore not as indicative of anomaly as large changes in ASN usage. Click the ☰ log icon to open a log of activities for the user from that ASN.

OTHER DEVICE SUMMARY SECTIONS

Alerts

Device summary lists all historic AI Analyst incidents and model breach alerts for the device, separated into acknowledged and unacknowledged. The default timeframe is one week but can be extended to 6 months, data dependent.



For model breaches, hover over the model name for a description tooltip. Click the ⚠ model icon to view the model breach log for the model and the selected user.

For AI Analyst incidents, click the AI Analyst icon to open the incident event.

Other Sections

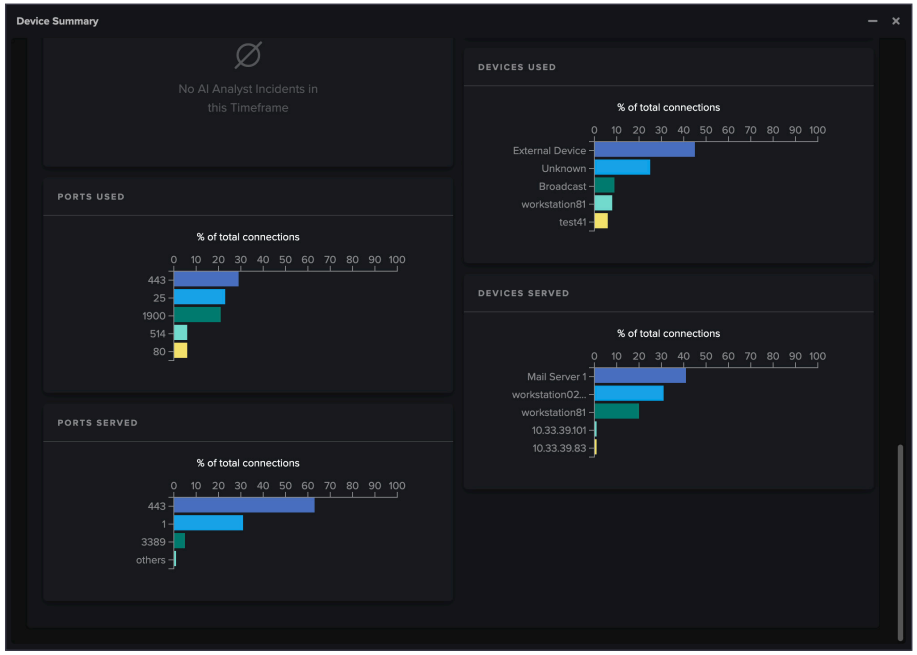
Similar Devices

A list of similar peer devices is returned for network and endpoint-monitored devices. Similar devices are devices Darktrace has derived as having a similar 'pattern of life' - the list is also accessible from "View Similar Device Map" and "View Similar Devices" in device omnisearch options.

To switch Threat Visualizer focus to one of these similar devices, click the 🔍 magnifying glass icon beside the device name. To view model breaches triggered by the peer device, click the ⚠ model icon.

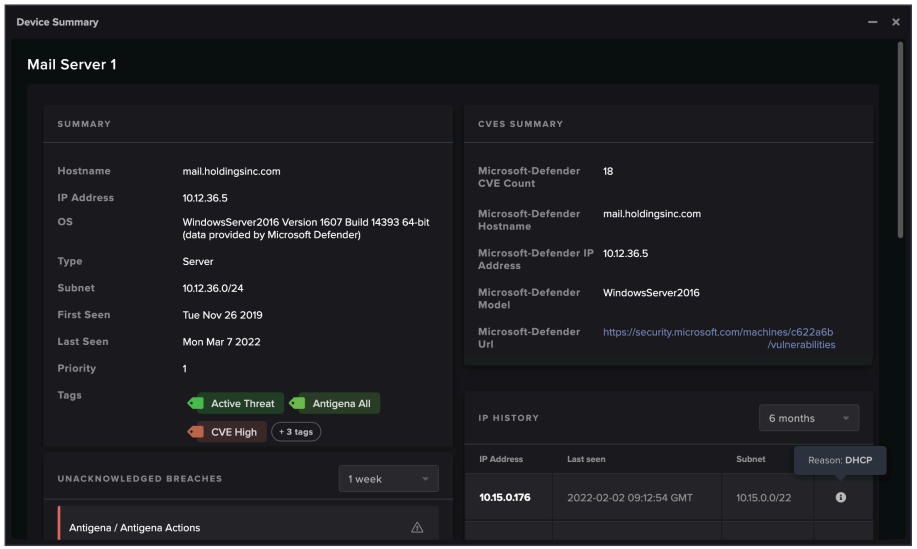
Connectivity

For network and endpoint devices, a breakdown of the top five endpoints and ports the device has connected to/from will be displayed. This data is calculated over the last month.



Vulnerability Data

If the Darktrace Defender Advanced Hunting integration or integrations with Tenable and Rapid7 vulnerability scanners are configured, CVE data for devices observed by both platforms will be populated in an additional device summary section.



Geolocation

For SaaS user devices, an additional tab is displayed. The location tab displays a map of the approximate geolocation for all IP addresses that have performed SaaS or cloud activities for this user in the last month.

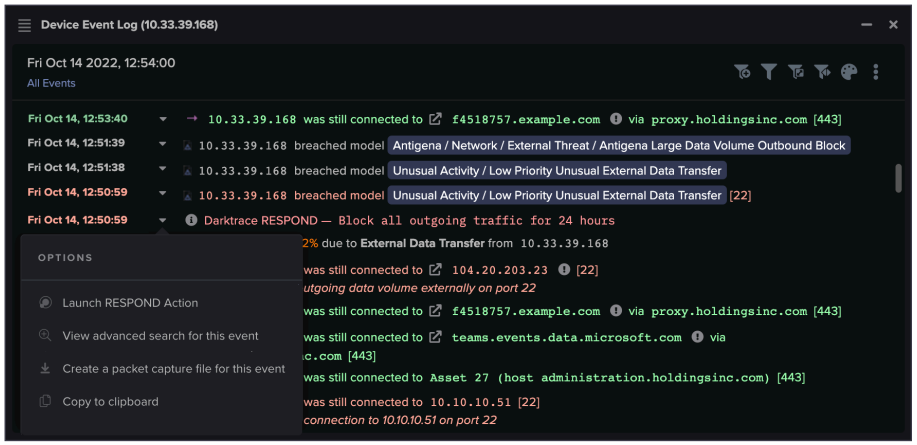
IP Addresses are grouped by ASN and the ASN location is represented as a point on the map. Light blue locations indicate sustained activity and dark blue locations are those which are rare for the user.

USING THE DEVICE EVENT LOG

The device event log is a detailed list of activity and connections in chronological order. When investigating an alert, reviewing the behavior of the device before and after the alert can be useful to acquire important context.

Understanding the Device Event Log

Each log entry is sorted in time order, working backwards from the time selected in the top right. Each log entry describes the device and the activity performed. For network devices, this might be a connection from the device IP to a specific host on a specific port. For SaaS user devices, the log is structured to show the action, the user and the location it was performed from. The event log also contains contextual information like changes to device information, commentary from Darktrace anomaly detection and model breaches triggered.



Between the time stamp and log line is an ▾ arrow icon with extra options:

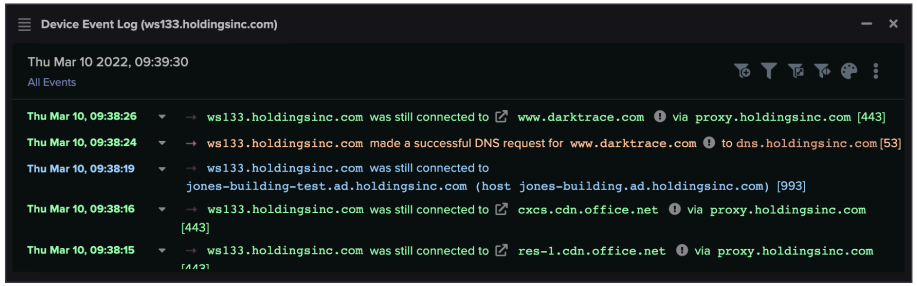
OPTION	DESCRIPTION
Launch RESPOND Action	For network and endpoint devices, trigger a RESPOND action populated with the details of the event log line. Particularly useful for the "Block Matching Connections" inhibitor type. See Manual Darktrace RESPOND Actions for more information on triggering these actions. Requires a relevant Darktrace RESPOND license to be present.
View advanced search for this event	This options opens the Darktrace Advanced Search interface to explore a more detailed version of the event or connection log for granular analysis For some events there may be more than one search result with additional contextual information. More information about Advanced Search can be found in Darktrace Advanced Search.
Create a packet capture file for this event	Opens a dialog to create a packet capture device for the specific chosen connection (see Creating Packet Captures). This option is not available for SaaS user devices.
Search Defender for Process Details	If the Darktrace Defender Advanced Hunting integration is configured, users can request process data from Microsoft Defender for connection events. This option will only appear for connections older than 15 minutes for devices observed by both Darktrace and Microsoft Defender. For more information, please see Darktrace Microsoft Defender Advanced Hunting Integration.
Search Defender for User Logins	If the Darktrace Defender Advanced Hunting integration is configured, users can request a list of users who logged into the device around the time of the activity from Microsoft Defender. This option will only appear for devices observed by both Darktrace and Microsoft Defender. For more information, please see Darktrace Microsoft Defender Advanced Hunting Integration.
Copy to clipboard	Copies the exact log line to the clipboard.
Visualize connection history	This option opens a 3D visualization of historic similar connections such as those on the same port, those to the destination device and those for the connection pair. This option is not available for SaaS user devices.
Filter Event Log by this [protocol/ port/IP]	Multiple options to filter the log by specific aspects of the connection will appear. Granular filters can be applied to the event log to focus on specific behavior; these are discussed in more detail in Filtering the Device Event Log. This option is not available for SaaS user devices.

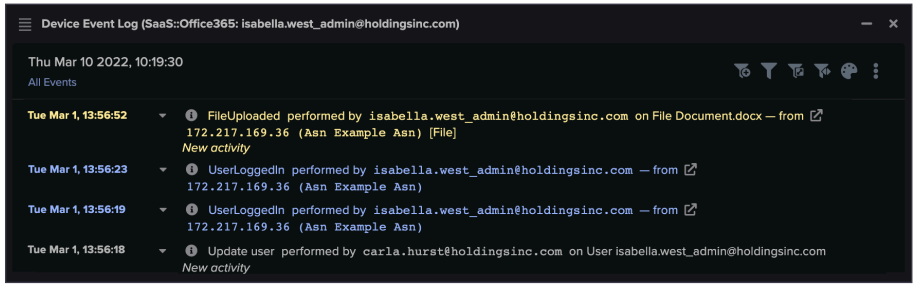
TYPES OF EVENT IN THE DEVICE EVENT LOG

Entry Types

There are seven “types” of event that can appear in the logs - these events can be filtered on or excluded to focus on specific activity.

TYPE	DESCRIPTION
Connections	Connection entries are indicated by an arrow. A ← left-facing arrow indicates the connection was outbound. A →right-facing arrow indicates the connection is inbound. An arrow with a cross indicates a failed connection. A flashing arrow means the connection is ongoing at the Threat Visualizer time selected.
Unusual Connections	Unusual connections are displayed in the same manner as connections but include commentary below the log line with the reason Darktrace has deemed the connection unusual.
New connections	New connections are displayed in the same manner as connections but include commentary below the log line describing why it has been identified as new.
Unusual Activity	Unusual activity entries are indicated by the Darktrace logo. Darktrace monitors behavior metrics for each device as part of the pattern of life. If the value of one of these behavior metrics is anomalous, an unusual activity entry may be added to the device event log. The metric is clickable and will open the metric graph.
Model Breaches	Model breach entries are indicated the triangle icon. If a device triggers a model breach during the timeframe of the log, an entry is created. The model name can be clicked to open the model breach log for further investigation. Some models can be triggered without creating a model breach; these models will also be returned but with limited information.
Notices	Notices are indicated by the info icon. Notice entries can be events or contextual information - SaaS user activity falls under this category. Darktrace RESPOND actions in the event log will also appear in this way. Clicking the info icon will open a detailed tooltip. Notices can be “popped out” to refer back to.
History	History entries are indicated by purple text. History events are created when information about the device changes, such as IP, contextual information and credential usage.



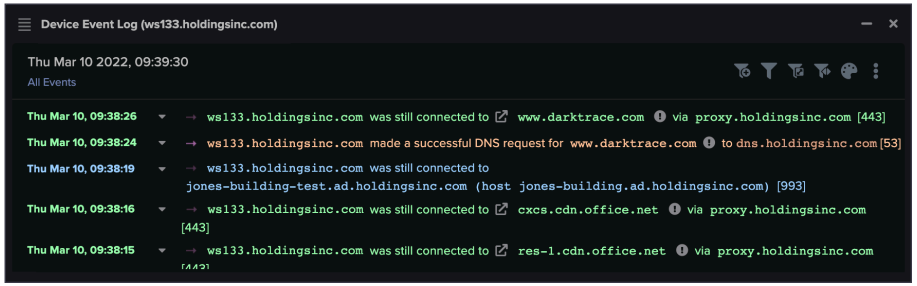


INVESTIGATING A CONNECTION IN THE DEVICE EVENT LOG

Investigating a Connection Entry

A typical network or endpoint device will have a log made up mostly of connection events. Connections are indicated by an arrow pointing in the direction of the connection. Each connection entry is broken up into important parts.

For example:



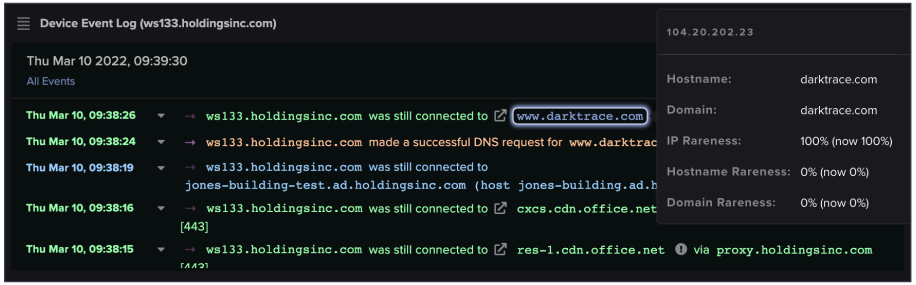
← ws123.holdingsinc.com was still connected to darktrace.com via proxy.holdingsinc.com [443]

First is the connection direction - a ← left-facing arrow indicates the connection was outbound, a → right-facing arrow indicates the connection is inbound. An arrow with a cross indicates a failed connection. A flashing arrow means the connection is ongoing at the Threat Visualizer time selected.

Click the direction arrow for detailed information about the connection such as source and destination IPs, data transferred and protocol. This info can be "popped out" in a separate window to refer back to.

Next, hover over the name of the device for a summary tooltip with basic information about the device. The first device in the log line should be the device currently focused on.

After the device is a simple description of the connection status. Examples include "connected to", "made a successful DNS request for", "was blocked from connecting to". The connection direction is part of this status: "was still connected to by" indicates our chosen device is the destination, "connected to" means it was the source.

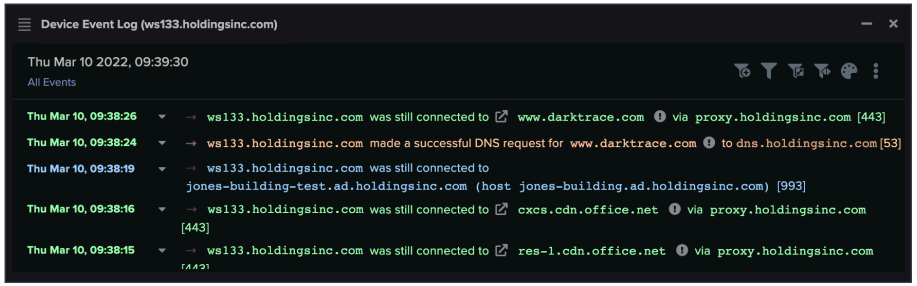


Following the status is the endpoint, device or host that was connected to (or the device was connected to by).

- If the source or destination is an internal device, hover over the name of the device for a summary tooltip with basic information about the device. Click the device name to change the Threat Visualizer focus to this other device.
- If the source or destination is an external hostname, multiple interactions are available:
  - Hover over the endpoint or host for a summary tooltip with basic information about the rarity of the endpoint in the context of the network.
  - Click the link icon to open the external sites summary. Similar to device summary, this summary gives information devices that have connected to the external location. More information is available in External Sites Summary and External Sites Summary for IPs and Hostnames.
  - Click the hostname or endpoint name to open an additional event log focused on that host.
  - If workflow integrations that allow lookups for external IPs and hostnames are enabled (such a VirusTotal), a pivot icon will appear for these entries.

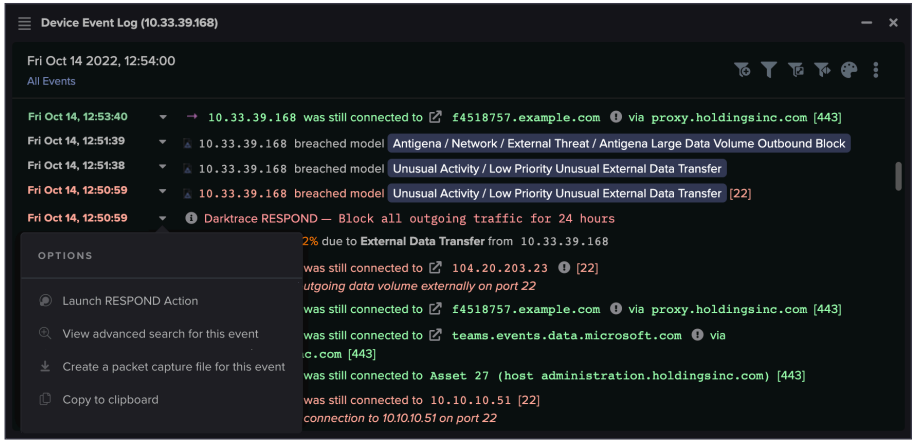
If a proxy is in use, it will appear here as a “via”.

Finally, depending on the color coding chosen, extra information such as protocol, port or rarity may be displayed at the end of the log line. For network connections, color coding is available by protocol, application protocol, source/destination port, by notice and by endpoint rarity. The default color coding is controlled in user Account Settings or can be altered using one of the filter options.



Darktrace RESPOND

Devices currently subject to Darktrace RESPOND actions will appear shaded green in pop out windows. Devices with pending actions will be shaded orange. Hovering over these devices will also show that they are either currently controlled or pending control.



If you wish to launch a RESPOND action against network entities in a connection, there are two ways to do so.

- From the hover-over summary, click “Launch RESPOND Action”
- From the ▾ down arrow, chose “Launch RESPOND Action” to create a “Block Matching Connections” action from the connection itself.

See Manual Darktrace RESPOND Actions for more information on triggering these actions.

Darktrace/Email

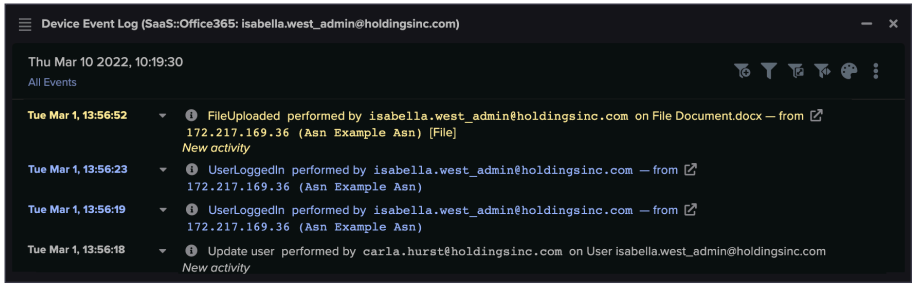
Darktrace/Email shares information about the endpoints it sees in organizational email traffic with other Darktrace DETECT and RESPOND components. If a connection entry includes a hostname Darktrace/Email has detected in an email in the last 12 months, an ✉ email icon will appear beside that connection entry.

Clicking this icon will attempt to retrieve relevant emails from Darktrace/Email that feature the hostname over the last seven days. As information about observed hostnames is retained for 12 months, it is possible that the relevant email communications are no longer present in Darktrace/Email.

INVESTIGATING SAAS ACTIVITY IN THE DEVICE EVENT LOG

Investigating a SaaS Activity Entry (Notice)

Log lines for SaaS and cloud users are structured slightly differently to connection events. The log lines take the format of event, actor, resource, source. For example:



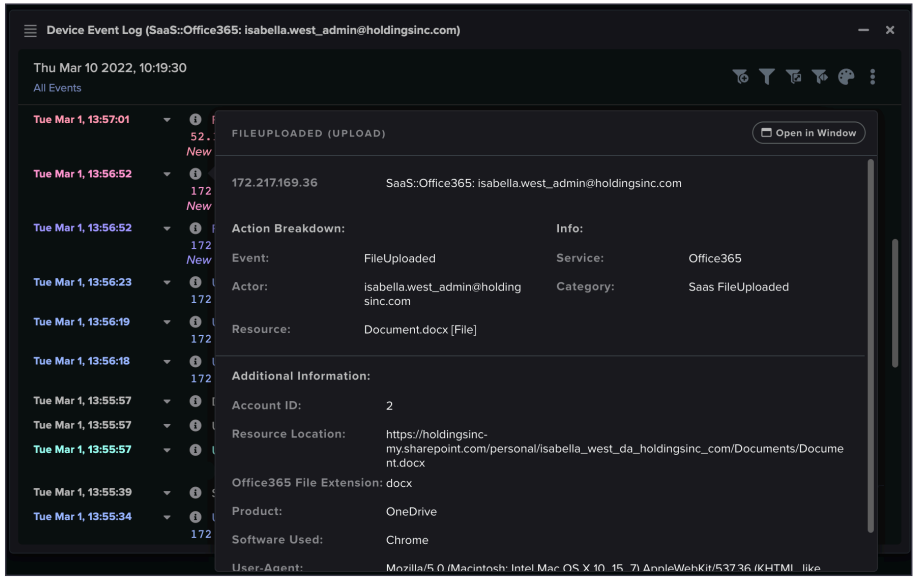
ⓘ Login performed by isabella.west_admin@holdingsinc.com - from 172.217.169.36 (ASN Example ASN)

ⓘ FileUploaded performed by isabella.west_admin@holdingsinc.com on Document.docx - from 172.217.169.36 (ASN Example ASN) [File]

ⓘ Update User performed by carla.hurst@holdingsinc.com on User isabella.west_admin@holdingsinc.com - from 172.217.169.36 (ASN Example ASN)

The ⓘ info icon that appears first can be clicked to open a tooltip with detailed information about the event, the user and the resource that was changed. This info can be "popped out" in a separate window to refer back to.

Users can perform a wide range of actions across different SaaS platforms; for ease, events are grouped into high level categories such as "SaaS Login" or "SaaS Resource Modified". Click the event to open an event log of all actions in the same category across all users. The SaaS-specific "event" color coding allows rare events to be distinguished amongst frequent logins or file actions.



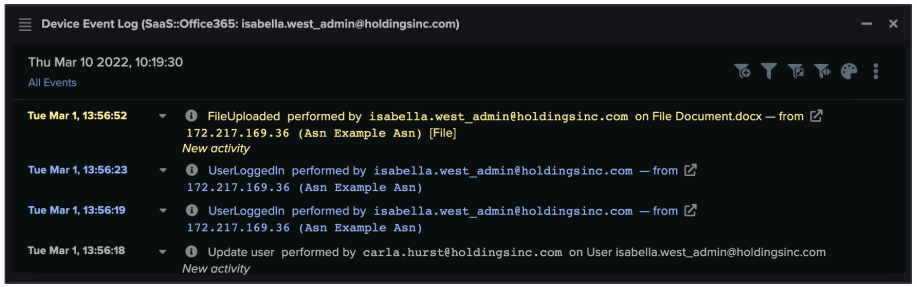
After the event is the *actor* - the user who performed the action ("performed by"). In most cases, this user will be the same user the Threat Visualizer is focused on. However, users can also change things about other users such as an admin changing their permissions. In this case, the SaaS user device currently focused on is the resource being changed. These events will appear with the user who performed the action as the actor, and the chosen device as the resource. Color coding the log by actor can be a useful way to identify these inbound events. Click the actor to change the focus of the Threat Visualizer.

If the action impacts a resource such as a folder, a calendar invite, a user, a file - this will appear next. The resource can also be the user currently focused upon if another user does something to affect them. Color coding the log by resource will also add the resource type in square brackets at the end of the line.

The final interactive part of the log is the source IP. Both the IP and the ASN are listed. Users using mobile devices or private home networks may have natural variation in source IP but a stable ASN. ASN variation is therefore a better gauge of anomalous location than change in source IP for users detected in external platforms (SaaS users).

- Click the link icon to open the external sites summary. For SaaS source locations, the external sites summary focuses on an ASN and user pair. More information is available in External Sites Summary and External Sites Summary for SaaS Activity.

- Click the IP itself to open an additional event log focused on that IP.



FILTERING THE DEVICE EVENT LOG

Filtering the Log Information

When you are familiar with the parts of each log line, the next step is to filter the log to focus on just the information and activity of interest. At the top of the event log several device-related options are available. The options that appear will vary depending on your environment and the makeup of network traffic or user activity returned by the log:

ICON	DESCRIPTION
	This icon allows a text filter to match against fields such as hostname, application protocol and ASN in the log entries. Only one filter of each type can be defined.
Remove filters	This icon will appear if filters are applied and can be used to remove the current filter set.
Choose which type of events to show in the log	The device event log includes many different types of event; the filter limits the log to just specific types of events, or filters out specific event types. The log can be filtered to show or exclude all network connections, new connections, unusual connections, connections blocked by Darktrace RESPOND, changes to device information ("history"), model breaches triggered by the device, and contextual information or activity ("notices"). SaaS user activity is always of the "notice" type.
Choose whether to show internal or external events in the log.	For connectivity, filter the log to show only events to/from internal endpoints, to/from external endpoints, or both.
Toggle incoming/outgoing events	Filter the log to only incoming connections, outgoing connections or both.
Color-code events by their properties	Color-code the event log lines by the specific filter. Doing so will add additional details after the event line. For example, coloring by port will add the port in square brackets (e.g. `[443]`). Default color coding is controlled in each user's account settings. Event logs for SaaS user devices will present different filtering options specific to SaaS activity.
Advanced options	Provides access to advanced options and filters. Refer to the table below for more information.

Advanced Filters and Options

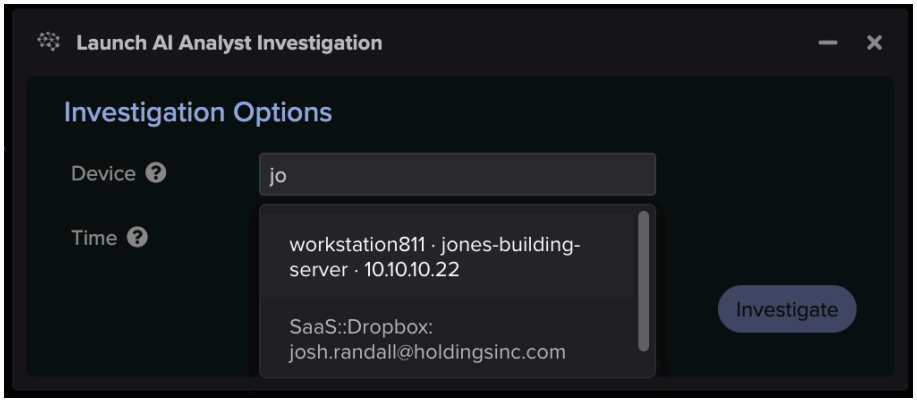
ICON	DESCRIPTION
Unsync from Threat Visualizer time	The device event log displays activity logs going backwards chronologically from the time selected in the top right. If the time is changed, the log will change. Unsyncing the log allows the time to be changed whilst still preserving the logs.
Unsync log from Threat Visualizer filters	The connection panel on the right of the workspace can be filtered to show specific ports or hostnames. These filters are also applied to the event log unless it is unsynced.
Hide duplicate connections	If a device repeatedly connects to an endpoint, this filter restricts the log to show only one entry per-endpoint and port.
Hide event descriptions	Darktrace anomaly detection may add commentary to events that are deemed unusual. This filter hides those comments inline in the event log.
> Any Size	Limit events in the log to those that transferred over a defined threshold of data.
View advanced search	This icon opens the Darktrace Advanced Search interface to explore a more detailed version of the event log for granular analysis. By default, the search is limited to the device IP over a period of 1hr before the Visualizer time selected. More information about Advanced Search can be found in Darktrace Advanced Search. This option is not available for SaaS or endpoint devices.
View packet capture	Opens a dialog to create a packet capture device for connections to/from the device IP over a chosen time window (see Creating Packet Captures). For devices monitored by Darktrace DETECT & RESPOND/Endpoint agents, the interface IP for which the capture should be created must be selected first. This option is not available for SaaS user devices.

TRIGGERED AI ANALYST INVESTIGATIONS

Investigating a User or Device

To trigger an AI Analyst investigation, a **device** or **SaaS user** and a **time** are required. The **device** is the subject of the investigation and the **time** is the starting point for analysis.

In the Threat Visualizer, the **AI Analyst Investigations** item is available from the main menu. The **Launch Investigation** item will open the investigation dialog without a device selected. Here, the field acts as a search bar which can be used to locate the device for investigation.



An AI Analyst investigation can be triggered on a specific device from the omnisearch bar in the Threat Visualizer with the AI Analyst icon. When an investigation is launched from a specific device or SaaS profile, the **Device** or **Account** field is already completed.

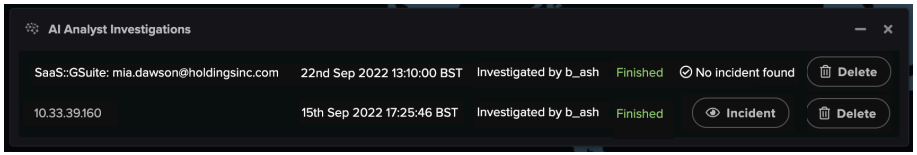
Click **Investigate** to start the investigation.

Investigations

Investigations display the device or user under analysis, the analysis time, the user who triggered the analysis and the current investigation state.

Investigations have three states - pending, processing, finished. A ***pending*** investigation waiting for AI Analyst to start the analysis. A ***processing*** investigation means analysis is underway and a ***finished*** investigation means the investigation has concluded.

Current and completed investigations can be reviewed by navigating to **AI Analyst Investigations** > **View Investigations** in the main menu.



When an investigation is triggered, AI Analyst will perform a close analysis of the activity for the device or user for a period of roughly one hour, using the time provided as a central point. It will then contextualize this behavior against historic activity and connectivity for the entity and its peers. If a finished investigation has the result “No Incident Found”, AI Analyst did not locate any identifiable anomalous activity around the investigation time. If anomalous activity is detected, one or more incident events will be created.

In the Threat Visualizer interface, click the **Incident** button to open the investigation result. Incidents created from triggered investigations follow the same format as AI Analyst incidents created automatically.

For details on triggering AI Analyst investigations from third-party alerts, please see AI Analyst Third Party Investigations or the Darktrace System Administration Guide.

OTHER DEVICE INVESTIGATION WINDOWS

Other Investigation Windows

After consulting the device summary and event log, many users move on to close examination of activity using the Darktrace Advanced Search interface. However, there are additional investigation views for a device that can be useful for gathering further context of normal behavior.

Metric Graph

The graph can be used to plot device behavior over time for a number of metrics and against model breaches the device has triggered. The default metric for network and endpoint devices is “Connections”, for SaaS user devices it is “SaaS All” (all SaaS activity).



Up to two metrics can be displayed on the graph. To add an additional metric, click the ✚ plus icon. For metrics with in/out values such as data transfer and connectivity, click the 👁 eye icon beside the direction to show or hide it.

The time range defaults to center around the time selected in the top right of the Threat Visualizer. The window can be changed using the dropdown or by clicking and dragging across the graph. Clicking on the graph when a time tooltip is shown will change the time in the selector to the clicked time.

If model breaches are shown on the graph, these points are clickable and will open a model breach log for the model breach chosen.

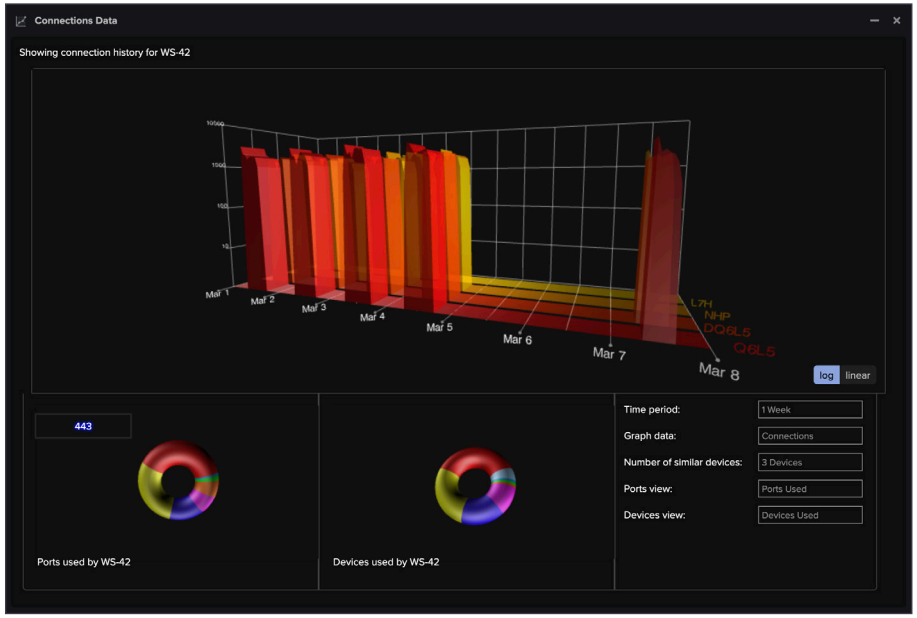
Connections Data Visualization

This option visualizes historic connectivity for the device. It is only available for network and endpoint devices.

The primary graph displays the device behavior over the last week - number of connections or data transfer volume - in the context of peer devices. Device similarity decreases from front to back.

Two additional graphs display port and device breakdown for the connectivity visualized. The main graph can be further filtered by clicking on ports or devices represented on the two donut charts.

The time period, number of devices, and display mode of the graphs can be customized.



CHAPTER 4 - DARKTRACE MODEL BREACHES

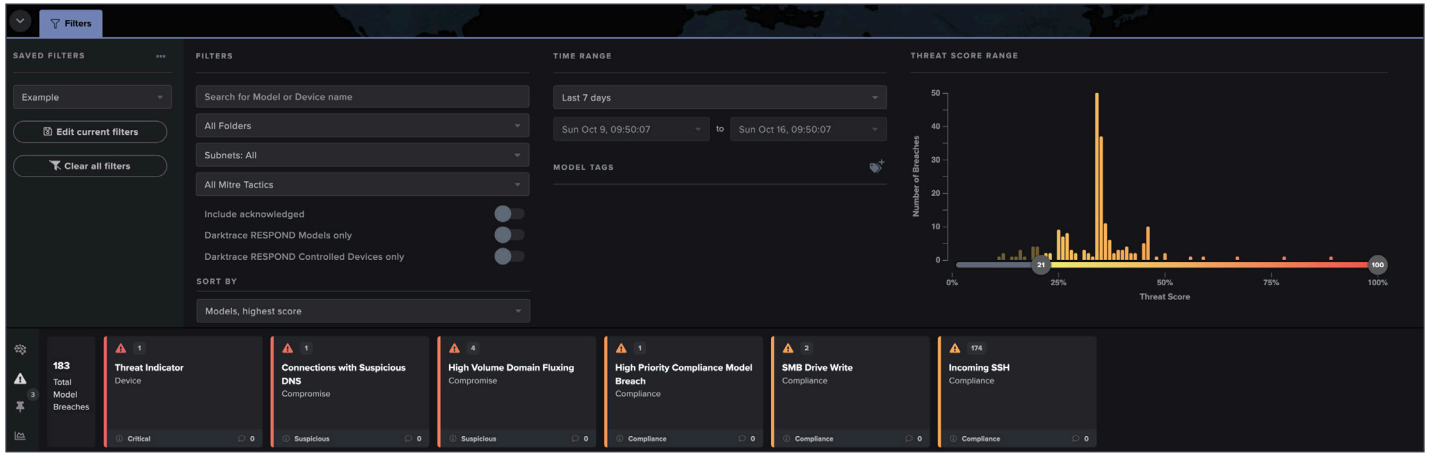
MODEL BREACH ALERTS

DARKTRACE MODEL BREACHES 58
MODEL BREACH ALERTS IN THE THREAT TRAY 59

MODEL BREACH LOGS

OPENING THE MODEL BREACH LOG 60
MODEL BREACH LOG OPTIONS AND FILTERS 61
ENTRIES IN THE MODEL BREACH LOG 62
UNDERSTANDING WHY A MODEL BREACH WAS CREATED 64
UNDERSTANDING EXAMPLE MODEL BREACH LOG ENTRIES 65
ACTIONS FOR MODEL BREACH LOG ENTRIES 66
EXPLORING THE MODEL BREACH EVENT LOG 68

DARKTRACE MODEL BREACHES



What are Model Breaches

Darktrace models are a logic framework that allows operators to interact with the output from 'pattern of life' and anomaly classifiers. Models are primarily used to identify and alert on anomalous and potentially malicious behavior. When conditions for a model are met, this creates a model breach and can trigger an alert or action.

AI Analyst already investigates every model breach that occurs on the system and reports only the most interesting activity as incidents. Users who wish to dive down deeper and investigate specific behaviors after investigating AI Analyst incidents can move to the model breach view of the threat tray.

Threat Tray

When the threat tray is switched to model breach view, model breach alerts appear as tiles. Each tile will list the number of model breaches that have been grouped together.

The total number of model breach alerts in the tray (after filters) is displayed on the left.

Model breaches have a behavior category and an overall score from 0-100%. Behavior categories are high level filters that allow an operator to focus in on specific levels of severity or behavior. There are four categories for model breaches: Critical, Suspicious, Compliance and Informational.

The score can be used to filter on a sliding scale; where a yellow line and icon indicate a low score and red a high score. The score can be a useful way to prioritize within a behavior category - for example, starting with "critical" alerts and investigating from highest to lowest score before moving on to "suspicious".

The information displayed on the tile will change depending on how model breaches are grouped. If model breaches are grouped by device or user, the highest category and highest score of the underlying model breaches for the device will be displayed.

For more information about other filter options available, please see Filtering Alerts with Behavior Categories, Filtering Alerts with Basic Filters, and Filtering Alerts with Advanced Filters.

MODEL BREACH ALERTS IN THE THREAT TRAY

Model breaches can be grouped in a number of different ways - this is controlled by the **Sort By** option in the filter panel. There are three options:

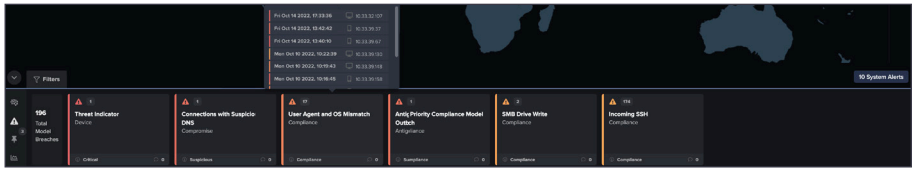
- **Devices** groups model breaches by device.
- **Models** groups model breaches by the model that was breached.
- **Users** groups model breaches by the credential associated with the device (if applicable) at the time the model breached. It also restricts the results returned to just model breaches for devices that had a credential associated within the timeframe selected.

Within each option the results can be ordered alphabetically, by score, time, quantity and number of user comments. When grouped, each tile will show the highest score and most severe behavior category for the underlying alerts.

Individual Tiles

The information displayed on the tile will change depending on how model breaches are grouped. Tiles will always display the highest underlying category, highest underlying score, number of grouped model breaches and the number of comments made by users on underlying model breaches.

- If sorting by model is chosen - indicated by the ⚠ icon - the name of the model that was triggered is shown on the tile.
- If sorting by user (credential) is chosen - indicated by the 👤 user icon - the credential that was associated with the device when it triggered the model breach is displayed.



Hover over a tile to open a tooltip with more information about the model breaches or devices that were grouped:

- If sorting by device or user, the hover will show the names of models that were triggered and the time that the model breach occurred.
- If sorting by model, the names of the network or SaaS user devices that triggered the model breaches are listed with the time that the model breach occurred.

  If the device or user is subject to an active or pending Darktrace RESPOND action at the current time, a Darktrace RESPOND icon - ◉ - will appear beside the device name on the tile. The device name is highlighted green if the action is currently active or orange if currently pending.

Each entry in the tooltip is clickable to open a model breach log focused on only that one occurrence. If the tile itself is clicked, a model breach log with all grouped model breaches is opened instead.

The next step is to understand the model breaches and the log itself.

OPENING THE MODEL BREACH LOG

Model Breach Logs

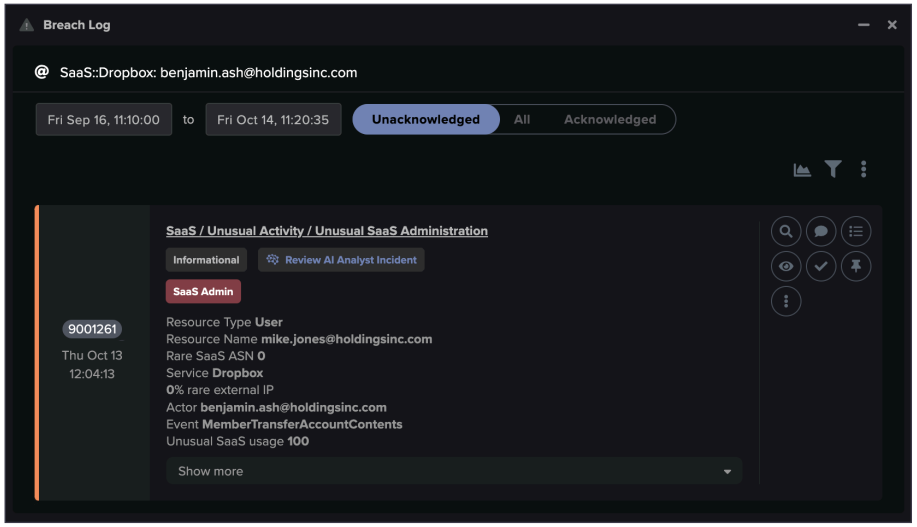
Models are primarily used to identify and alert on anomalous and potentially malicious behavior. Models can also be used to profile a device, assign tags, highlight misconfigurations or trigger autonomous responses from the Darktrace RESPOND framework.

Model breach logs contain a list of model breaches for a model, a device or a subnet over a chosen time window. Each entry in the log is an instance where the criteria for a model were met by a entity in your digital environment.

Reviewing the model breach log for a device can, for example, be valuable in understanding the low-level anomalies that it displayed before a significant incident. Similarly, reviewing the model breach log for a compliance model is a simple way to view all non-compliant devices.

Breach Log

Starting from the threat tray, click on one of the breaches to open the Breach Log. The breach log can be opened from any location - including from a Cyber AI Analyst incident - where the ⚠ icon appears.



In this example, the log has been opened from a threat tray tile for the device "SaaS::Dropbox: benjamin.ash@holdingsinc.com". The log has opened filtered on that specific device. If a model breach log is opened from a tile grouped by device, it will open focused on the device. If alerts are grouped by model, clicking a tile will open a model breach log focused on the individual model.

- A model breach log filtered on a device will display the device name and type in the top left as a heading. Click the device to change the visualizer focus to the device.

- A model breach log filtered on a model display the name of the model as a heading. If enabled in account settings, the model description will appear below. Collapse the description using the ^ arrow icon.

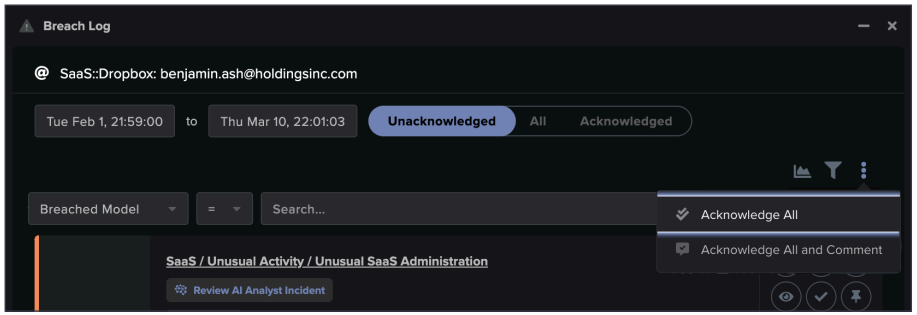
Underneath is the time range for model breaches to be included in the log. The default range is taken from the current threat tray filter but can be altered using the time selectors.

The log can also be filtered to show only unacknowledged model breaches, only acknowledged model breaches, or both states of model breach. By default, acknowledged alerts are removed from the returned results. Most investigation workflows include acknowledging an alert after it has been investigated. For a device, it can be helpful to review anomalous behavior already investigated by other analysts when trying to understand the background to a more severe model breach.

MODEL BREACH LOG OPTIONS AND FILTERS

Log Options

Above the log entries are a series of filters and actions. These options are particularly useful to filter the log when there are many model breach entries returned.



Some options will only appear in a specific mode.

ICON	DESCRIPTION
Add Model Defeats	Defeats are a way to control model behavior without impacting automatic updates from Darktrace analysts. One click defeats allow users with the "Edit Models" permission to quickly exclude specific conditions from triggering a model breach. A detailed description of defeats is covered in the guide to the Model Editor (Model Defeats). This option only appears if the model breach log is filtered to a single model.
View breaches graphically	Clicking this icon will open a graph to the right of the log of model breach score against time. The timeframe will reflect that for the log. The device/model, score and time are displayed in hover.
View model	This icon opens the Model Editor interface focused on the model. The model editor interface is introduced in The Model Editor. This option only appears if the model breach log is filtered to a single model.
Only show discussed	Users can comment on model breaches to provide context or indicate to other users they are working on the alert. Comments remain within the Threat Visualizer interface or any configured alert outputs - to discuss a model breach with Darktrace analysts, please use Ask the Expert or the Darktrace Customer Portal. This option will only be displayed if model breaches in the list have comments.

ICON	DESCRIPTION
Add/remove custom filters	Focus the log on a specific device, model breach ID or metric with custom free text or Regex filters.
Additional actions	Open a dropdown with additional actions.
Acknowledge All	Clicking this icon will acknowledge all model breaches currently in the log. Acknowledging model breaches removes them from display by default - it is a good way to indicate that an alert has been successfully investigated. This option appears under Additional actions.
Unacknowledge All	This icon will only appear when the log is filtered to show "Acknowledged" model breaches. On click, it will unacknowledge all model breaches currently in the log. This option appears under Additional actions.
Acknowledge All and Comment	Clicking this icon will acknowledge all model breaches currently in the log and add a comment to all those acknowledged. This can be helpful if, for example, a custom model developed by a user has over-triggered alerts. This option appears under Additional actions.
Add Comment to All	Add the same comment to all mode breaches currently returned in the breach log.

One Click Defeats

Defeats are a way to control model behavior without affecting automatic updates from Darktrace analysts. A detailed description of defeats and their purposes is covered in the guide to the Model Editor (Model Defeats).

Users with appropriate permissions (**Edit Models**) will see a toggle at the top of the log to "Add model defeats". When enabled, a remove icon will appear beside every field in the model breach entry.

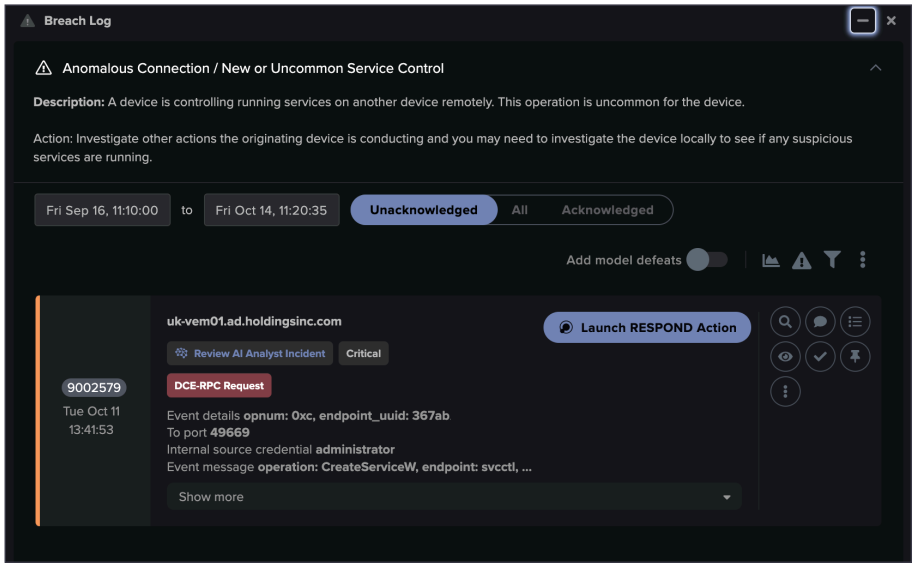
To add a defeat, click the blue icon beside the desired value. A popup will appear describing the change to the model. Confirm to add the defeat.

ENTRIES IN THE MODEL BREACH LOG

Log Entries

Model breach logs contain a list of model breaches for a model, a device or a subnet over a chosen time window. Each entry in the log represents an instance where a device performed activity that met the criteria of a model and triggered a model breach.

In this example, the log has been opened from a threat tray tile for the model “Anomalous Connection / New or Uncommon Service Control”. The log has therefore opened filtered on that specific model.



Entries are sorted chronologically with newest first. The colored bar on the left indicates the severity of the model breach on a gradient from yellow (low) to red (high). In this case, the model has a medium severity,

Next to the bar is the model breach ID - this numeric ID is a unique reference to this model breach in this threat visualizer instance. Click the ID to copy a shareable link to the model breach to the clipboard. In Unified View environments, model breach IDs are prefixed with a number to indicate the source master.

Underneath the model breach ID is the time at which the model breach occurred.

Entry headings

The main part of the log entry includes key information about the device and model breach.

- If the model breach log is filtered on a specific model, each log entry will display the device that triggered the model breach at the top. Hover over the name of the device for a summary tooltip with basic information about the device. Click the device name to change the visualizer focus to the device.
- If the model breach log is filtered on a specific device, each log entry will display the name of the model that was triggered. Click the model name to open the model editor interface for the chosen model in a new tab.
- If the model breach log is filtered on a subnet or a user, both the model name and the device that triggered the model breach are displayed.

If the device or user that triggered the model action is subject to an active or pending Darktrace RESPOND action at the current time, a Darktrace RESPOND icon - - will appear beside the device name on the entry. The device name is highlighted green if the action is currently active or orange if currently pending.

Threat Behavior Category

Behavior categories are high level filters that allow an operator to focus in on specific levels of severity or behavior. There are four categories: Critical, Suspicious, Compliance and Informational. Categories are used across both AI Analyst and Darktrace models.

The model breach category appears under the model or device name. Please see Filtering Alerts with Behavior Categories for more info on categories.

AI Analyst Incidents

AI Analyst investigates every model breach that occurs on the system and reports only the most interesting activity as incidents. If a model breach has triggered an AI Analyst incident, this can be quickly accessed by clicking the "Review AI Analyst Incident" text that appears on the entry.

Darktrace RESPOND Actions

Certain models can also trigger Darktrace RESPOND actions; if a model breach has triggered an action, "Darktrace RESPOND triggered" will appear on the entry.

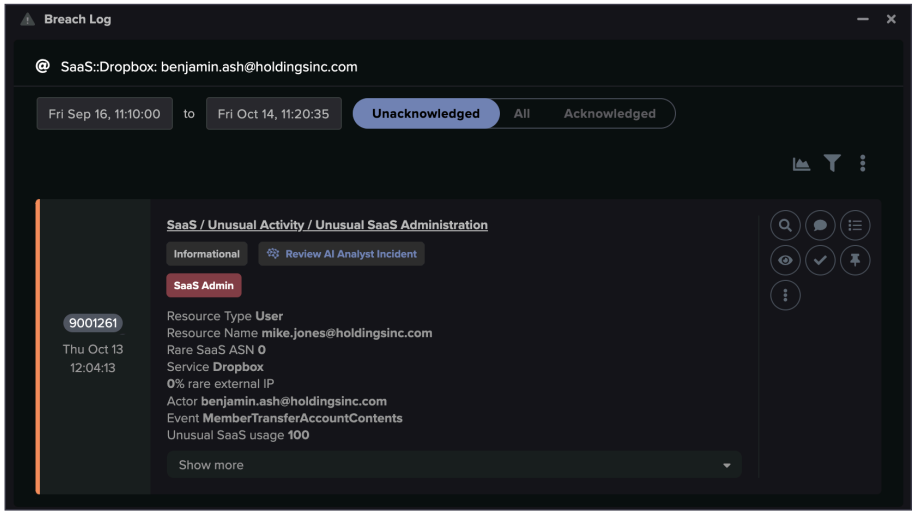
- If this text is highlighted green, the triggered action is currently active.
- If this text is highlighted orange, the triggered action is currently pending - click to access the pending actions and confirm.
- If this text is shaded grey, the triggered action has expired.

Click this text to open the Darktrace RESPOND Actions window filtered to only actions triggered by the chosen model breach.

UNDERSTANDING WHY A MODEL BREACH WAS CREATED

Each log entry gives context as to why the device activity met the model criteria. In bold is the activity ("metric") that the device performed that caused the model breach to be triggered.

For example, this could be a connection, a data transfer over a certain size, a SaaS Admin action, or any activity which is considered new or unusual in the context of the devices' 'pattern of life'. Continuing the example of "Anomalous Connection / New or Uncommon Service Control", the device performed a DCE-RPC Request that met the requirements of the model.



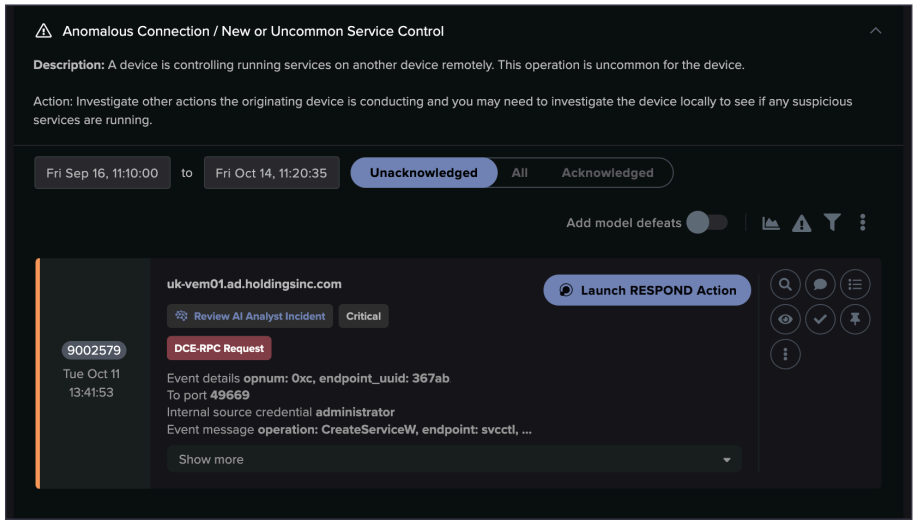
In this example, the SaaS user device performed a SaaS Admin action that met the criteria.

Click the words in bold to open a graph of that metric for the device over time.

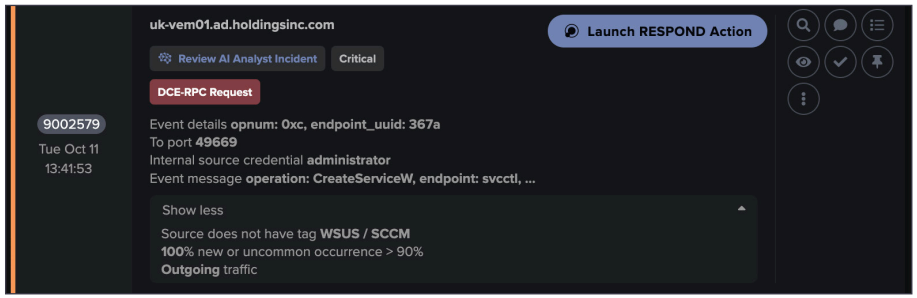
Metrics and Filters

Underneath the key metric is further information about the activity and criteria it met to trigger the model breach. This information is useful context for an analyst to understand the activity and potentially why the activity was anomalous.

For example, it is helpful to know the resource that was changed by the administrative action in the SaaS environment or the DCE-RPC operation that was performed. The information included is chosen when the model is created ("display fields"). Therefore, the amount of information can vary greatly between models.



Not all SaaS Admin or DCE-RPC Requests will have met the model criteria - models have additional restrictions that the activity must meet. For example, the model may be looking for unusual connections from ports other than 443 or SaaS activity from an ASN that is 100% unusual. If there are criteria which were not included by the model creator in the main info section, a "Show More" box will appear. This box contains details of these extra criteria that were met by the activity.

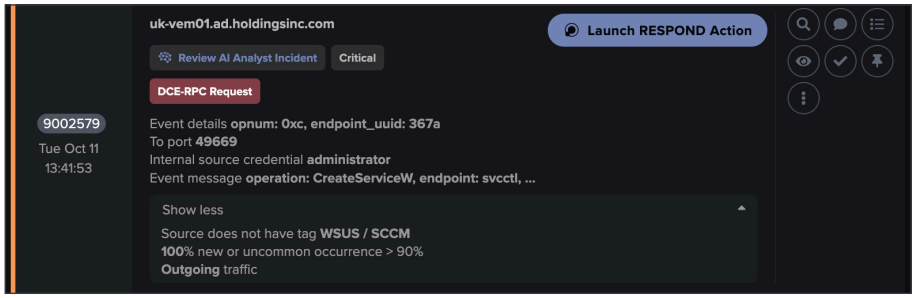


Some models may look for multiple types of activity so these elements may appear more than once on a single log entry.

If workflow integrations that allow lookups for external IPs and hostnames are enabled (such a VirusTotal), a ❗ pivot icon will appear for these entries.

UNDERSTANDING EXAMPLE MODEL BREACH LOG ENTRIES

Worked example - Network

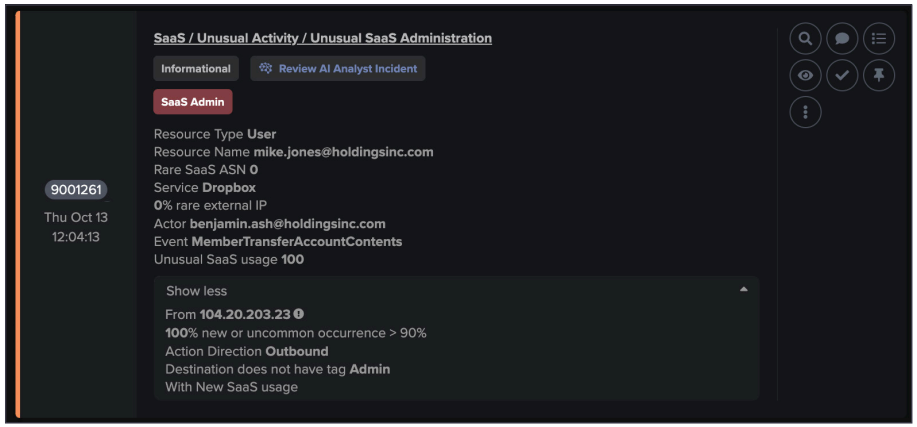


Taking all the sections discussed already together, we can see that the device "vem01.ad.holdingsinc.com" performed a DCE-RPC request containing a "CreateServiceW" function to port 49669 using the credential "administrator". Using the additional model criteria, we can see this was 100% unusual for the device.

This activity triggered the "informational" category model "Anomalous Connection / New or Uncommon Service Control". AI Analyst subsequently investigated and raised an incident for "Suspicious DCE-RPC Activity".

This activity was linked with multiple new and unusual credential usage on the same device, suggesting a wider incident.

Worked example - SaaS



Performing the same process on the SaaS example, the device "SaaS::Dropbox: benjamin.ash@holdingsinc.com" performed a "MemberTransferAccountContents" action on another user "carla.hurst@holdingsinc.com" in Dropbox. Although the ASN and IP for the action are not unusual, it is outside the 'pattern of life' for the user to perform admin actions and this is the first time this user has performed this specific action at all.

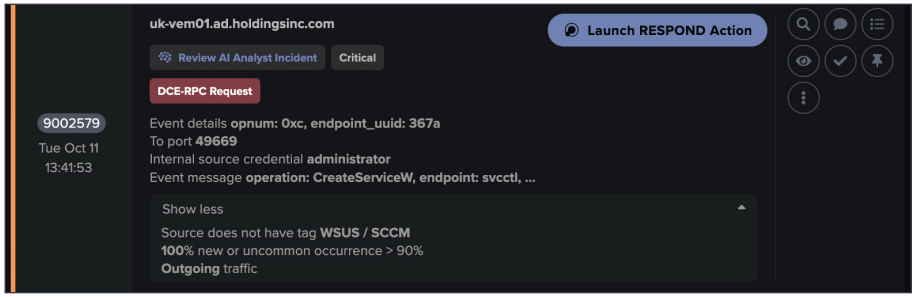
In conclusion, the behavior suggests they may be trying to elevate permissions or access resources they are restricted from and has triggered the "informational" category model "SaaS / Unusual Activity / Unusual SaaS Administration". AI Analyst subsequently investigated and raised an incident for "Possible Hijack of Dropbox Account".

ACTIONS FOR MODEL BREACH LOG ENTRIES

Once the reason for the model breach is understood, a number of actions and further investigation steps can be taken.

- On the body of the entry, the "Launch RESPOND Action" button appears for eligible device or user types. Click to launch a manual action against the device that triggered the model breach.

  See Manual Darktrace RESPOND Actions for more information on triggering these actions.



Each entry also has a series of actions and options on the right:

OPTION	DESCRIPTION
View this model breach in visualizer	This icon changes the focus of the main visualizer workspace to the device that triggered the model breach at the time of the model breach. It is one of the ways to access the device investigation tools introduced in Focusing on a Device.
Discuss this model breach	Users can comment on model breaches to provide context or indicate to other users that they are working on the alert. Comments on a model breach that triggered an AI Analyst incident are also pulled through to the incident discussion. Comments remain within the Threat Visualizer interface or any configured alert outputs - to discuss a model breach with Darktrace analysts, please use Ask the Expert or the Darktrace Customer Portal.

OPTION	DESCRIPTION
View model breach event log	The model breach event log is a filtered event log, similar to the device event log introduced in Understanding the Device Event Log. The log contains only the activity that met the model criteria and caused a model breach to trigger.
Analyze this model breach	One click analysis is a quick way to understand the events that triggered the model breach. Clicking the icon will open an event log with the events that triggered the model breach and a graph displaying relevant metrics that the model itself was looking for.
✔ / ✖ Acknowledge / Unacknowledge this model breach	Acknowledging model breaches removes them from display by default - it is a key part of many investigation workflows. Click this icon to acknowledge a model breach. If the model breach has already been acknowledged, hover the icon to see the user who acknowledged it and when, or click to unacknowledge.
Pin this model breach	Pinning a model breach means it can be retrieved at a later date from the pin icon in the bottom left of the threat tray. Click again to unpin the model breach. Pinned model breaches are specific to each user but will persist between sessions (***requires local storage***).
Filter graph by this model breach	When the model breaches in the log are displayed graphically, this filter allows just model breaches for the chosen model to be displayed. This option will only be shown if the graph is open.
Additional actions	Open a dropdown with additional actions. These advanced options are described below.

Advanced Options

OPTION	DESCRIPTION
⌀ Ignore any future breaches...	This option adds the device to a list, preventing it from ever triggering another model breach for the model. For more information, please see Other Model Tabs for more information.
☁ View this model breach in the SaaS Console	The Darktrace SaaS Console is a purpose built interface for investigating SaaS and cloud user behavior. This icon opens the equivalent model breach in the SaaS console to continue investigation there.
⚠ View Model in Model Editor	This icon opens the Model Editor interface focused on the model. The model editor interface is introduced in The Model Editor.
View MITRE ATT&CK mappings for this model breach	Models maintained by the Darktrace analyst team have been mapped to the MITRE ATT&CK framework where applicable. This mapping is a valuable tool to understand coverage and for teams with internal playbooks for how to address each technique. Click to open a window listing the mapped techniques.
ℹ Explain metrics used in this model breach	The metrics and filters of a model define the activity the model is looking for. When these criteria are met, the model breach is triggered and the criteria are displayed on the log entry in the format “Hostname google.com” or “Source has tag Antigena All”. To understand what these metrics and filters are looking for, click this icon to review definitions written by Darktrace analysts for metrics and filters relevant to the model breach.
Copy to Clipboard	Copies the breach details to the clipboard including the device, the model breach display fields and a link to the model breach.

EXPLORING THE MODEL BREACH EVENT LOG

Investigate a Breach Log Entry

The investigation process for model breaches is very similar to AI Analyst investigations:

1. Understand the activity that caused the alert
2. Focus on the device that triggered the alert
3. Gather context about the way it behaves
4. Deep dive into detailed metadata
5. Action the alert internally if necessary
6. Acknowledge the alert

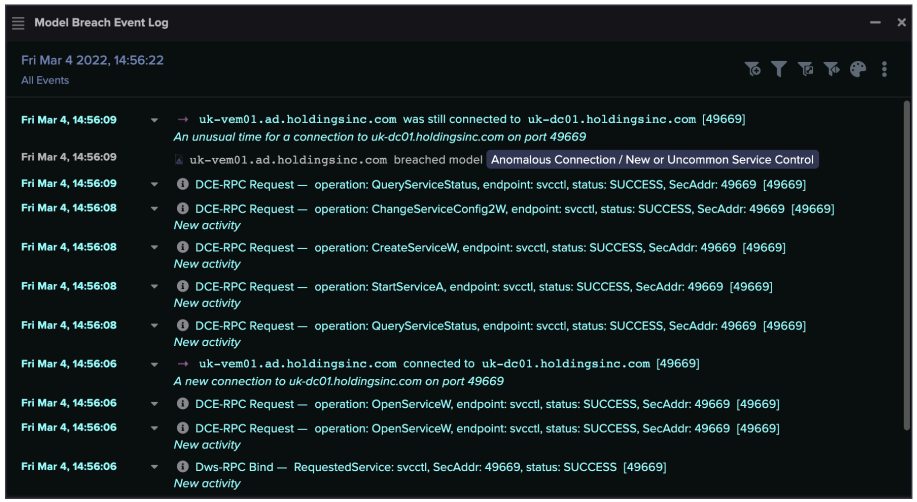
AI Analyst performs a significant amount of these actions for the user and presents the information as part of the AI Analyst incident. For model breaches, these steps must be performed by moving through logs and investigation dialogs in the Threat Visualizer interface. Many of these features such as the device event log and device summary have already been discussed in detail and should be familiar.

To start, open the ☰ model breach event log from the breach log entry.

Model Breach Event Log

The model breach event log is a filtered event log; the log contains only the activity that met the model criteria and caused a model breach to trigger.

For example, for a model breach relating to unusual SSH connectivity, the model breach event log might display a single anomalous SSH connection (or possibly more than one). In contrast, for a model breach relating to excessive HTTP errors or port scanning, the event log will display far more events, because for these types of behavior it is a number of events in combination that all contribute to the detection of an anomaly.



For most models, the event log will contain a small number of entries that are relevant. The following actions are useful to understand this activity better:

- For events or notices, click the ℹ info icon to open a detailed pop-up.
- For connections, click the arrow to see details of the connection. Use the ⬇ packet capture functionality to create a replayable version of the connection at a packet level.
- Use the ▼ arrow icon to pivot to the Advanced Search interface to see detailed metadata about the entry.

The model breach event log is very similar to the device event log and is easy to navigate if the latter is already understood. The filter options used in the model breach event log are the same as the device event log (see Filtering the Device Event Log). The investigation workflow for model breach event logs also mirrors that of device event logs and is covered in detail in Using the Device Event Log, Types of Event in the Device Event Log, Investigating a Connection in the Device Event Log and Investigating SaaS Activity in the Device Event Log.

The next step, as with AI Analyst investigations, is to now focus on one or more key devices that triggered the model breach. Click "🔍 View this model breach in visualizer" from the breach log entry to focus the workspace on the device and investigate its behavior as described in Focusing on a Device.

CHAPTER 5 - OTHER INVESTIGATION TOOLS

EXTERNAL SITES SUMMARY

EXTERNAL SITES SUMMARY 70
EXTERNAL SITES SUMMARY FOR IPS AND HOSTNAMES 71
EXTERNAL SITES SUMMARY FOR SAAS ACTIVITY 72

ALTERNATIVE APPROACHES

DYNAMIC THREAT DASHBOARD 73
ASK THE EXPERT 75
CREATING PACKET CAPTURES 76

APPLYING TAGS

ADDING AND REVIEWING TAGS 77
COMMONLY APPLIED TAGS 78

EXPLORE

EXPLORE MODE 80
FILTERING THE EXPLORE WORKSPACE 82

REPORTING

CYBER AI INSIGHTS REPORTS 84
EXECUTIVE THREAT REPORTS 86

EXTERNAL SITES SUMMARY

External Sites Summary

External sites summary provides statistics on the historic interaction between internal devices and external endpoints. The summary can appear in two formats, depending on how it was launched. There are three tabs - information, rarity and location.

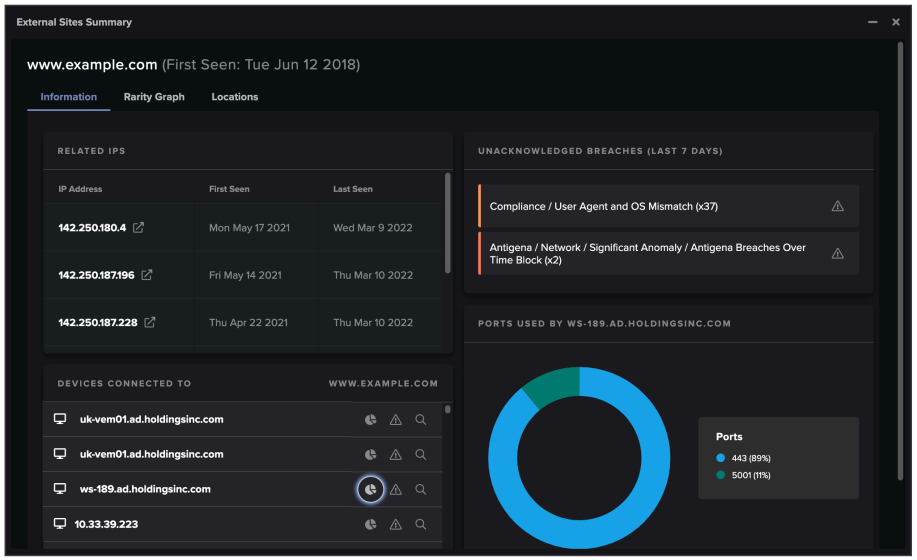
To open the external sites summary, click the globe icon from the omnisearch bar or the external link icon in event logs.

Most frequently, external sites summary is launched from omnisearch or connection event logs. In this scenario, it will provide statistics on internal devices that have connected to the external location and the general rarity of the location in the network.

If external sites summary was instead launched from a SaaS or cloud action, it will visualize the relationship between the ASN / IP and the SaaS user from the event.

Information Tab

The external sites summary will open on the information tab to display statistics on the external endpoint and interaction with internal devices. The contents of the information tab differs when the summary is opened from connectivity information or from SaaS activity.



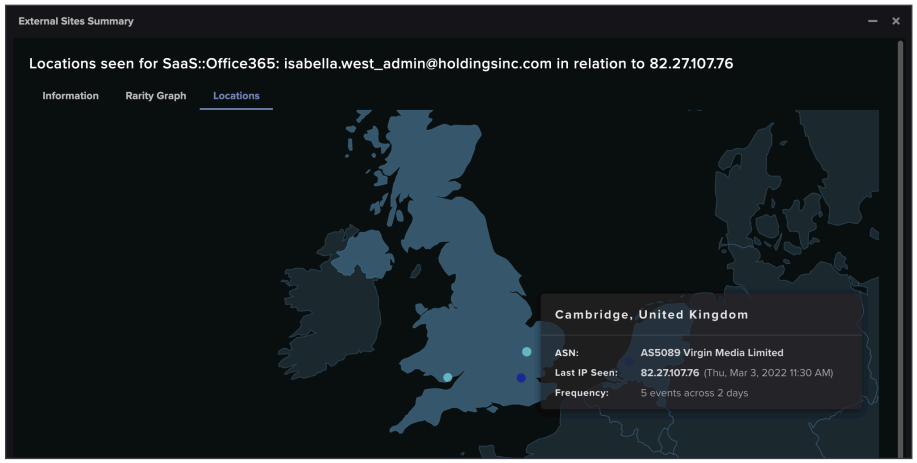
The timeframe for alerts and lists of devices or users is taken from the current time range of the threat tray.

Rarity Tab

The rarity graph tab reports how common the external endpoint is in the context of network behavior or global SaaS activity over time. Endpoints on the “Trusted Domains” list will always be reported as 0% rare.

Location Tab

For IPs and hostnames, the locations tab displays the approximate location of the external host on a map using ASN geolocation data. For SaaS activity, the locations tab instead displays the approximate geolocation for all IP addresses that have performed SaaS activities for this user. In both cases, the ASN for the IP under investigation will pulse to highlight its location.

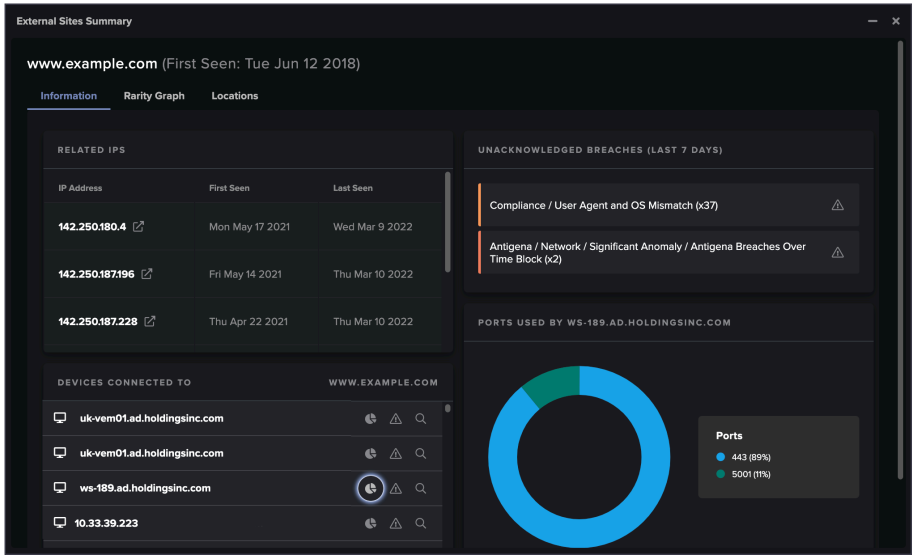


Click on any location marker to see the associated IP addresses, the ASN and - for SaaS activity - the frequency at which the activities were performed. Both the ASN or IP in the tooltip can be clicked to launch event logs.

Locations for SaaS user activity are also colored in this mode to display rarity: light blue locations are observed frequently for the user, dark blue locations rarely.

EXTERNAL SITES SUMMARY FOR IPS AND HOSTNAMES

When the summary is focused on network connections to and from the external location, the hostname or IP appears as heading in the top left. The first time the location was seen in connections observed by Darktrace is also displayed. This will persist between tabs.



If the summary is focused on an external IPs, a section will appear with related hostnames that have been associated with the IP in DNS queries observed by Darktrace. If the summary is focused on a hostname, the IPs seen in DNS will instead be displayed. New external sites summary windows can be opened for these related IPs and hostnames using the external link icon.

On the right is a list of unacknowledged, potentially uninvestigated, alerts where the external location was involved in the activity that triggered the model breach. Click the triangle icon to open the model breach log.

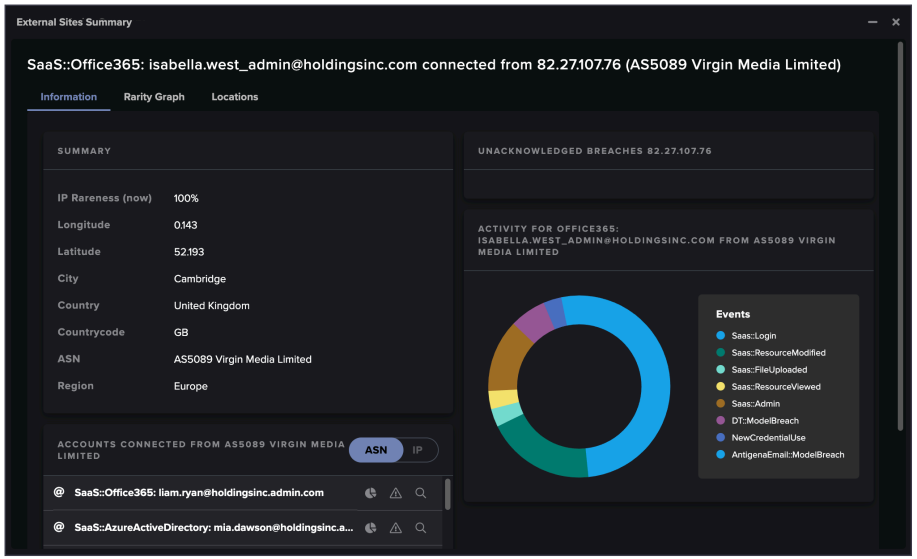
Underneath the related IPs or hostnames is a list of internal devices that have been observed connecting to the external location. The panel on the right will not populate until a device has been chosen using the chart icon.

ICON	DESCRIPTION
View graph	This icon populates the panel to the right with a chart breaking down the ports that the device used to contact the external location.
Open device model breach log	Opens a model breach log focused on the device. Model breach logs are covered in more detail in Exploring the Model Breach Event Log.
View device in visualizer	Switches the focus of the main workspace to this device. Focusing on a device is covered in more detail across a series of articles, introduced in Focusing on a Device.

EXTERNAL SITES SUMMARY FOR SAAS ACTIVITY

SaaS

When the summary is opened from SaaS or Cloud activity, the information displayed is focused on the combination of user and external location. This user and IP/ASN pair appears as heading in the top left.



The summary section on the left contains information about the external IP that was used to perform SaaS activity, including approximate geographic location and rarity in the context of the network.

On the right is a list of unacknowledged, potentially uninvestigated, alerts where the external location was involved in the activity that triggered the model breach. Click the ⚠ triangle icon to open the model breach log.

The bottom right displays a summary of activity the user has performed from the same ASN over the last month in chart-format. This activity is grouped into high-level categories for clarity.

On the left are other users who have connected to a SaaS service using the same ASN or IP. The activity chart can be switched to one of these other users by clicking the ◔ chart icon. Changing the filter between ASN and IP will also change the chart.

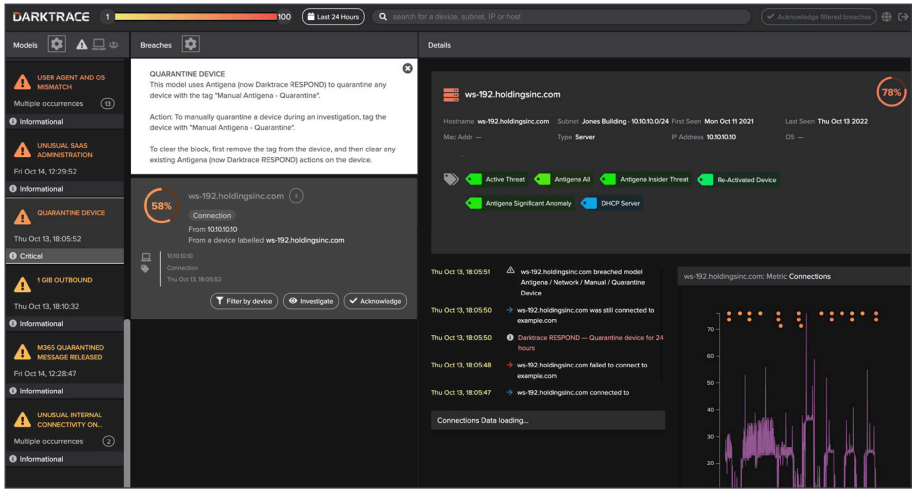
ICON	DESCRIPTION
◔ View graph	This icon populates the panel to the right with a chart breaking down the activities the user performed when they connected to a SaaS service from the IP or ASN selected.
⚠ Open device model breach log	Opens a model breach log focused on the user device. Model breach logs are covered in more detail in Exploring the Model Breach Event Log.
🔍 View device in visualizer	Switches the focus of the main workspace to this device. Focusing on a device is covered in more detail across a series of articles, introduced in Focusing on a Device.

DYNAMIC THREAT DASHBOARD

Accessing the Dashboard

The Dynamic Threat Dashboard view dashboard can be accessed from the main menu, or navigated to directly from the URL. It provides a full screen interface for viewing model breaches and key information for each breach. This screen is intended for extremely rapid triage with a minimum of interaction. Note that it will scroll through incidents if left unattended which can be useful on SOC TV display screens.

This page allows for the quick sorting, viewing, and acknowledging of breaches for triage purposes. If a breach requires additional investigation, the user can pivot back to the 3D View to view further information. Selections flow from left to right across the dashboard, and keyboard arrow keys can be used for navigation.



The contents of the dashboard can be globally filtered by the search bar on the top right.

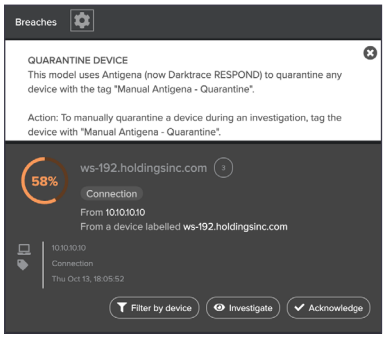
Search Bar

The search bar allows for powerful filtering by multiple search terms.



- Free text searches are converted to filter terms upon pressing return.
- Typing additional text and pressing return constructs a second filter term. This process can be continued to add additional filters to narrow results down. Each subsequent filter is AND-ed to the previous ones.
- Filters can be removed by clicking the “x” in each “pill” in the search bar.

Left-hand Panel



This panel provides the top-level navigation for breach information.

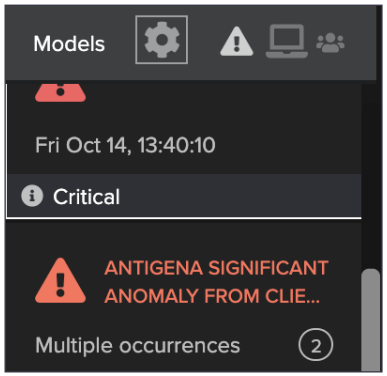
The selector icons in the title bar switches between model, device and user grouping. Selecting an entry populates the breaches panel with the breaches for that model, device or user.

In **model** grouping the panel lists the models that have breached during the specified time frame, sorted by total score for that model. Each entry includes either the time of breach (for models with only one breach) and the count of the number of breaches for the model.

In **device** grouping, the panel lists the devices that have breached models during the specified time frame, sorted by total score for that device. Each entry includes either the time of the breach (for devices with only one breach) and the count of the number of breaches for the device.

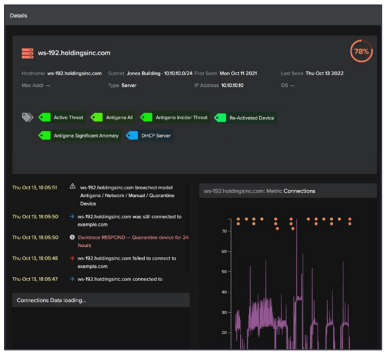
In **users** grouping, the panel lists user credentials that were associated with model breaches within the specified time frame. Not all model breaches will have an associated user credential; this panel will not encompass all model breaches, only those where a credential can be associated. Each entry includes either the time of the breach (for users with only one breach) and the count of the number of breaches for the user.

Breaches Panel



The panel lists individual breaches for the selected device or model. The entries are sorted by breach score, displayed in the radial icons. The entries will be labelled with either the name of the model that breached (in device view), or the name of the device that breached the model (in model view), and the number of other breaches for that model or device. Each breach entry has buttons to acknowledge or filter.

Clicking the filter button will populate the search bar with a filter term to match the other breaches for that model or device. All model breaches returned by the filter can be acknowledged using the "acknowledge filtered breaches" button in the top right. Automatically produced filters can be combined like any other filter within the search bar. The filter can be removed by clicking the x beside the filter term in the search bar.



Details Panel

The right-hand side panel displays summary information for the selected breach. This information is designed to allow for quick decision making about the events in question.

The top section displays overall threat score and information about the device in question. The sections below contain dynamically populated content showing connections and visualizations that contributed to the breach.

The most relevant information will be automatically displayed.

If the displayed information is inconclusive, clicking the "Investigate" button will navigate to the 3D View with the device, time, and breach event log for the current breach selected.

ASK THE EXPERT

What is Ask the Expert

Ask the Expert allows for the creation of incidents which can be submitted to Darktrace analysts for feedback. Creation of an incident is done within the Threat Visualizer by entering text and dragging relevant data into the creation window. Submitted incidents become open questions which can be viewed, and commented with updates both by Darktrace analysts and local user accounts. Questions have associated statuses and function as "tickets" that can be closed when resolved.

Ask the Expert is available from the Main Menu, under Help – the **"Ask the Expert"** menu item allows for the creation of new questions, and the **"View Questions"** menu item allows for the viewing of existing questions and their statuses.

Viewing Questions

The **"View Questions"** pop up lists all previously submitted questions with their timestamp, owner, and status. Possible question statuses are:

- New
- Responded
- Closed

The toolbar in the popup allows for filtering and highlighting of the incidents by their attributes, as well as searching by author or name.

Clicking on a question opens a new popup that allows the viewing of Darktrace responses and the addition of new comments. Questions can be closed or deleted – closing will mark the status of the question as closed and resolved, without actually removing from the question list.

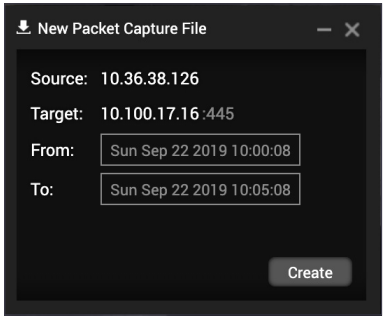
CREATING PACKET CAPTURES

Packet Captures

Packet Captures can be generated from connections seen in the Darktrace Threat Visualizer for further investigation. There are two ways of viewing raw packets (in PCAP format): firstly, by creating the PCAP from the dropdown next to any connection in the device event log, secondly by creating it from the device event log when a device is selected.

If you have a distributed deployment, raw traffic is stored on the probes, which can cause a time delay in retrieving PCAPs. Any PCAPs created are then saved on the master appliance.

Be aware that PCAPs are truncated due to storage requirements so you may not see the entirety of the data flow between the devices in question especially where large amounts of data (e.g., files) are transferred. However, the PCAP typically gives you enough information to get an idea of what is happening.



Creating Packet Captures

For advanced users who wish to access raw packets relating to device communications there are options within the device event log for creating a packet capture file.

***Tip**: If creating a packet capture file for a single device from the device event log, the Source and Target fields are not shown. The field shown instead is 'IP' (referring to the IP address of the device) – no destination is specified.*

Source: This shows the source IP address pre-populated with the value from the event which has been selected.

Target: This shows the target IP address pre-populated with the value from the event which has been selected.

From: The start time of the packet capture. This is pre-populated from the event which has been selected. The start time can be altered by clicking on the gray 'From' box.

To: The end time of the packet capture. The end time can be altered by clicking on the gray 'To' box. The time range defaults to Darktrace's best guess as to which range will capture the connection of interest (for example if it knows that the connection ended at a certain time).

Create: This option starts the packet capture.

Cancel: This option closes the window.

Use Connection: This option is only available when creating a packet capture for an event. Clicking this will further filter the selected packets by the connection. The time fields default to the start and end second of the connection selected and the source and destination ports are as specified in the connection.

Viewing the PCAP File

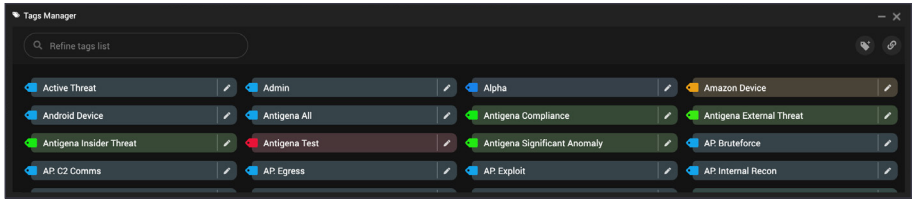
The file can be viewed by clicking the PCAP option from the main menu. PCAPs can be viewed in the Threat Visualizer interface or downloaded for use in external tools.

ADDING AND REVIEWING TAGS

Tags offer a simple labelling system for network devices, credentials and users detected in external platforms (SaaS users) which can be used to identify important resources, control eligibility for Darktrace RESPOND and alter model logic. Tags can be automatically applied as a model action, used in model defeats to exclude devices, or factored into model logic. For users detected in external platforms (SaaS users), tags may be also be applied automatically by the module to identify the roles assigned to the user - these tags are identified by "(CG)".

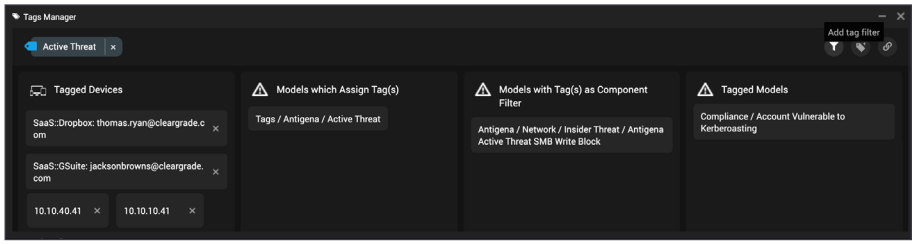
A small number of tags are restricted for system use only and are used to indicate devices excluded from monitoring or behaving erratically. See Commonly Applied Tags for an explanation of commonly applied tags and their purpose.

Viewing and Editing Tags



Tags can be managed in the tags manager, accessible from both the SaaS Console or the Threat Visualizer, or via the API (/tags). Select **Tags** from the **Main Menu** to open the manager.

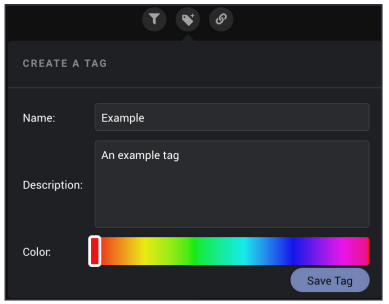
A list displaying all current tags will appear, optionally filterable with the **search**. In the top right, the link icon pivots to Explore Mode in Explore Tags mode.



To review a tag, click on it. A summary page will specify (where applicable) the devices tagged, models that apply the tag, models dependent on the tag for breach logic and models that have the tag applied to them. All models and devices or users can be clicked. To further refine the list, the filter icon can be used to add an additional tag, filtering only by models/devices that share both tags.

To edit or delete a tag, click the edit icon beside the tag name.

Add a New Tag



To add a new tag, click **Add Tag** from the Tags Manager.

When creating a new tag, a description can be added to help identify the purpose of the tag. Selecting a color assists identifying the tag when assigned to a device. Click **Save**.

The new tag will appear on the tag manager.

Tags can also be created via the API.

Tag a Device

In the Threat Visualizer, tags can be added when a device is focused upon from the global omnisearch or from an event log. The tags associated with the device are listed below the search.

Click the plus icon to open a dialog to add an additional tag. In the pop-up window, locate the tag and click on it to add to the device.

When the device has many tags, these are collapsed and can be shown in a popup, searchable dialog by clicking the "+[#] tags"

Scoped Tags

Scoping is a special syntax that can be used to create groups of tags (or *"scopes"*). Scoped tags are mutually exclusive; only one tag of a specific scope can be present on a device at once. Attempting to add another tag from the same scope will replace the existing tag.

For example, a device has the tag `Team::Finance` (which has the scope `Team`). If a user or model attempts to add the tag `Team::Marketing` to the device, doing so will *remove* `Team::Finance`. This is because the two tags have the same scope - `Team` - and only one tag from the `Team` scope can be present on the device at once.

This can be useful when tags are used to indicate a status or membership of a group where it would not make sense to have more than one entry. For example, a set of tags could be created to describe the employment status of users: `EmployeeStatus::AwaitingStart`, `EmployeeStatus::Current`, `EmployeeStatus::Leaver` & `EmployeeStatus::Left`. A user can move between these states but would never occupy more than one state.

To create a scoped tag, use the syntax `[a]::[b]` when creating the label for the tag, where `a` is the scope and `b` can be any text.

A double colon - `::` - must be used as part of the tag label or the tag will not be scoped.

COMMONLY APPLIED TAGS

The Threat Visualizer supports a flexible label system called Tags to allow analysts to be able to rapidly label and refer to groups of devices within the platform. One use case for this feature is to enable monitoring of high-risk users or devices such as departing employees or key assets. Tags can be applied manually, applied automatically by models as a breach action or applied as part of the Darktrace mathematical modeling.

The following tags are commonly used within the Darktrace platform:

TAG	DESCRIPTION
Admin	Device will be excluded from common admin related models (for example, unusual admin session models).
Security Device	Device is a known security device which can affect the behavior of other devices by making thousands of inbound connections - excludes it from related models for such activity (for example, network scanning models).
Exclude Data Transfer	Device should not breach on unusual data transfer models (for example, backup servers).
Model Excluded	Device will not breach any models, but its connections are still recorded and visible as normal.
High Risk	Device has breached a network scan or attack tools model in the past 48 hours.
Active Threat	Device has breached a Darktrace RESPOND/Network (Antigena Network) model (excludes Compliance models) with a breach score greater than 25% in the past two hours.
Key User	Device has authenticated using a key user account (requires configuration of key users in the corresponding model).
Conflicting User Agent	Device has been detected using at least three different user agents in individual connections within the last 4 minutes (for example, Linux, Windows and iOS).
DHCP Relay	Device has been identified to relay local DHCP requests to a DHCP server.

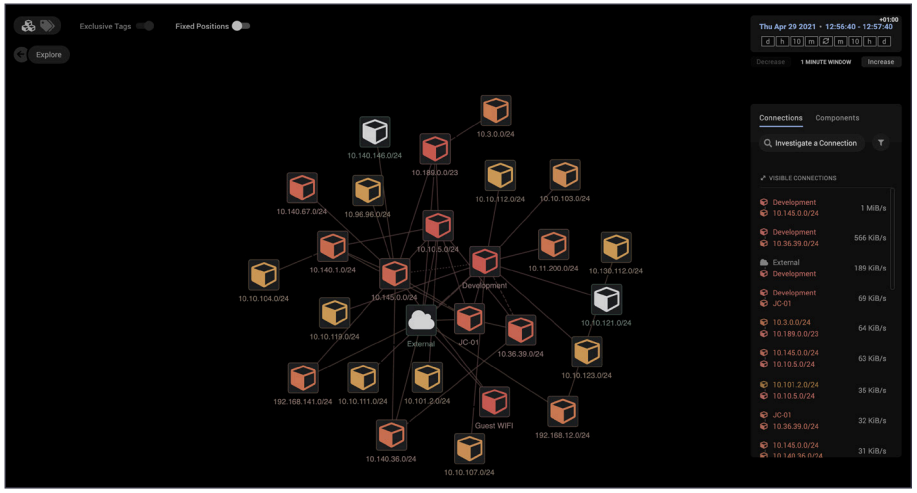
TAG	DESCRIPTION
Unique Clients	This tag is applied to client devices which are unable to be assigned to a cluster by mathematical modelling due to having a unique fingerprint of network behavior. It is allocated by Darktrace's machine learning and is part in of the clustering of devices into 'peer groups' or 'similar devices'. This tag is applied by the system and not controlled by models.
Unique Servers	This tag is applied to server devices which are unable to be assigned to a cluster by mathematical modelling due to having a unique fingerprint of network behavior. It is allocated by Darktrace's machine learning and is part in of the clustering of devices into 'peer groups' or 'similar devices'. This tag is applied by the system and not controlled by models.
Unusual Connectivity Excluded	A device has performed such an anomalous volume of connections that it is no longer tracked for unusual connectivity, as its pattern is so erratic / high volume.
Domain Authenticated	Device has authenticated with a domain controller within the last 0.5 days.
Internet Facing System	Device has received incoming external connections but is not tagged as a server (i.e. internet exposed device).
Re-Activated Device	A device was inactive for at least the past 4 weeks and has re-appeared on the network.
External DNS	Device has performed an external DNS connection.
No Device Tracking	Devices within IP ranges that have no tracking will be tagged (requires configuration of the corresponding model). Devices with this tag will be excluded on unusual activity breaches.
DNS Server	Devices performing the functionality of DNS servers - will be excluded from models which look for unusual DNS activities.
Mail Server	Device was seen sending or receiving lots of email. Devices with this tag will be excluded from models which look for unusual volumes of emails sent.

TAG	DESCRIPTION
WSUS / SCCM	A device is operating as a Windows Server Update Service (WSUS) / System Centre Configuration Manager (SCCM). A device with this tag is excluded from breaching on some common activity patterns such as unusual internal data transfers to new devices. This tag should be added automatically by the WSUS or SCCM Tag model but in some circumstances, it may not be possible to automatically detect the service, and the tag can be added manually.
Inoculated Device	Devices that trigger an inoculation breach receive an Inoculation tag for 1 week.

EXPLORE MODE

Opening Explore

The Explore feature can be accessed from the Main Menu item **Explore**, where the user may select from **Explore Tags** or **Explore Subnets**. Alternating between Tags and Subnets is also easily performed within Explore itself. The pivot icon is also available in the main Threat Visualizer when investigating a specific Subnet, or from the Tag Manager.



On the **Explore Tags** and **Explore Subnets** pages respectively, tags and subnets are clearly labelled and color coded to enhance the user experience. Zooming in and out is possible in order to improve visibility.

On the right-hand side of both the **Explore Tags** and **Explore Subnets** pages, there is a panel with a list of **Connections** and **Components**, which can be accessed by clicking on the respective title.

Connections

The **Connections** subsection gives a list of the visible connections together with information about the average rate of transfer in MiB/s or KiB/s. These connections can be further explored by clicking on them.

Components

The **Components** subsection gives a list of the visible Subnets and Tags and further information about the average rate of transfer in MiB/s or KiB/s. These components can be further explored by clicking on them.

Refer to **Filtering the Explore Workspace** section for further information about exploring **Connections** and **Components**.

The Explore function utilizes the same time selector as the main Threat Visualizer to adjust the point in time represented, and the **increase window** feature allows for expansion of the time window itself. Explore presents a cross-section of the network at a specific point in time, therefore it does not display real-time connection information.

Using Explore

Communication between subnets or tags is indicated by a dashed line, where dash length correlates with the amount of data transferred, i.e. **the shorter the length and more numerous the dashes present** in a connection, **the more data transferred** and respectively, **the longer the length and less numerous the dashes present** in a connection, **the less data transferred**.

LINE	KEY
——— ——— ———	Less data transfer.
— — — — — —	
----------	More data transfer.

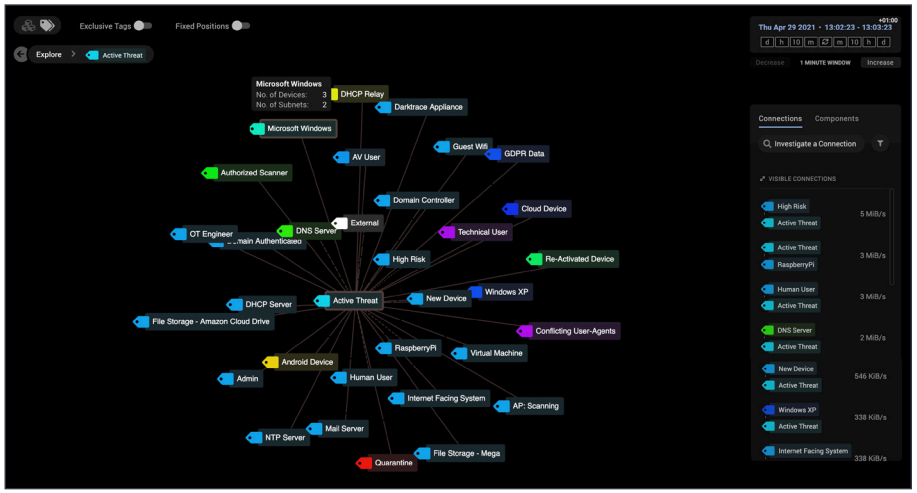
represented visually as concurrent connections in both directions may be observed at the same moment between numerous devices in the particular Subnet or represented by the specific Tag.

The severity of any alerts generated by devices in a Subnet during the time frame is indicated by the color of the icon, where white indicates no breaches and a gradient from light yellow to red indicating severity. The Subnet color is determined based on the highest scoring device of all breaches for all devices in that Subnet taking account for the last seven days.

Tags can be customized in the main Threat Visualizer Tag Manager.

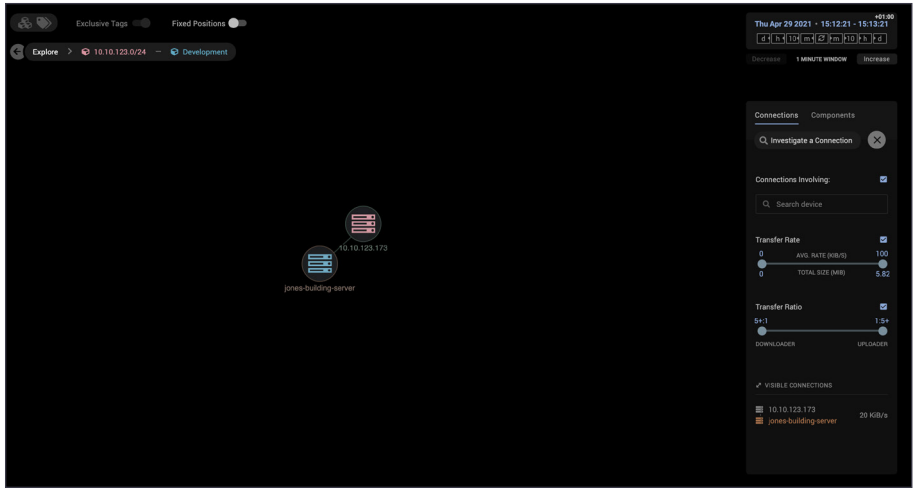
Investigating A Device

- When in **Subnets** or **Tags View**, hovering over a subnet or tag will display key information.
- When in **Subnets View**, hovering over a subnet will display the subnet's name, the number of devices it encompasses and the number of tags applied to devices within it.
- When in **Tags View**, hovering over a tag will display the tag's name, the number of devices within that tag and the number of subnets that devices with the tag are part of.



From either **Subnets** or **Tags View**, clicking on a particular line of connection (i.e. clicking on a visual dash connecting subnets or tags) will reveal much more detailed information about the chosen connection. While in this **Device View** mode, hovering over any of the icons representing various components provides relevant information such as name, vendor, mac address, operating system, type, subnets and tags.

From this **Device view**, clicking further on a specific line of connection between two devices will open the **Event Log** on the right hand side of the page. The log for connectivity connection between the two devices is represented within the specified timeframe. The log includes time, date and connection direction. By hovering over the arrow, further information is available, including start time, source, destination, and protocol.



To return to previous levels of connection, click on one of the available options that appear in the top left hand side corner of the screen after **Explore >**. This submenu represents the route taken to reach the currently observed **Event Log**.

FILTERING THE EXPLORE WORKSPACE

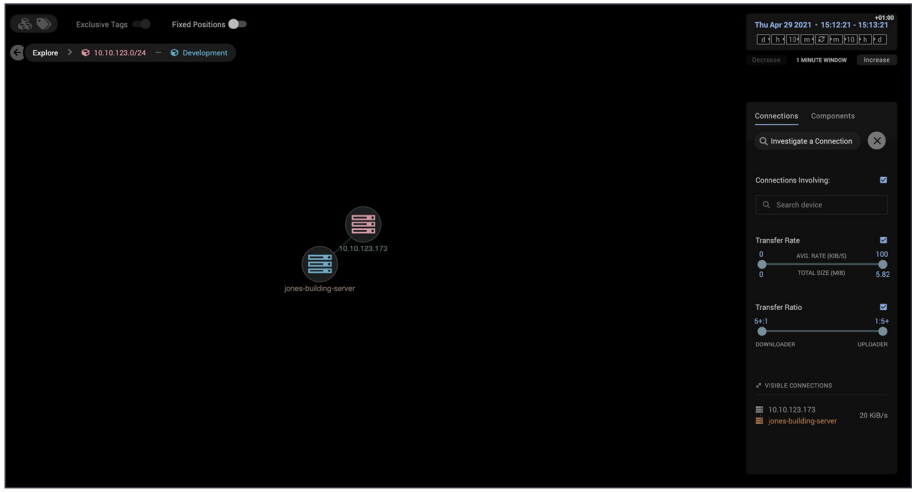
Filtering the Explore page

When in **Subnets** or **Tags view**, on the right hand side of the screen where the **Connections** and **Components** can be explored, clicking on the filter icon next to **Investigate a Connection** will open a filter pane. Filtering will focus the workspace by removing connections that do not match the given conditions. The available filters:

Connections Involving

If the **Connections Involving** box is ticked, the "Search device" field underneath it will allow free text input to search for a device. Once text is inserted, this will remove **Subnets** or **Tags** from the workspace that do not contain the search text nor are in connection with the devices that contain the search text. Elements which are in connection with the devices that contain the search text but do not match the search will appear slightly faded.

Filtering with a search term is a helpful function in narrowing down a specific device or devices and clearly observing their connections.



Average rate (KiB/s)/ Total size (MiB or GiB)

If the **Transfer Rate** box is ticked, the option to search for connections based on the average rate of transfer in Kibibyte per second **(KiB/s)** is available. By using the slider, the workspace can be customized based on the transfer rate. The **Subnets, Tags** or **Devices** (depending on where this search is being performed) that are applicable will stay in the workspace. Items which do not match the criteria will disappear from view.

In the same way, by using the same slider, the workspace can be customized based on total size of data transferred, in **MiB** (Mebibytes) or **GiB** (Gibibytes), which is displayed underneath the sliding pointer.

Transfer Ratio

If the **Transfer Ration** box is ticked, the option to search for connections based on the transfer ratio is available. The transfer ratio filters **Subnets** or **Tags** on the workspace by the ratio of data uploaded versus that of data downloaded during the time frame. The left slider controls the downloaded amount and the right slider controls the uploaded amount. At the highest level, this controls the total amount of uploaded:downloaded between the two subnets. At a lower level, this ratio is on a device to device level.

Where **Subnets** or **Tags** that fall within the slider bounds have connected to/been connected to other elements that do not fulfil the filter conditions, these elements will remain but will be greyed out. Other elements which do not meet the criteria and have not connected to these nodes will disappear.

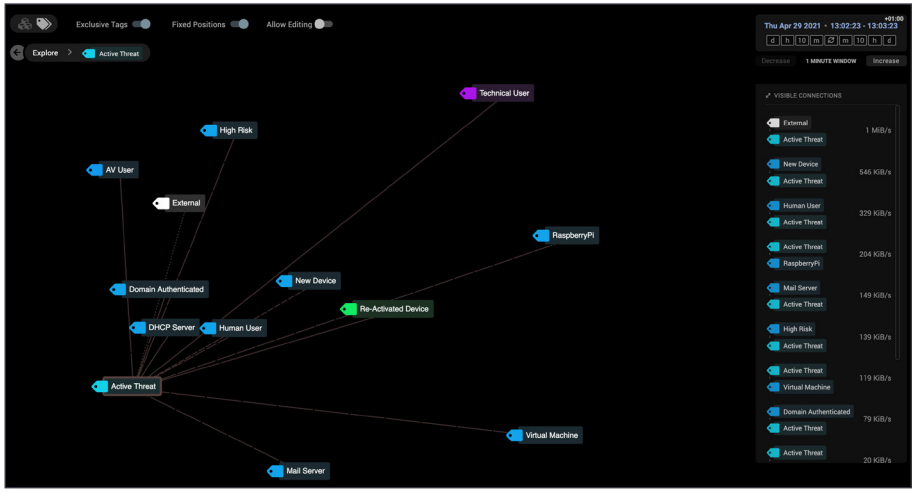
Explore Layout

The distribution of **Subnets** and **Tags** in the **Explore screen** dynamically updates to position connection events in the clearest format; initial locations are derived only from the communications observed during the time window. At the top left of the workspace are view toggles.

Fixed Positions

The **Fixed Positions** toggle locks subnet and tag nodes in place, preserving their location on the workspace when deep-diving into particular connection events and returning to the main overview.

Once the **Fixed Positions** toggle is enabled, another toggle appears next to it - **Allow Editing**. If the toggle **Allow Editing** is enabled, the nodes can be individually dragged and dropped to alter their persistent position. This function allows nodes to be places according to the user's understanding of the network layout. Each user maintains their own local layout.



Exclusive Tags

To use **Exclusive Tags**, activate the relevant toggle on the main workspace. When enabled, the **Exclusive Tags** toggle restricts the connections visualized between tags. If a device has more than one tag in common with the device it is connecting with, when tags are exclusive, Explore will not render a connection between the tags that are shared between the devices.

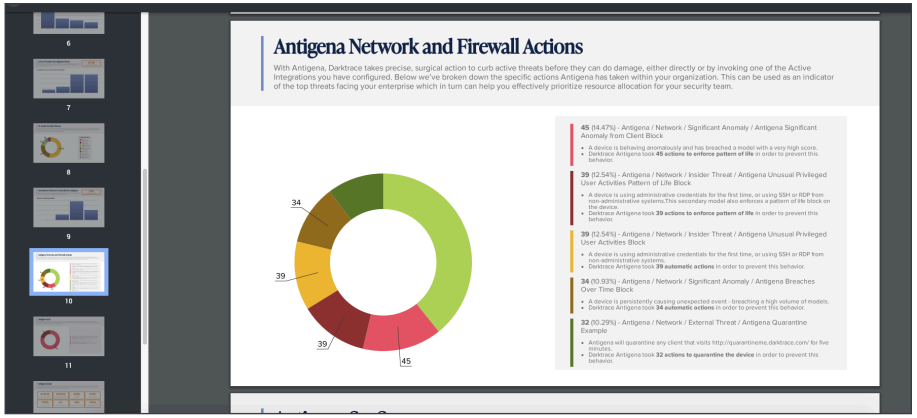
For example, if two devices (a & b) are tagged with ***iOS device*** and ***iPhone*** and Exclusive Tags was disabled, the behavior of Explore would draw a connection between ***iOS device*** and ***iPhone*** and ***iPhone*** to ***iOS device*** to represent this communication. With Exclusive Tags, no connection line would be drawn between these shared tags. If device b is also tagged with ***New Device***, exclusive tags would render the connection between ***iOS Device*** and ***New Device*** and the connection between ***iPhone*** and ***New Device***, but not the connection between the shared tags.

CYBER AI INSIGHTS REPORTS

Cyber AI Insights reports can be generated within the Threat Visualizer and present an overview of coverage, user engagement and value provided by your Darktrace deployment. The report displays currently deployed components, integrations, probes and modules against those generally available.

The data contained in each report is specific to your organization and includes trends in processed network traffic, the categories of threat that triggered Darktrace RESPOND response and a breakdown of events detected by AI Analyst by threat stage. To measure user engagement, the average time-to-acknowledge for human operators is calculated across the timeframe.

Reports can cover a single day or up to 12 months (data-dependent) and can be customized to include estimate costs savings from AI-augmented incident response and triage. Where applicable, reports will also include data on threat trends and response from Darktrace/Email. Please note, this data is restricted to the last 4 weeks, regardless of overall report timeframe.



Cyber AI Insights reports can also be automatically generated and sent to recipients by email; this is configured on the System Config page. For more information on scheduling reports, please refer to Scheduling Automatic Report Generation.

Report Options

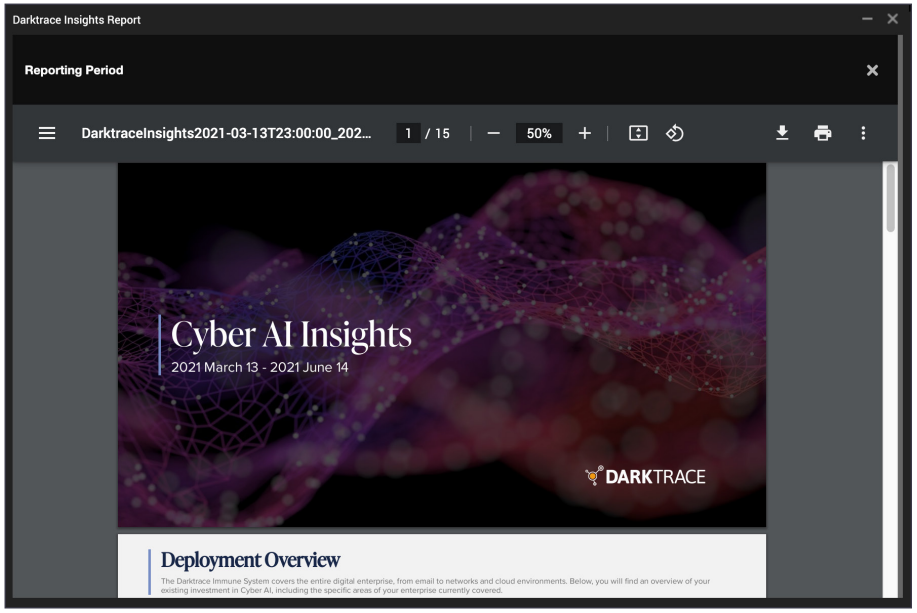
To generate a report, navigate to **Cyber AI Insights Report** in the main menu (**Reporting** > **Cyber AI Insights Report**). A new window will open. The following configuration options are available.

OPTION	DESCRIPTION
Time Period	Accepts a time range for the report period. The actual date period covered will be displayed in the two boxes beside the selector. Reports can generated for any period of less than 1 year, data-dependent.
Timezone	Alters the timezone of the report to that specified.
Estimated Incident Response Cost	To produce an estimate of the savings generated by Darktrace RESPOND, provide the expected cost in the chosen currency (default USD) of a successful cyber incident to your organization.
Analyst Hourly Wage	To produce an estimate of cost savings created by AI Analyst investigations, provide the expected hourly wage of a cyber analyst in the chosen currency (default USD).
Currency	Select a currency for monetary fields.
Include Human Response Data	Information about response times to model breaches (time to acknowledgment) can be optionally included if acknowledgment is part of your organizational workflow.
Behavior Visibility	Behavior categories are high level filters that allow an operator to focus in on specific levels of severity or behavior. There are four categories: Critical, Suspicious, Compliance and Informational. Select the categories to restrict data on the optional "Mean Time to Acknowledgement" page. Requires "Include Human Response Data": On.
Human Response Minimum Breach Score	For operator acknowledgment statistics on the optional "Mean Time to Acknowledgement" page, set a minimum model breach score for response time to be calculated against. This allows acknowledgement time to be focused on high-level events. Requires "Include Human Response Data": On.

Generate a Report

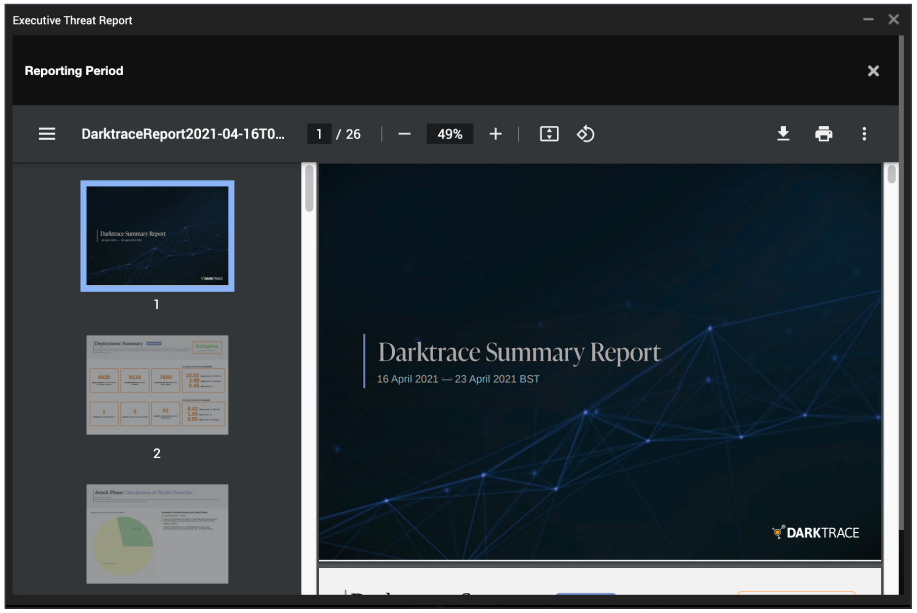
Once all desired options have been selected, click the **Generate Report** button to create the report. The generated report will appear within the popup window, where it can be downloaded as a PDF.

Previously generated reports and those over a larger timeframe can be retrieved from the **Download Reports** page (**Reporting** > **Download Reports**). Users with the "Download My Reports" permission can review reports they have generated there. Users with the "Download All Reports" permission can review reports generated automatically or by manually all users.



EXECUTIVE THREAT REPORTS

Executive Threat Reports can be generated within the Threat Visualizer and present a simple visual overview of model breaches and activity in the network environment. The report is separated into a graphical representations of network traffic, model breaches and Darktrace RESPOND response and an optional detailed breakdown more suitable for advanced audiences.



Reports can cover a single day or up to 12 months (data-dependent) and can be customized to include extra details or restricted to high level information. Executive Threat Reports can also be automatically generated and sent to recipients by email; this is configured on the System Config page. For more information on scheduling reports, please refer to Scheduling Automatic Report Generation.

Report Options

To generate a report, navigate to **Executive Threat Report** in the main menu (**Reporting** > **Executive Threat Report**). A new window will open. The following configuration options are available.

OPTION	DESCRIPTION
Report Period	Accepts a time range for the report period. The actual date period covered will be displayed in the two boxes beside the selector. Reports can generated for any period of less than 1 year, data-dependent.
Timezone	Alters the timezone of the report to that specified.
Subnets	Reports can be generated for a single subnet, or for multiple subnets. By default, all subnets are included. Multiple entries can be selected from the dropdown. Please note, enabling **SaaS Report** will clear this field.
Filter Breaches	Alters the report to include only ***unacknowledged*** model breaches (default), only ***acknowledged*** breaches or both ***acknowledged*** and ***unacknowledged***.
Behavior Visibility	Behavior categories are high level filters that allow an operator to focus in on specific levels of severity or behavior. There are four categories: Critical, Suspicious, Compliance and Informational. Select the categories to restrict included Model Breaches to only those that match the category filter.
Minimal Report	Generates only high level summary pages, excluding the detailed appendix.
SaaS Report	Restricts the content of the report to only SaaS user devices and SaaS breaches. The Darktrace RESPOND page is also restricted to RESPOND actions from platform environments such as Darktrace RESPOND/Zero Trust and Darktrace RESPOND/Apps only.
Include Comments	Includes any comments made by users or analysts in the detailed appendix.
Send Report Via Email	Emails the report to the users specified in the "Executive Threat Report Recipients" field. This field is configured on the Modules section of the Darktrace System Config page (**Workflow Integrations > Report Scheduler > Executive Threat Report**)

Generate a Report

Once all desired options have been selected, click the **Generate Report** button to create the report. The generated report will appear within the popup window, where it can be downloaded as a PDF.

Previously generated reports and those over a larger timeframe can be retrieved from the **Download Reports** page (**Reporting** > **Download Reports**). Users with the "Download My Reports" permission can review reports they have generated there. Users with the "Download All Reports" permission can review reports generated automatically or by manually all users.

CHAPTER 6 - THE SAAS CONSOLE

GETTING STARTED

THE SAAS CONSOLE 89
GETTING STARTED WITH THE SAAS CONSOLE 90
MAIN MENU 92
UNIVERSAL ELEMENTS 94

THE DASHBOARD

THE DASHBOARD 94
USING THE THREAT TRAY 96
EVENTS TAB AND LOGS 97
OTHER MODEL BREACH TABS 98

CYBER AI ANALYST

CYBER AI ANALYST FOR SAAS & CLOUD 99
TRIGGERED AI ANALYST INVESTIGATIONS IN THE SAAS CONSOLE 101
SAAS AND CLOUD INCIDENTS 102

PROFILES

USER PROFILES 103
TAGS IN THE SAAS CONSOLE 105

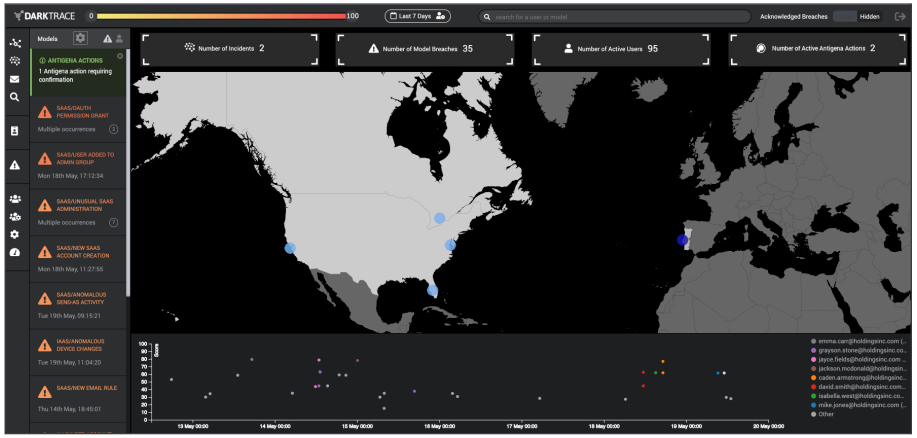
OTHER FEATURES

SAAS EXECUTIVE THREAT REPORTS 106
ASK THE EXPERT IN THE SAAS CONSOLE 107

THE SAAS CONSOLE

Introduction

Many organizations use SaaS (Software-as-a-Service) solutions as part of their everyday workflow for file storage, collaboration, and centralized access control, or Cloud solutions for their virtualized infrastructure and R&D. In many of these solutions, user interactions take place and organizational data is hosted in a third-party managed environment with a separate audit log and security interface; keeping track of alerts and activity across multiple SaaS and Cloud environments can be a challenge for many security teams.



The SaaS Console is a specialized user interface for investigating anomalous activity and user behavior in and across SaaS and Cloud environments, surfacing the events and analysis from Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules in one centralized location. The console is powered by the Cyber AI Analyst and Darktrace's unique 'pattern of life' anomaly detection; the interface is purpose built for monitoring and analysis in these environments whilst also maintaining existing workflows for operators already familiar with the Darktrace Threat Visualizer.

Please note, the SaaS Console is focused on administrative activity within Cloud environments. If you wish to extend visibility over network traffic in your virtualized infrastructure, please discuss the deployment of Darktrace virtualized sensors with your Darktrace representative.

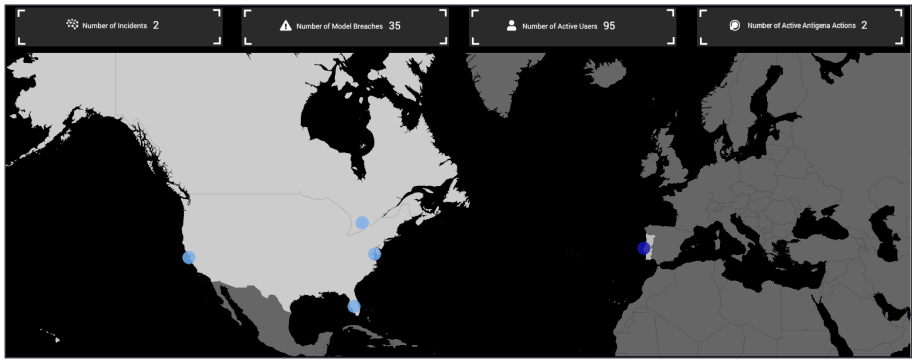
Access

Where Darktrace DETECT is monitoring SaaS, Zero Trust or Cloud environments and Darktrace/Email is monitoring your organizational email domains, the SaaS Console can be selected from the main menu of the Email Console. Pivot points to Darktrace/Email are also provided throughout the user interface and the Email Console is available from the SaaS console menu for easy movement between the interfaces.

Organizations that have Darktrace DETECT coverage over both network devices and users in SaaS and Cloud environments can access the SaaS Console from the main menu. The console is a dedicated interface for investigating model breaches and incidents in the SaaS environment; users can continue to investigate SaaS alerts within the main 3D Visualizer user interface if desired.

GETTING STARTED WITH THE SAAS CONSOLE

There are multiple entry points in the Darktrace SaaS Console through which analyst teams can begin seeing and reacting to identified alerts or incidents produced by Darktrace.



Model Breaches

A model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event: Darktrace model breaches are the most fundamental alert type. A model can be formed of specific triggers, specific thresholds for 'pattern of life' anomaly or multiple combinations of the two. The severity of a breach is calculated from the priority of the device, the level of anomaly associated with the breach, the priority of the model and the frequency of similar breaches.

The majority of Darktrace operators review, investigate, remediate and acknowledge model breaches as the basis of their workflow. Users have many different preferences as to how breaches are sorted and the order in which they are investigated; we recommend that new operators begin with the highest severity and work downwards.

Model breaches can be accessed from the threat tray on the left of the dashboard, the *Number of Model Breaches* statistic above the global map, or individually by clicking a node on the breach graph at the bottom of the dashboard.

AI Analyst

Cyber AI Analyst will review and investigate all Model Breaches that occur on the system as a starting point for its analysis process. It will then form incidents - a collection of one or more related events - centered around a SaaS user. The Cyber AI Analyst, with its global awareness and machine-speed investigation time, performs the heavy-lifting of the analysis process.

The AI Analyst will only surface incidents that are considered high priority - this can be a useful place to begin investigation and familiarization with the Darktrace environment.

AI Analyst can be accessed from the main menu or the *Number of Incidents* statistic above the global map in the dashboard.

Profiles

The profiles page provides an overview of all SaaS and Cloud users detected by Darktrace, as well as real-time feeds of actions performed and source locations for activity. Accounts will only be added to the profiles when activity is observed by the relevant DETECT module. Profiles are also grouped cross-platform where possible to give the fullest picture of user activity.

Operators are more likely to interact with individual profiles as part of an investigation workflow or when seeking further information on a specific user or platform.

The Profiles page can be accessed from the main menu or from the *Number of Active Users* statistic.

Suggested Workflow from a Cyber AI Analyst Incident

1. Access the SaaS Console and click the Cyber AI Analyst icon from the menu or the *Number of Incidents* statistic above the global map in the dashboard. Review the incidents created.

2. Select the most severe incident or the most interesting to you based on your knowledge of the business and network setup. Review the summary created by Cyber AI Analyst to quickly understand what each incident may involve.

3. Review the summary within the incident and understand how the events relate chronologically.

4. Scan the detailed event information to gauge the involved activity and actor. Optionally click on the account to review the user profile and historic locations for SaaS activity.

5. If the activity of interest relates directly to a model breach, investigate and optionally acknowledge the model breaches.

6. Optionally create a PDF report describing the events.

7. Acknowledge each event as the investigation is concluded.

Suggested Workflow from a Model Breach

When investigating an alert, a typical workflow will involve starting with summary information. As an example:

1. Access the SaaS Console and filter the threat tray to include a manageable number of Model Breaches - filtered by user or by model - using the severity slider and adjusting the selected time window.

2. Investigate the alerts that seem most interesting to you based on your knowledge of the business and SaaS platforms under monitoring. Regardless of how anomalous the activity is, you may prefer to begin with key users or SaaS and Cloud platforms, based upon the priorities of your investigation.

3. Review the breach description to understand what the alert concerns.

4. Identify the SaaS user performing the activity and review the pattern of life on the graph - has the event that triggered the breach been performed before?

5. Look at the expanded SaaS activity in the log that triggered the breach and the further information about the event shown on the right.

6. Confirm whether the user has any anomalous locations for access and activity.

7. Review the chain of behavior graphically and confirm any additional anomalous activity.

8. Look at whether the device has related anomalies at the time or previously.

9. Optionally pivot to Advanced Search for further investigation.

MAIN MENU

Threat Investigation

Home

The dashboard provides a global map of all SaaS logins, important statistics about SaaS and Cloud environments, and access to model breaches for triage and investigation.

AI Analyst

Cyber AI Analyst will review and investigate all Darktrace Model Breaches that occur on the system as a starting point for its analysis process. It will then form Incidents - a collection of one or more related events - centered around a device or SaaS user.

Email Console

Pivots to the Darktrace/Email interface for investigation and autonomous actions on your email domains. Please see the documentation on Darktrace/Email for more information.

Advanced Search

Opens the Advanced Search interface. Refer to the Darktrace Advanced Search introduction for more information.

Users, Actions and Reporting

Profiles

The profiles page provides an overview of all SaaS and Cloud users detected by Darktrace, as well as real-time feeds of actions performed and source locations for activity.

Darktrace RESPOND Actions

Presents currently active and expired actions taken by Darktrace RESPOND against users detected in external platforms and allows the actions schedule to be configured. Current and historic actions can be exported as a CSV. Refer to Reviewing Darktrace RESPOND Actions for more information.

Tags Manager

Tags offer a simple labelling system for users detected in external platforms (SaaS users) which can be used to identify important users, display relevant contextual data and alter model logic. Tags can be automatically applied as a model action, used in model defeats to exclude devices, or factored into model logic. Refer to Tags in the SaaS Console for more information.

Download Reports

Allows Threat Intelligence reports and Executive Threat Reports to be retrieved.

Executive Threat Report

Executive Threat reports are high level, visual overviews of activity and model breach events. Refer to SaaS Executive Threat Reports for more information.

Models

Model Editor

A Model is used to define a set of conditions which, when met, will alert the system to the occurrence of a particular event. Refer to The Model Editor for more information.

System Configuration and Administration

User Admin

User permissions can set access and view on a per-user basis. See Guide to User Privileges or the Darktrace System Administration Guide for a full list of User Permissions.

User Groups

Allows permissions to be controlled in groups which are then assigned to multiple users, instead of on a per-user basis.

System Config

Provides all configuration settings for the Darktrace Threat Visualizer application including Darktrace RESPOND settings and authentication for SaaS modules. Alerting options can be configured here.

System Status

System metrics such as ingested traffic quality, master/probe reachability, and protocols observed can be reviewed. The visible statistics on the Status page will depend on your environment.

Help

Ask the Expert

Ask the Expert allows for the creation of incidents which can be submitted to Darktrace analysts for feedback. Refer to Ask the Expert in the SaaS Console for more information.

User Settings

Account Settings

Change settings for the current account including changing password, enabling Accessibility Mode or registering the Darktrace Mobile App.

UNIVERSAL ELEMENTS

A severity slider is displayed at the top left - by default, all model breaches with a score between 0 and 100 will be shown. Moving the controls will filter out breaches which do not have a score within the range selected. In user view, users will be filtered by the scores of the model breaches associated with them. In User view, all users will be displayed when the slider is set to 0.

The sensitivity slider does not affect any underlying processing, only what is currently displayed in the dashboard.



The timeframe to display model breaches and AI analyst incidents for is located in the top right. The available options are: *Today*, *Yesterday*, *Last 24 Hours*, *Last 7 Days*, *Last 30 Days* or *This Month*. Clicking the user icon ( ) will provide options to change the current time zone.



The **search** function can be used to view user profiles or filter the dashboard by a specific SaaS service or a specific user.

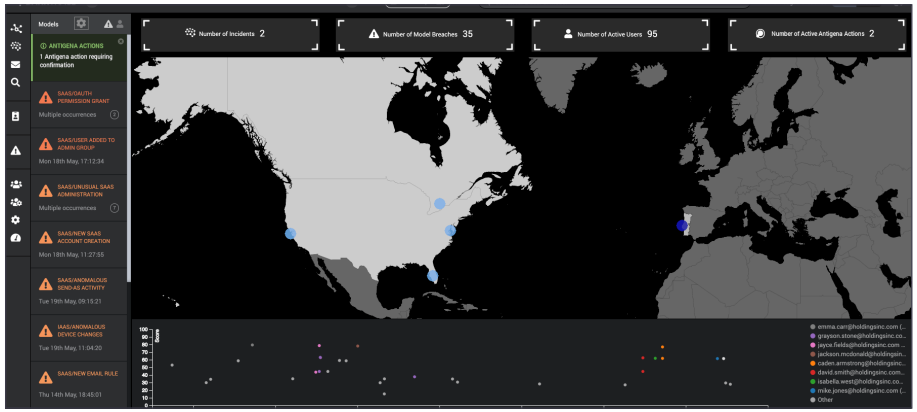
- The AI Analyst icon opens a dialog to launch an AI Analyst investigation on the user.
- The profiles icon will open the profile for the specified user.
- The filter icon will filter all data by the platform or user specified. Filters can be cleared by clicking the ✖.

For example, searching for "Office 365" and clicking the filter icon will filter the dashboard (including model breaches) and the profiles page by only users within office 365. A user filter can be applied at the user level (e.g., benjamin.ash@holdingsinc.com) which will filter for their accounts across multiple platforms, or platform-specific (e.g., "benjamin.ash@holdingsinc.com (Office 365)").

Using the **View Acknowledged** toggle will hide (*Hidden*) or include (*Visible*) acknowledged model breaches and AI Analyst incidents. This will also affect the *Number of Incidents* and *Number of Model Breaches* statistics displayed on the dashboard.

To **logout**, click the ( ) button in the top right of the window.

THE DASHBOARD

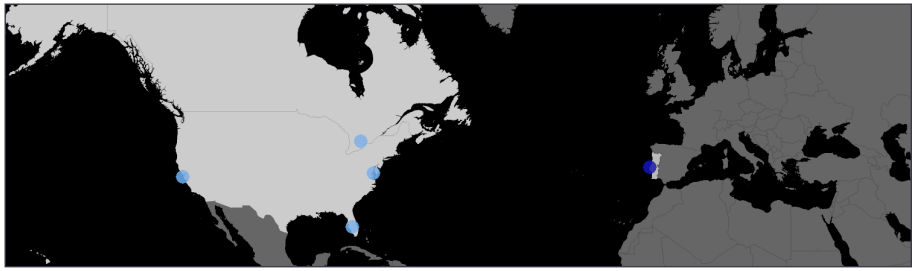


Four headline statistics can be reviewed at the top of the dashboard; clicking on a statistic will open the relevant interface.

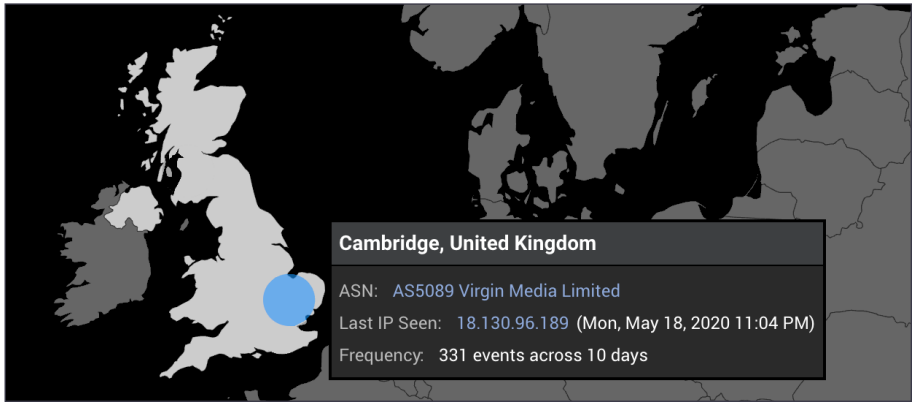


- **Number of Incidents** opens AI Analyst for the timeframe selected.
- **Number of Model Breaches** will expand the model breach with the highest score in the timeframe selected.
- **Number of Active Users** opens the Profiles page.
- **Number of Active Antigena Actions** (*SaaS actions only*) opens the Antigena Actions page.

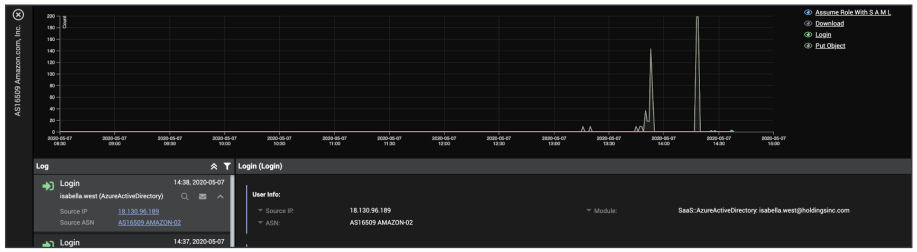
The dashboard displays a **global map of all SaaS and Cloud login activity** over the last three weeks, grouped by the ASN that the activity is performed from.



An ASN is a group of IP addresses controlled by one Internet Service Provider or entity. As SaaS activities are often performed by users on mobile devices or tethering to mobile devices, their IP may change regularly but the ASN will stay consistent (usually the mobile provider). Greater variation in IP address is therefore not as indicative of anomaly as large changes in ASN usage.



Lighter locations indicate sustained activity and darker locations are those which are rare for the organization. Click on any location marker to see the associated IP addresses, the ASN and the frequency at which the activities were performed. Within this information pop-up, click the **IP address** or **ASN** to open a log overlay with events seen from that location within the timeframe selected.



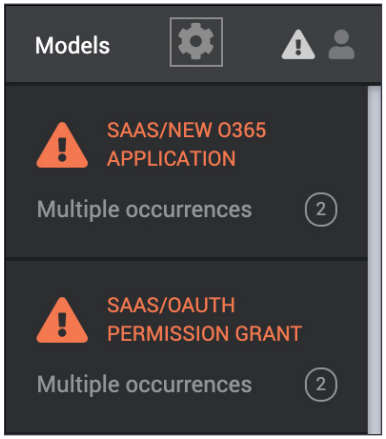
The breach graph displays all model breaches within the time frame and the associated breach score. Nodes are colored by user, with only users with the highest scoring model breaches during the timeframe individually represented. Clicking on a user will open the relevant entry in the Profiles page.



Hovering over a node will produce a tooltip with the exact time and model that was breached. Clicking a node will open the breach investigation window, focused on the specific breach.



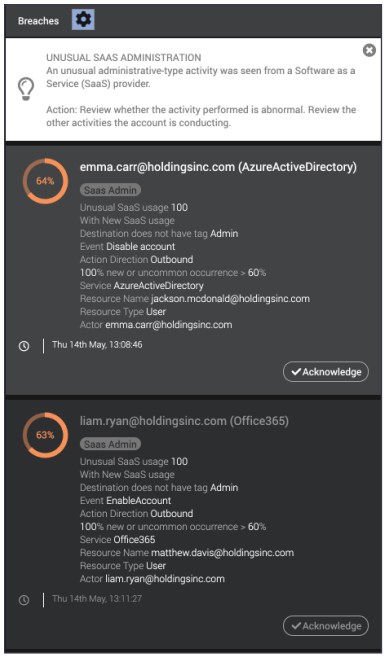
USING THE THREAT TRAY



On the left-hand side of the dashboard is the **threat tray**, displayed per *Model* or per *User*. This can be altered by clicking the model ⚠ button or the user 👤 button to switch view.

Clicking the cog ⚙ icon will display sorting options for the returned breaches. Model breaches can also be filtered by the parent folder that the model is contained within.

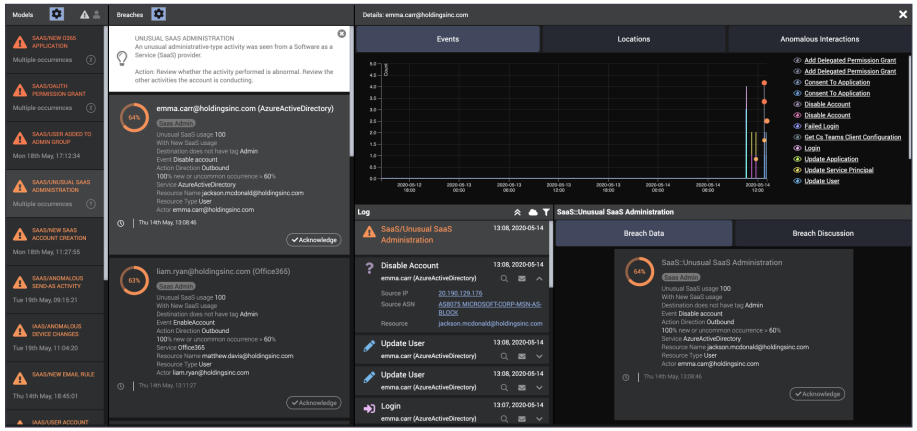
Clicking on a model breach or a user will open the **breach investigation overlay**, centered on the selected model or user. The investigation overlay can be closed by clicking the ✖ or the model/user entry once more.



The central column contains an entry for every model breach; each entry contains important information about the breach event including the time, the breach ID, the SaaS service, the user involved and any ***display fields*** that have been assigned to the model in the Darktrace model editor. Where a model has breached multiple times, or a user triggered multiple model breaches, this column will contain multiple entries - the cog ⚙ icon will display sorting options for these entries. A description of the model and a suggested action to take in response to the breach are displayed at the top.

When a model breach is selected, the log will open and the relevant entry for the action that triggered the breach will be expanded to show additional metrics. By default, the breach investigation window will open on the first sorted entry in **Breaches** list.

Details Pane



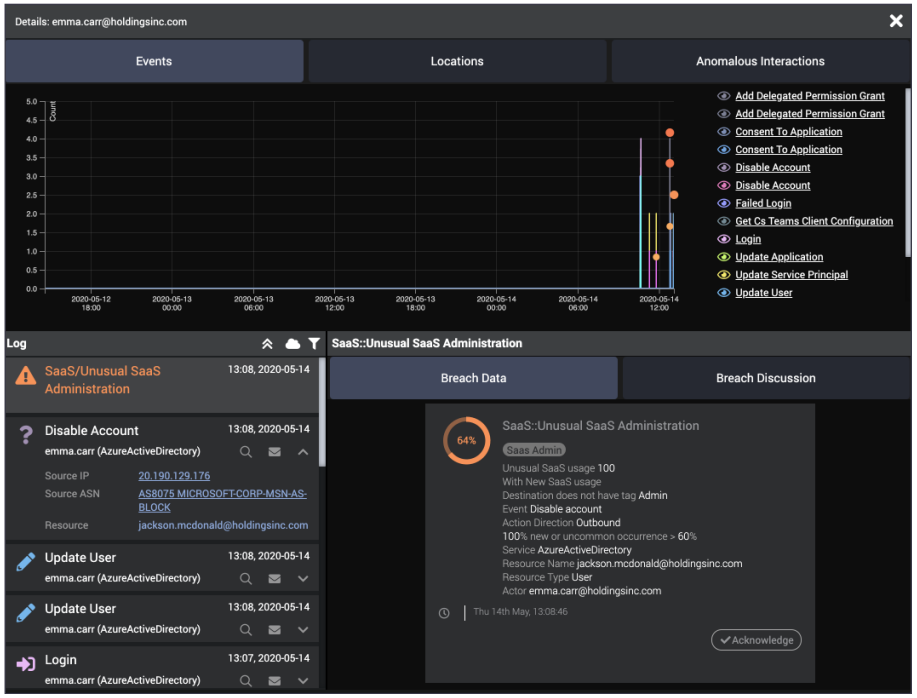
On the right side of the investigation overlay is the details pane. This window contains important information about the breach, the user, their historic locations for SaaS activity and the specific interactions that triggered the breach.

The user under investigation will always be displayed in the top left of the window, for example "Details: carla.hurst@holdingsinc.com".

There are three tabs available for the details pane: **Events**, **Locations** and **Anomalous Interactions**.

EVENTS TAB AND LOGS

Events Tab

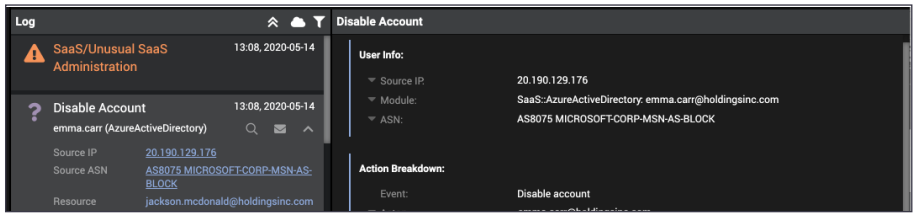


The Events tab will display a graph of all activity by the user under investigation from the start of the timeframe until the time of the model breach. This is the default view. The graph and the logs can be focused on a specific timeframe by clicking and dragging a window on the graph itself.

In addition to user activity, logs will include events added by Darktrace pattern of life modeling (such as "Unusual Activity"), Darktrace RESPOND actions created for the user, and high severity Darktrace/Email model breaches (***select Darktrace/Apps Modules only***).

Listed beside the graph are all event types currently displayed. Clicking the name of an event in the key will hide it from the graph; hiding an event type in the graph will also hide it in the log and vice versa. Shift-click the name of an event to hide all other types apart from that event.

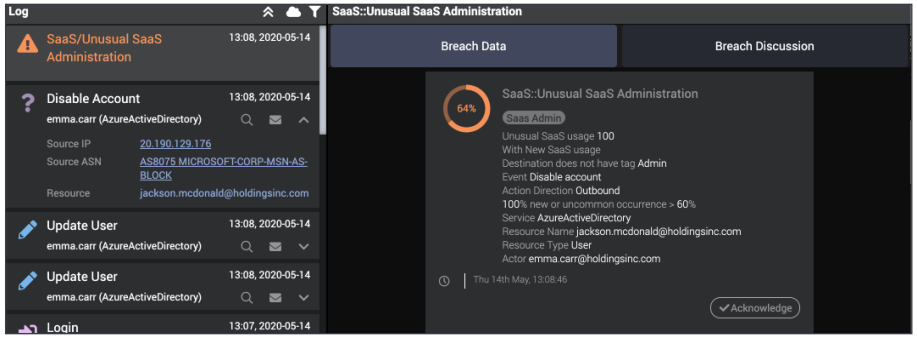
Logs



At the top of the log pane are three icons:

- The double arrow icon ⌃⌃ collapses all log entries so no additional metrics are visible.
- The cloud icon ☁ allows the log entries to be filtered by a specific SaaS service.
- The filter icon ▼ controls the types of event shown in the log. Individual (and multiple) events can be toggled, and model breaches and unusual activity notices can also be enabled or disabled.

Log Entries



Each log entry is composed of an icon describing the kind of action or event, a short description of the entry, the time it occurred, the user that caused it and the SaaS or Cloud platform it was performed on. When the log is opened from a model breach, it will open with the model breach selected, and the relevant entry for the action that triggered the breach will be expanded to show additional metrics. Clicking 'Load More' will load additional event log entries until the end of the selected timeframe is reached.

Beside each event, one or more of the following icons may be displayed:

- The search icon pivots to the exact entry in Advanced Search.
- Where Darktrace/Email is also monitoring the email environment, the email icon pivots to the Darktrace/Email console filtered on the selected user.
- The down and up arrows show or hide additional metrics in the logs. All log lines can be collapsed by clicking the double arrows icon at the top of the log.
- Where Ask the Expert is enabled, the clipboard icon will appear against log entries. This icon allows log entries to be copied to a clipboard functionality for inclusion in Ask the Expert questions. See Ask the Expert in the SaaS Console

Please note, some buttons may not be visible due to user permissions.

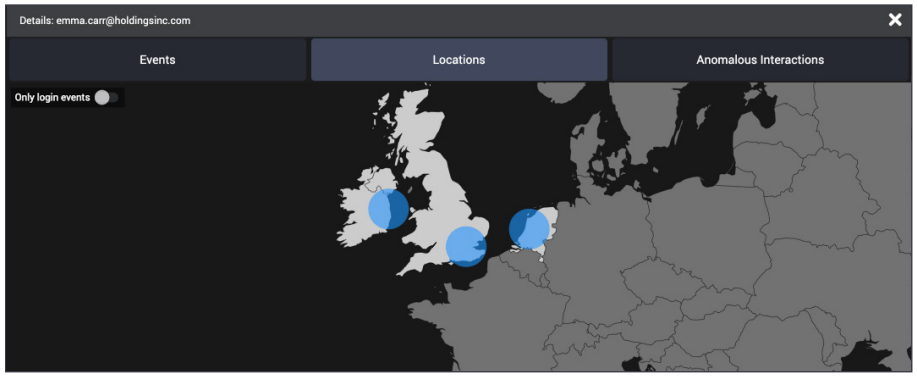
The pane at the right of the log will always show further information about the action, breach, IP or ASN selected:

- When a model breach is selected, the right pane can be toggled between the breach summary and discussion, where comments on the breach can be made or reviewed.
- When an action (log entry) is selected, the right pane will display further metrics about the activity. A small arrow beside a metric () provides additional options such as filtering the logs by the selected metric (for example, by a source IP) or reviewing more detailed information in Advanced Search.

  Any filters added via this method can be removed through the filter icon above the logs.
- When an ASN or IP is selected, the pane will contain a summary of activity and model breaches associated with that source over the selected timeframe. The option "View Logs" will open a new, focused log window with activity and logs from that specific IP or ASN across all users in the selected timeframe.

OTHER MODEL BREACH TABS

Locations Tab

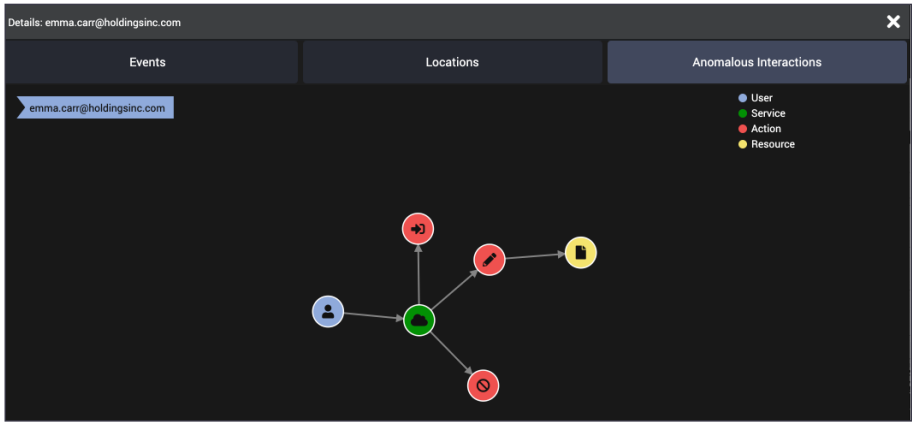


The locations tab displays a global map of SaaS and Cloud activity for the user under investigation from the start of the selected timeframe, grouped by the ASN that the activity is performed from. Lighter locations indicate sustained activity and darker locations are those which are rare for the organization.

The map can be toggled to show login events only using the slider in the top left.

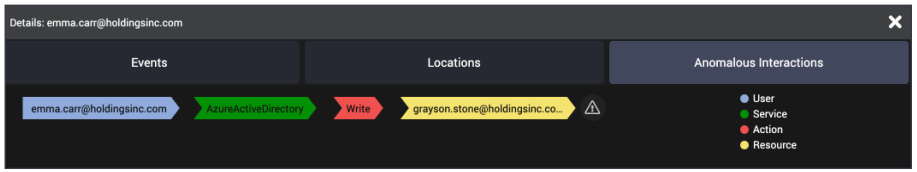
Clicking on any location marker will produce a pop-up with the associated IP addresses, the ASN and the frequency at which the activities were performed. Within this information pop-up, click the IP address or ASN to open an additional event log with events seen from that location within the timeframe selected. These events are restricted by user.

Anomalous Interactions Tab



The anomalous interactions tab provides a simple visualization of all activity involved in the model breach. Each action is visualized as a series of connected nodes with directional arrows. Clicking on an end node will fill in the chain of activity at the top left of the tab. The key distinguishes resources, actions, services and users. Every chain of behavior relates to a model breach.

An example chain of behavior is represented by a user (central node) acting upon a SaaS service (cloud icon node connected to user node), performing a delete (trash can icon connected to the cloud node) on a file (file icon node connected to the trash can).



Each node uses a simple graphical representations to give an at-a-glance awareness to a cyber-analyst, allowing them to quickly establish the chain of anomalous behavior. For example, a large number of file icons under a download icon indicate a user potentially about to exfiltrate data - this can be seen quickly and without diving into the details.

Introduction

The Darktrace Cyber AI Analyst investigates, analyzes and triages threats seen within your Darktrace environment. From a global position within the environment, the Cyber AI Analyst accelerates the analysis process, continuously conducting investigations behind the scenes and operating at a speed and scale beyond human capabilities.

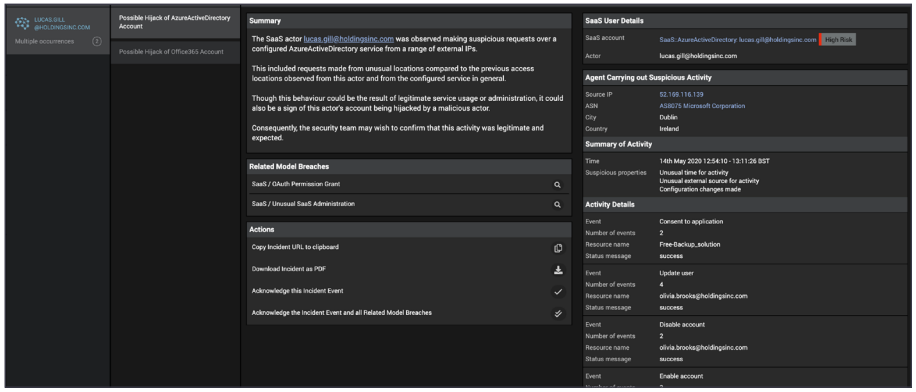
The Cyber AI Analyst will review and investigate all Darktrace Model Breaches that occur on the system as a starting point for its analysis process. It will then form Incidents - a collection of one or more related events - centered around a device or SaaS user. Incidents involving multiple devices will be classified as 'cross-network'. Incidents created by Cyber AI Analyst do not encompass all model breaches in the environment, only those deemed particularly unusual or in need of further investigation.

AI Analyst incidents displayed within the SaaS Console are filtered to only show those specific to SaaS and Cloud platforms. Network incidents may be reviewed from the Threat Visualizer homepage.

Please note, the Cyber AI Analyst will only investigate model breaches of default Darktrace models. Custom models are not currently supported.

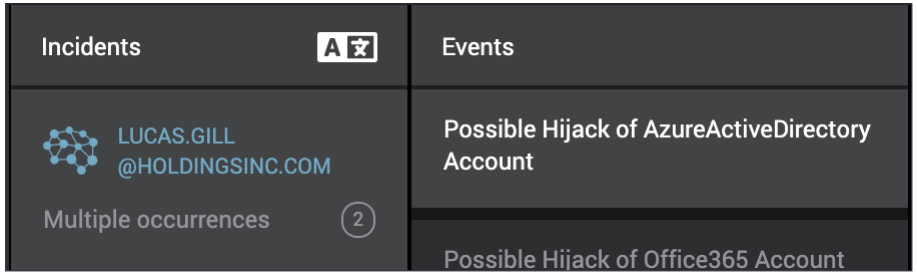
Cyber AI Analyst View

AI Analyst for SaaS and Cloud can be accessed via the Incident stats on the dashboard or by clicking the AI Analyst icon from the main menu. Incidents are categorized from blue to white, where white indicates the highest level of threat.



Click the language icon to select an alternative language for AI Analyst incidents.

Click the AI Analyst icon to switch between **Incidents** and **Investigations**. Incidents are generated automatically in response to model breaches. Investigations are triggered manually and may result in an incident if anomalous behavior is found.



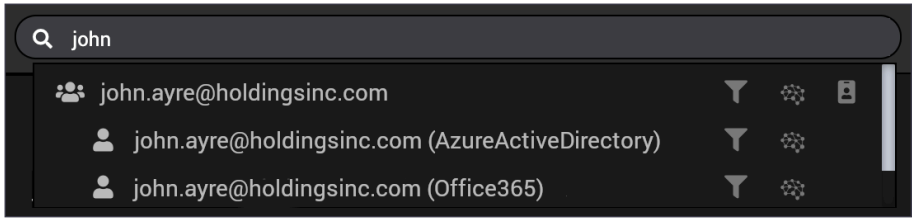
An incident is composed of one or more events; events are a group of anomalies or SaaS and Cloud user actions investigated by Cyber AI Analyst that pose a likely cyber threat. Click on an Incident to review the available “Events” and the “Details” for those events.

TRIGGERED AI ANALYST INVESTIGATIONS IN THE SAAS CONSOLE

Investigating a User

To trigger an AI Analyst investigation, a **device** or **SaaS user** and a **time** are required. The **device** is the subject of the investigation and the **time** is the starting point for analysis.

An AI Analyst investigation can be triggered on a specific device from the omnisearch bar with the AI Analyst icon, or from a specific **Profile** in the SaaS Console with the "**Investigate**" button.



When an investigation is launched from a specific device or SaaS profile, the **Device** or **Account** field is already completed.

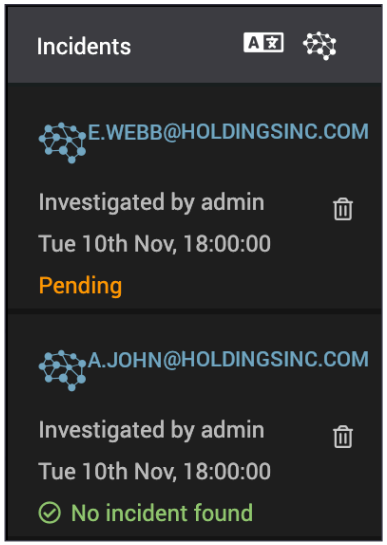
Click **Investigate** to start the investigation.

Investigations

Investigations display the device or user under analysis, the analysis time, the user who triggered the analysis and the current investigation state.

Investigations have three states - pending, processing, finished. A *pending* investigation waiting for AI Analyst to start the analysis. A *processing* investigation_ means analysis is underway and a *finished* investigation means the investigation has concluded.

In the SaaS Console, the AI Analyst view can be toggled between **Incidents** and **Investigations** using the AI Analyst icon at the top of the first column (Incident tray) to see current and completed investigations.

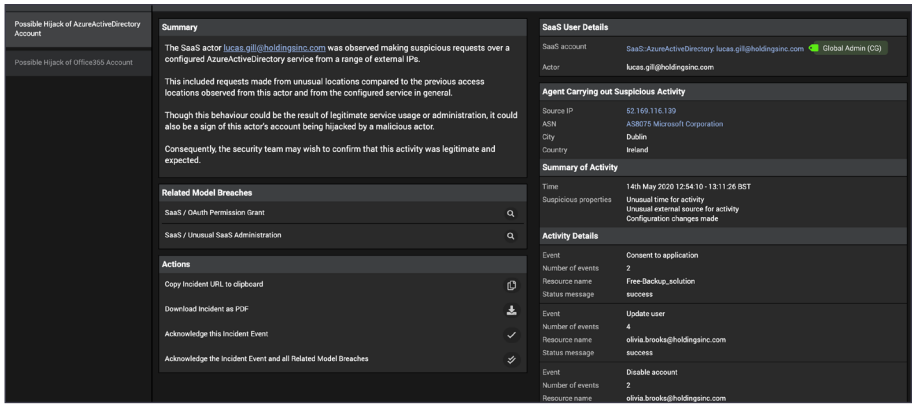


When an investigation is triggered, AI Analyst will perform a close analysis of the activity for the user for a period of roughly one hour, using the time provided as a central point. It will then contextualize this behavior against historic activity and connectivity for the entity and its peers. If a finished investigation has the result "No Incident Found", AI Analyst did not locate any identifiable anomalous activity around the investigation time. If anomalous activity is detected, one or more incident events will be created.

Click the incident tile to open the investigation result. Incidents created from triggered investigations follow the same format as AI Analyst incidents created automatically.

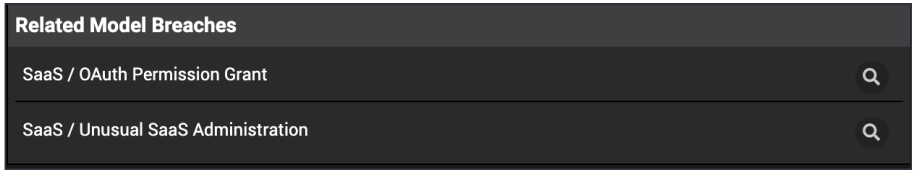
SAAS AND CLOUD INCIDENTS

Selecting an incident will display the events that comprise the incident. Further information for each event displays on the right-hand side in the “Details” pane.



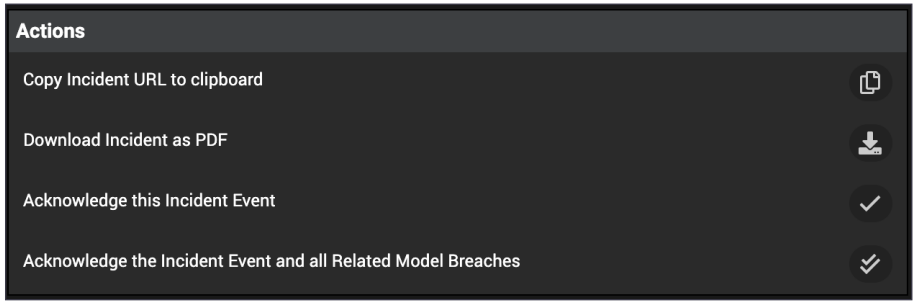
The summary gives an accessible overview of the activity AI Analyst has detected and a potential response to the activity an operator may take. Model breaches that triggered Cyber AI Analyst investigation will be listed as related model breaches. The Cyber AI Analyst is not reliant on model breaches, they primarily provide a starting point for its deeper analysis of the user activity around the anomaly time.

- The search icon 🔍 will open the breach investigation overlay, centered on the specific model breach.



Like a human, the Cyber AI Analyst uses a model breach as a starting point to investigate the device. The behavioral analysis it performs may discover anomalies or patterns of activity that were not the original trigger point for the model breach but are worthy of investigation. Consequently, the event period may not correspond with the breach time. Additionally, some model breaches require sustained behaviors such as repeated actions before breaching, so the final breach trigger may be later than the action of interest.

Further investigation by an operator should prioritize the anomalous activity signposted by Cyber AI Analyst first and the model breach conditions second if the two do not directly align.



Actions

- The download icon creates a downloadable PDF report of the Incident.
- The copy icon copies a link to the incident to the clipboard.
- The check icon ✔ acknowledges the incident event. If all events are acknowledged, the incident will be shown/hidden depending on the user settings.
- The double check icon acknowledges the incident event **and** all related model breaches.

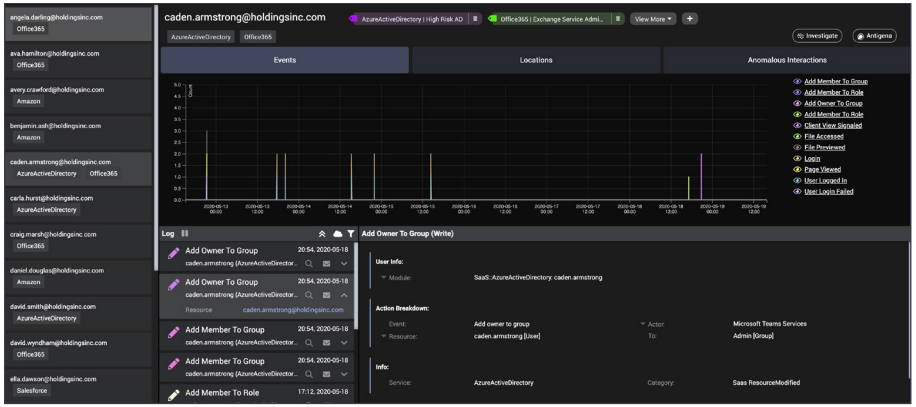
Finally, the right-hand panels will breakdown key elements of the event and the involved users detected in external platforms (SaaS users); the data is specific to each event type.

SaaS User Details	
SaaS account	SaaS::AzureActiveDirectory: lucas.gill@holdingsinc.com High Risk
Actor	lucas.gill@holdingsinc.com
Agent Carrying out Suspicious Activity	
Source IP	52.169.116.139
ASN	AS8075 Microsoft Corporation
City	Dublin
Country	Ireland
Summary of Activity	
Time	14th May 2020 12:54:10 - 13:11:26 BST
Suspicious properties	Unusual time for activity Unusual external source for activity Configuration changes made
Activity Details	
Event	Consent to application
Number of events	2
Resource name	Free-Backup_solution
Status message	success
Event	Update user
Number of events	4
Resource name	olivia.brooks@holdingsinc.com
Status message	success

USER PROFILES

The profiles page lists all SaaS and Cloud users detected by Darktrace; only users that have performed an action whilst Darktrace monitoring was enabled and during the specified time window will be listed. Here, user activity can be reviewed and investigated in more detail and across SaaS and Cloud platforms.

The full page can be accessed from the main menu icon or via the Active Users stat on the dashboard. Individual profiles can also be reviewed by using the search bar or as part of an investigation workflow, for example, by clicking on an actor within an AI Analyst incident summary.



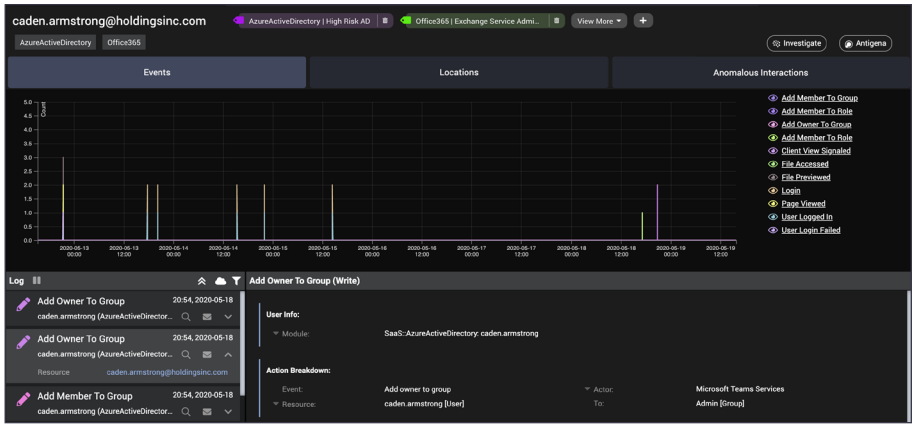
Each entry lists the user and SaaS or Cloud platform(s) they have been detected on by the relevant Darktrace/Apps, Darktrace/Cloud or Darktrace/Zero Trust module. The first listed profile will open by default.

Using the search functionality to filter the profiles by SaaS or Cloud platform can be a useful way of limiting the returned profiles to a specific subset of users or accounts.

User Pane

Profiles contain the same elements as the breach investigation window: a log and graph of user activity, further information about those actions, historic locations for SaaS activity and any anomalous interactions seen for the user. Unlike the breach window, information in these elements is displayed from the start of the specified timeframe until the current time and is updated in real time when new events are seen.

The selected user is displayed in the top left of the window - SaaS and Cloud platforms they have been detected upon are listed below. Tags applied to the user across their detected SaaS accounts are also displayed and additional tags can be added. See Tags in the SaaS Console for further information on adding and editing tags.



On the right are the action buttons:

- The "**Investigate**" button opens a dialog to launch an AI Analyst investigation for the user.

- The "**Launch RESPOND Action**" button opens a dialog to perform manual Darktrace RESPOND actions on the user.

  For users detected by multiple Darktrace RESPOND-enabled modules, selecting "All Eligible Modules" allows a "Disable User" RESPOND action to be triggered for the user against all modules at once.

Tabs

- In the **Events** tab, all activity the user has performed in the selected timeframe is displayed on the graph and in the logs. Please see Events Tab and Logs for further information about interacting with the logs.

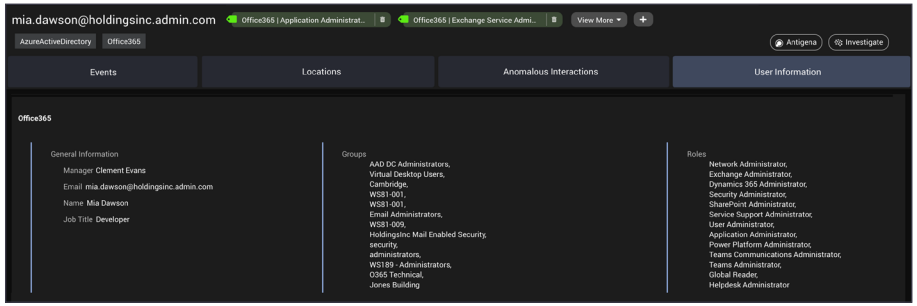
  Profile logs are updated in real time - the log can be paused or resumed with the pause ⏸ and play ▶ icons.

- The **Locations** tab contains all historic locations that the user has performed SaaS or Cloud activity from over the specified timeframe. The map can be toggled to show login events only. More information about interacting with the locations tab can be found in Other Model Breach tabs.

  Logs for ASN or IP are filtered by the specific user when accessed from the user profile. To review logs for all users, please use the 'View Logs for this IP' or 'View Logs for this ASN' buttons in the summary.

- The **Anomalous Interactions** tab displays anomalous chains of behavior that triggered a model breach within the specified timeframe. Please see Other Model Breach tabs for further information.

- The **User Information** tab (*select SaaS Modules only*) displays contextual data about the user including group membership, assigned roles and assigned licenses, retrieved from the SaaS environment. This contextual data provides valuable insight when investigating user behavior and potentially anomalous access.

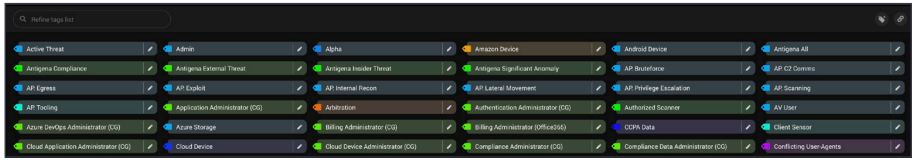


TAGS IN THE SAAS CONSOLE

Tags offer a simple labelling system for users detected in external platforms (SaaS users) which can be used to identify important users, display relevant contextual data and alter model logic. Tags can be automatically applied as a model action, used in model defeats to exclude devices, or factored into model logic. For SaaS users, tags may be also be applied automatically by the module to identify the roles assigned to the user - these tags are identified by "(CG)".

A small number of tags are restricted for system use only and are used to indicate devices excluded from monitoring or behaving erratically.

Viewing and Editing Tags



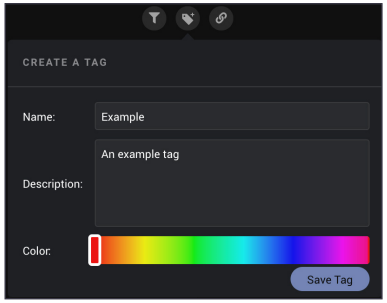
Tags can be managed in the tags manager, accessible from both the SaaS Console or the Threat Visualizer, or via the API (/tags). Select **Tags Manager** from the **Main Menu** to open the manager.

A list displaying all current tags will appear, optionally filterable with the **search**.

To review a tag, click on it. A summary page will specify (where applicable) the devices tagged, models that apply the tag, models dependent on the tag for breach logic and models that have the tag applied to them. All models and devices or users can be clicked. To further refine the list, the filter icon can be used to add an additional tag, filtering only by models/devices that share both tags.

To edit or delete a tag, click the edit icon beside the tag name.

Add a New Tag



To add a new tag, click **Add Tag** from the Tags Manager.

When creating a new tag, a description can be added to help identify the purpose of the tag. Selecting a color assists identifying the tag when assigned to a device. Click **Save**.

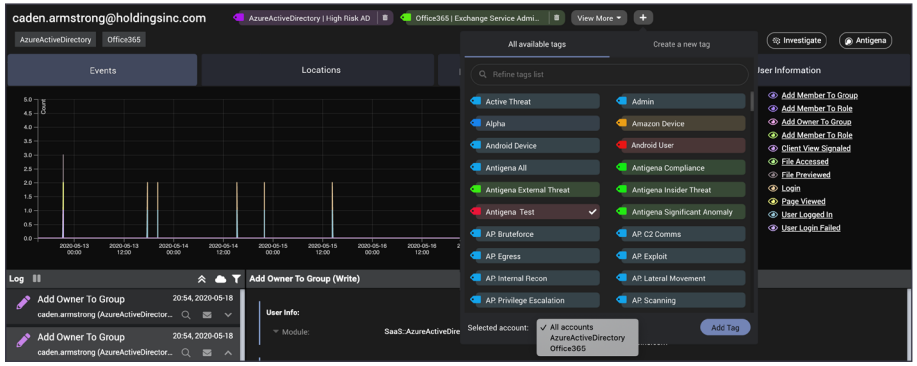
The new tag will appear on the tag manager.

Tag a SaaS User

In the SaaS Console, tags are added on the profiles page for the desired user.

Tags are applied to and removed from users on a per-module basis. As Darktrace may observe a user on multiple SaaS platforms, tags on the user's profile indicate which environment they are associated with. For example, "Office365 | Application Administrator (CG)".

Click the plus icon to open a dialog to add an additional tag. In the pop-up window, locate the tag and click on it to add to the device. Where the user has multiple SaaS platforms associated, the account to be tagged can be altered with the dropdown in the bottom left.



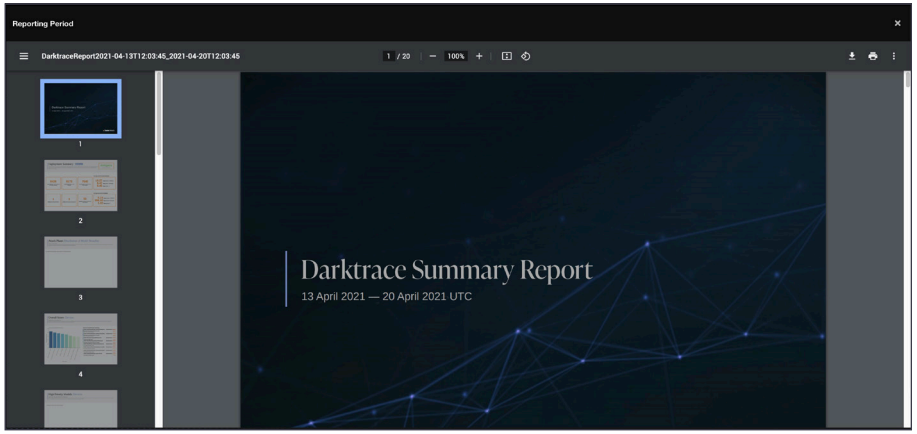
When the user has many tags, these are collapsed and can be shown in a popup, searchable dialog by clicking "View More"

SAAS EXECUTIVE THREAT REPORTS

Executive Threat Reports can be generated within the SaaS Console and present a simple visual overview of model breaches and activity in the SaaS environment. The report is separated into a graphical representations of network traffic (***environments with network and SaaS only***), model breaches and Darktrace RESPOND response and an optional detailed breakdown more suitable for advanced audiences.

Reports can cover a single day or up to 12 months (data-dependent) and can be customized to include extra details or restricted to high level information.

In the SaaS Console, the content of the report is restricted to only SaaS user devices and SaaS breaches. The Darktrace RESPOND page is also restricted to RESPOND actions from platform environments such as Darktrace RESPOND/Zero Trust and Darktrace RESPOND/Apps only. This is the equivalent of the "SaaS Report" option in the main Threat Visualizer.



Report Options

To generate a report, navigate to **Threat Report** in the main menu - the following configuration options are available.

Report Period

Accepts a time range for the report period. The actual date period covered will be displayed in the two boxes beside the selector. Reports can generated for any period of less than 1 year, data-dependent.

Timezone

Alters the timezone of the report to that specified.

Filter Breaches

Alters the report to include only ***unacknowledged*** model breaches (default), only ***acknowledged*** breaches or both ***acknowledged*** and ***unacknowledged***.

Minimal Report

Generates only high level summary pages, excluding the detailed appendix.

Include Comments

Includes any comments made by users or analysts in the detailed appendix.

Send Report Via Email

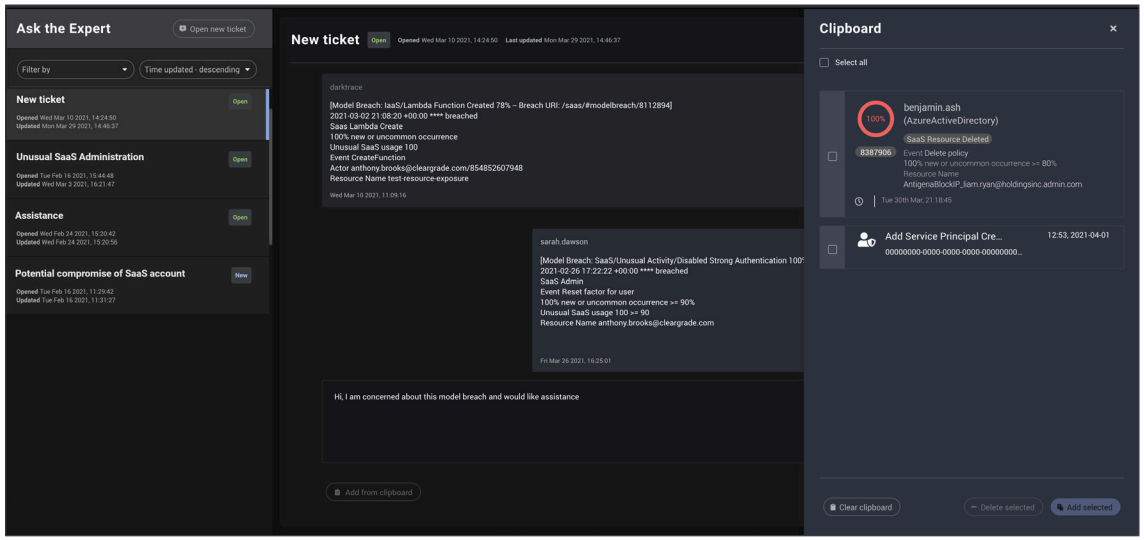
Emails the report to the users specified on the System Config "Executive Report Recipients" field. This can be found beneath the configuration for Email Recipients, in the Workflow Integrations section of the Modules page.

Generate a Report

Once all desired options have been selected, click the **Generate Report** button to create the report. The generated report will appear within the window, where it can be downloaded as a PDF.

Previously generated reports and those over a larger timeframe can be retrieved from the **Download Reports** page.

ASK THE EXPERT IN THE SAAS CONSOLE



What is Ask the Expert

Ask the Expert allows for the creation of incidents which can be submitted to Darktrace analysts for feedback. Submitted incidents can be reviewed and commented upon both by Darktrace analysts and local users. Questions have associated statuses and function as "tickets" that can be closed when resolved. In the SaaS Console, **Ask the Expert** is available from the Main Menu.

Incidents can be created with just textual input, or specific log lines and model breaches can be copied to a special clipboard for use in an incident. Where the clipboard icon appears, the element can be copied to the Ask the Expert clipboard. This element, with any other copied element, can then be inserted into a new or open Ask the Expert ticket using "Add from Clipboard".

Open Tickets

Within the Ask the Expert page, previously submitted questions are listed with their timestamp, owner, and status. Tickets can be filtered by status or sorted alphabetically or chronologically.

Clicking on a question allows Darktrace responses to be reviewed and new comments or log elements to be added. Questions can be closed or deleted – closing will mark the status of the question as closed and resolved, without actually removing from the question list.

CHAPTER 7 - ADVANCED SEARCH AND INTEL

ADVANCED SEARCH

DARKTRACE ADVANCED SEARCH 109
NAVIGATING ADVANCED SEARCH 110
SEARCHING ADVANCED SEARCH 111
CHANGING THE LAYOUT AND FILTERING ADVANCED SEARCH RESULTS 113
BUILDING A COMPLEX ADVANCED SEARCH QUERY 115

THREAT INTEL

TAXII CONFIG 117
WATCHED DOMAINS 120
TRUSTED DOMAINS 121

DARKTRACE ADVANCED SEARCH



What is Advanced Search?

Darktrace analyzes network traffic through Deep Packet Inspection; each connection is processed and logs containing key metadata about the connection are produced. Similarly, Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules and other integrations analyze activity data from third-party platforms and create log entries for each event. The Threat Visualizer interface displays a relevant subset of this metadata in device event logs and breach logs, however as an investigation progresses, the need can arise to review connections and events in more detail.

The Advanced Search interface provides searchable access to the detailed metadata logs produced by network traffic and event analysis. For advanced users, complex query syntax and data aggregation analysis are available. For those just getting started, simple queries can be built up slowly and saved for future use.

Launching Advanced Search

There are two ways to access Advanced Search: from an event of interest or from the main menu.

Main Menu

Select "**Advanced Search**" from the main menu of the Threat Visualizer, or "**Advanced Search**" from the SaaS Console sidebar.

Event of Interest

To review a Threat Visualizer event log entry in Advanced Search, click the ▾ triangle icon beside the entry and select the option "View advanced search for this event". This applies to all log types, including the model breach log, device event log, and event log. This option may not be available for all connections or events.

For the SaaS Console, Advanced Search can be accessed from any event log line with the search icon. Additionally, the ▾ triangle icon may appear beside detailed event properties such as actor or IP address. Selecting "Search Advanced Search for [this value]" from the options will open a simple search for that characteristic.

NAVIGATING ADVANCED SEARCH



Overview

By default, all events from the last 15 minutes are displayed in order by time, with the most recent events at the top of the screen. If Advanced Search is accessed for specific event, the time will be focused on the event time and only relevant log lines will be shown.

The Darktrace icon in the top left will reset the page to the current time with no search and the default timeframe selected.

Searches can be manually typed out in the search bar or constructed from the building blocks available - how to perform a search will be explained in more detail later.

To the left of the search bar is the current search timeframe - this can be selected from one of the preset options or entered as a custom timeframe using the date selectors above the graph. The search period can also be moved forwards or backwards using the arrows below the graph.

The graph shows the results that match the search criteria across the timeframe selected, grouped by the unit specified in the dropdown. The grouping defaults to larger units like Day or Hour when the search period is longer. Hovering over a bar of the graph will display how many events it contains. Clicking within the graph and dragging to create a window will zoom in on that timeframe and filter the log events accordingly.

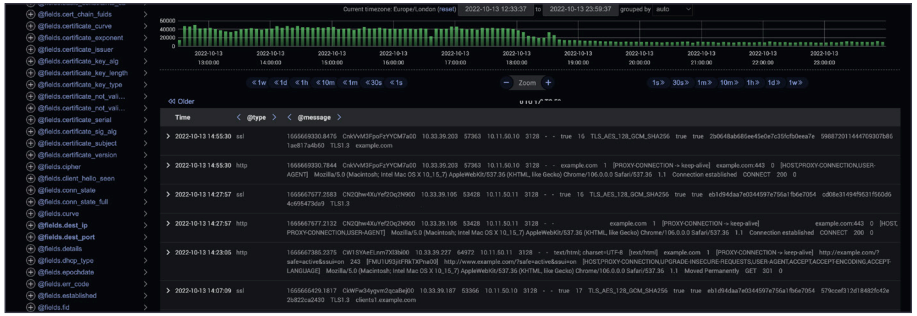
Each page displays 50 results, sorted by time from newest to oldest. The "**◀◀ Older**" and "**Newer ▶▶**" buttons can be used to move through the results, or the **▶| return** button to return to the start. Matching results can also be exported as a CSV file.

The notepad functionality can be used to note down interesting connections or analysis comments. Local storage must be enabled to use this feature, as well as saved layout and searches.

Default Layout

By default, the columns displayed are (from left to right):

- **Time**: Date and time that the log entry was created, in the time zone of the device used to access Advanced Search. Dates are shown in ISO format: **YYYY-MM-DD hh:mm:ss**.
- **@type**: The log entry type - may be the protocol that was inspected, an event such as a connection or refer to the module that created the entry. For example: `@type:conn` indicates a connection, `@type:kerberos` indicates the protocol Kerberos was inspected and `@type:Office365` indicates the entry was analyzed by the Darktrace/Apps Microsoft 365 module.
- **@message**: A tab-separated summary of the fields contained within the log entry. Some fields are included in the detailed log entry but not in the message field. For `@type:notice`, a list of possible messages is available.



Making a Basic Query

A query consists of a field and a value, for example `@type:"dns"` where `@type` is the field and `"dns"` is the value. Simple searches for values like IP addresses or SaaS credentials, as well as wildcards (*) are also supported.



Multiple fields and multiple values can be connected together with comparators, for example: `@type:"dns" AND @fields.source_ip=10.1.23.2`. These additional conditions can be added manually, or by using the following buttons:

- The 〓 equals button adds a new AND condition to the current search.
- The ≠ does not equal button adds a new AND NOT condition to the current search.

As a search term is entered, Advanced Search will attempt to match it to saved searches (name and query) and the last 10 historic searches. Matching searches will appear in a list below the bar.



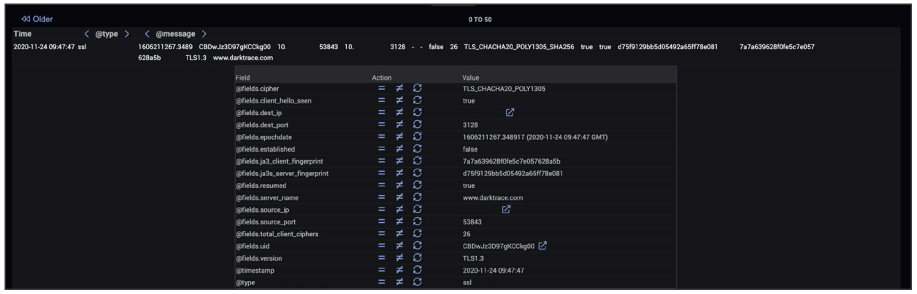
Saved searches are listed by the name given and indicated by the 🖫 save icon. Historic searches are indicated by the 🕘 history icon.

Please note, these features require local storage to be enabled. If local storage is disabled, historic and saved searches will be lost.

Set the timeframe to 15 minutes, enter `@type:"conn"` in the search bar and click **Search**. All connections processed by Darktrace in the last 15 minutes will be returned. For deployments with no network traffic, such as SaaS or Cloud only, substitute `conn` for a log type that can be seen in your environment such as `@type:office365` or `@type:amazon`.

Investigating an Entry

From the search results, click an entry to expand the event and show a table with a list of fields and their values.



In the left column are fields - distinct pieces of data or metadata that are produced during analysis. Each log contains universal fields such as the time the entry was created (`@timestamp`), a unique identifier (`@fields.uid`) and type specific fields which are only produced for that protocol or log type (`@fields.saas_actor`, `@fields.query_class`). Some fields, such as source and destination IP fields, are shared across connection log types and others are optional - appearing only where a value can be found in the processed traffic or event.

In the right column are the values that were derived for each fields during processing and analysis. Where a value has a link icon, the Threat Visualizer interface can be opened with a focus on that value. For example, focused on a specific connection (`@fields.uid`) or device (`@fields.dest_ip`) in the event log.



Each row of the table can also be used to refine the current search or start a new search based on the log entry.

- The **equals** button adds the field and its value as a new AND condition, narrowing the query to only those entries which match that value. For example, clicking this button in the `@fields.proto` row would add `AND @fields.proto:"udp"` to the query and limit the valid results to only those connections using UDP.

- The **does not equal** button adds the field and its value as a new AND NOT condition to the current search, effectively results which have that value in the specific field. For example, clicking this button in the `@fields.service` row would add `AND NOT @fields.service:"dns"` to the query and exclude any DNS queries from the results.

- The **circular arrows** button does not modify the current search and instead starts a new search using the criteria in the entry. When this button is clicked in a row, a new search is created with the Source IP and Destination IP of the entry, as well as the field and value in the row where it was selected. For example, the row `@fields.dest_port` would create a new query: `@fields.source_ip:"10.0.18.224" AND @fields.dest_ip:"192.168.72.4" AND @fields.dest_port:"53"`.

CHANGING THE LAYOUT AND FILTERING ADVANCED SEARCH RESULTS

Changing the Layout

On the left of the page is a list of fields which appear in the results; the list can be collapsed with the **hide** button. Each field listed can be added as a column in the results, making it easy to compare the value across all entries. This can be particularly useful to identify unusual log entries within a wide search criteria, or where the connection of interest is not yet known.



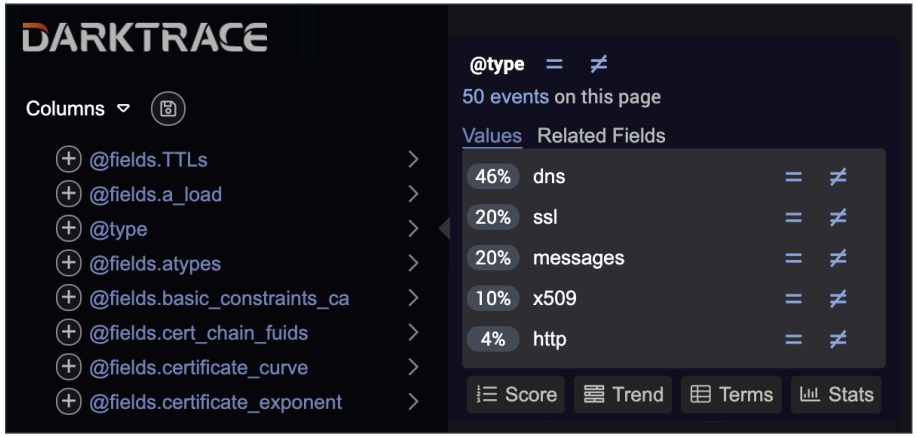
Adding, as an example, `@fields.dest_port` as a column allows an operator to quickly skim a list of connections for anomalously high ports or those unusual for the protocol under analysis.

- The **plus** icon adds the field as a column in the results. The default `@type` and `@message` columns are removed.
- The **minus** icon removes a column from the results layout.

Useful layouts can be saved by clicking the save icon and providing a name for the layout template (*Requires local storage*). Existing templates can be loaded by clicking "**Columns**" and selecting from the dropdown.

Field Summaries

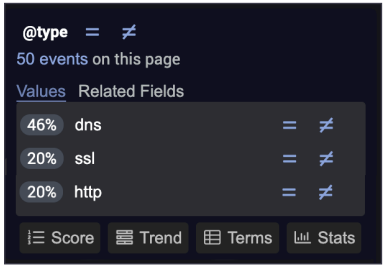
When the name of a field is clicked, a pop-up will appear with a breakdown of values seen and options to perform data aggregation queries.



Some fields are optional and will not return for all log entries - optional fields can be added as a requirement or an exclusion from the current search query:

- The **equals** button beside the name of the field will add a condition to the query that field must be present in each result. This takes the format `AND _exists_:"@fields.example"`.
- The **does not equal** button beside the name of the field will add a condition that the field must not be present. This takes the format `AND NOT _exists_:"@fields.example"`.

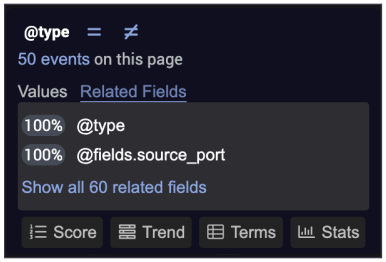
The number of results that include the field in question can also be highlighted by clicking "**# events on this page**".



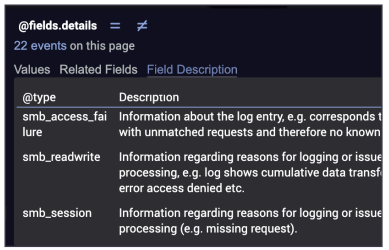
The **Values** tab shows a summary of the values across the 50 results currently shown on the page, sorted by frequency of occurrence. The summary allows an operator to get a sense of the logs returned and spot outlier results which may be unusual. Beside each result is the percentage of log entries that include that value in the specified field - clicking a value highlights the log entries it appears in.

- The **= equals** button beside each value adds the field and its value as a new AND condition, narrowing the query to only those entries which match that value. Where the value is '*Blank*', a `AND NOT _exists_:"@fields.example"` condition will be added instead.

- The **≠ does not equal** button adds the field and its value as a new AND NOT condition to the current search, effectively results which have that value in the specific field. Where the value is '*Blank*', a `AND _exists_:"@fields.example"` condition will be added instead.



The **Related Fields** tab shows other fields that frequently appear in the same log entries as the selected field. Related fields are ordered by the percentage frequency that they occur alongside the selected field out of all currently displayed log entries where the selected field appears.



The **Field Description** tab shows a brief description of the information the field contains. The same field may appear in multiple protocols but perform a different function - in this case, all possible type/description combinations are listed.

Data Queries

At the bottom of each field summary are four buttons: "**Score**", "**Trend**", "**Terms**" and "**Stats**". These buttons perform specific operations on the data returned by the search query and can be very useful to summarize large numbers of results.



- The **Score** action ranks values by frequency across the most recent 10,000 log entries that match the overall search query. Up to 100 values will be ranked in order of occurrence. In large instances, up to 100,000 results may be available for analysis.

- The **Trend** action provides an estimation of the change in frequency for values across the time period. A segment of matching entries at the start and end of the time period are analyzed and the change in value occurrence calculated.

- The **Terms** action also ranks values by frequency across log entries that match the query. This action is faster and performed against a larger number of results than the **Score** analysis, however it provides a less precise count. The number of entries analyzed where the field was missing will appear as a blank row above all other values.

- The **Stats** action is only available for fields with numeric values. This analysis produces statistical information about the field value across the timeframe including the average value, the sum, the count and the min/max limits. In this mode, each bar of the graph displays the mean field value during that time interval on hover.

BUILDING A COMPLEX ADVANCED SEARCH QUERY

Constructing a Search

Advanced Search allows for complex structured searches to be performed to retrieve detailed log information on connections. A wide starting search can be quickly focused using the buttons found across the interface to include and exclude values and fields.

Searches over a longer period will take longer to complete. Searches that take too long will be aborted to prevent negative impacts on the performance of other parts of the Darktrace system. Narrowing the scope of the search by time or by making a more specific query will help the search to return results more quickly.

Basic String Search

Words can be entered freeform and will be looked for across all values in the database. Enclosing a value in quotation marks will search for an exact string match and treat spaces as part of the search term.

Example:

```
"google"
```

A wildcard (*) can be used to represent any number of characters,

Examples:

```
*.darktrace.com
@fields.certificate_issuer:*.com AND @fields.certificate_subject:*.net
```

The latter example matches specific top-level domain extensions - regardless of the hostname - that are indicative of encryption certificates used by TOR nodes.

Searching by Field

To search for (or within) the value of a field, use the name of the field followed by a colon:

Examples:

```
@fields.host:www.darktrace.com
@fields.source_ip:10.10.*
```

For numeric values, greater-than (>) and less-than (<) operators can be used to identify values above or below a certain threshold.

Example:

```
@fields.orig_bytes:>0
```

The notation `TO` within square brackets designates a range of numbers.

Example:

```
@fields.dest_port:[6881 TO 6999]
```

The example above identifies a range of ports associated with the use of the BitTorrent file-sharing protocol.

Please note, this notation is not compatible with IP address ranges.

Combining Search Conditions

Conditions can be combined using the standard Boolean operators AND, NOT, OR. Most conditions added using the **= equals** and **≠ does not equal** buttons will take this format. Parentheses are also supported to group search conditions together

The default operator (if no operator is specified) between space-separated values is AND.

Examples:

```
@type:conn AND @fields.proto:"tcp"
@type:"azuread" AND (@fields.saas_metric:"Saas::Login" OR @fields.saas_
metric:"Saas::FailedLogin")
```

Boolean operators can also be used within parentheses to create shorter queries - the following two examples provide the same outcome.

Examples:

```
@fields.source_ip:10.10.4.3 OR @fields.source_ip:10.10.4.4 OR @fields.
source_ip:10.10.4.7
@fields.source_ip:(10.10.4.3 OR 10.10.4.4 OR 10.10.4.7)
```

Regular Expressions

Basic Regular Expressions (regex) can also be used for very specific searches.

A regex must be preceded and followed by a forward-slash character (`/`)

Example:

```
@fields.host:/w{3}\.darktrace\.co.+/
```

Regular expressions can also be combined with the previous syntax examples.

Complex Queries

Advanced Search supports the Lucene 4 Query syntax

TAXII CONFIG

Accessing TAXII Config

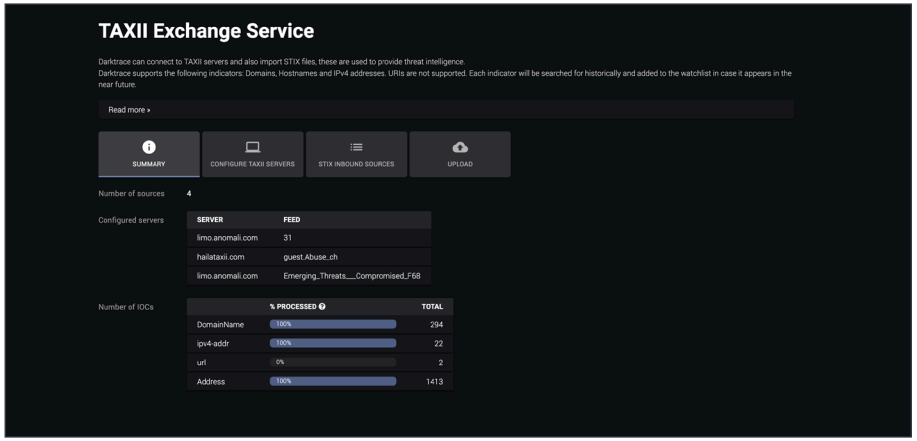
Darktrace can connect to TAXII servers and also import STIX XML files, these are used to provide threat intelligence. The **Threat Indicator** model will create model breaches when an observable derived from TAXII intelligence is seen.

Darktrace supports the following indicators: Domains, Hostnames and IPv4 addresses. URIs are not supported. Each indicator will be searched for historically and added to the watchlist in case it appears in the near future. TAXII services with multiple polling endpoints are not currently supported.

To access the TAXII Config page, navigate to the main menu and select **Intel** > **TAXII Config** from the available options. The page is also accessible from the **Modules** section of the System Config page, or from the sidebar of another page in the **Intel** category.

Summary

The summary page gives an overview of the TAXII servers and feeds subscribed to and the IoCs derived all sources.



Sources

The number of sources that have contributed IOCs - includes manual uploads and TAXII feeds.

Servers

Lists the hostnames of currently configured TAXII servers with details of the specific polled feed (collection).

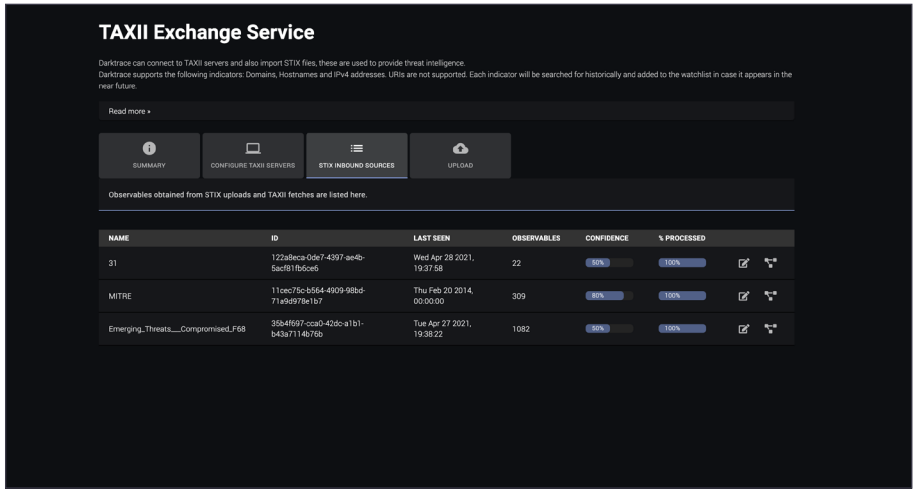
IOCs

The “total” value corresponds to the number of IoCs that have been collected from the TAXII feed.

IoCs that have been checked against the network and successfully added to the watchlist are considered “processed” - this is displayed by the “% Processed” value. If the *Valid Until* field is a date before the *Source Time* (the time when the IoC was received by Darktrace) the IoC will not be processed.

STIX Inbound Sources

The STIX inbound sources page displays details of the last 10,000 observables received by Darktrace. Each entry will describe the time of the last seen observable, the confidence and the processed observables from the feed. To view all ingested IoCs from each source, click the expand icon beside the source name.



Sources are derived from information contained within each STIX package - for STIX 1.2, this is typically the “Information Source” field. Information derived from STIX uploads and TAXII feeds can be viewed on the STIX Inbound Sources tab. For STIX packages uploaded without a defined “Information Source” field, the source will be ‘Ingested’. Where the STIX file originates from a defined feed and no source is provided, the source will default to the ***collection***. STIX files produced by Darktrace Inoculation will have the source ‘Darktrace’, or rarely ‘Inoculation’ in the case of older IoCs.

Confidence

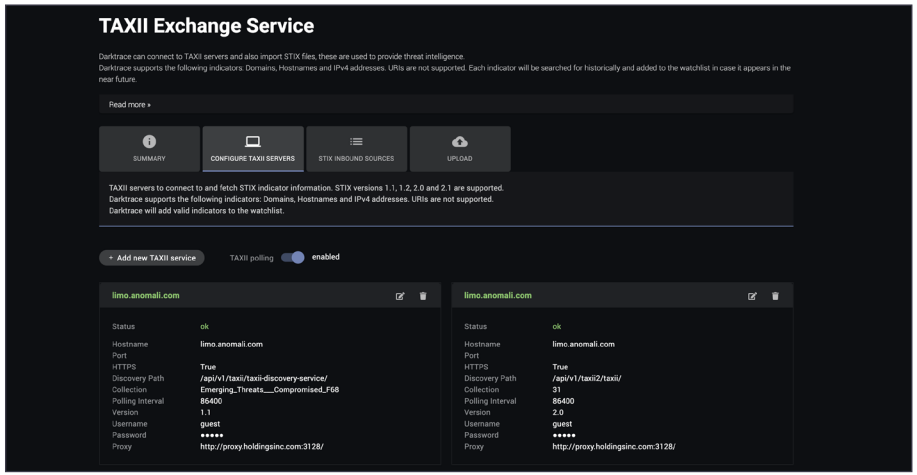
Each source will be assigned a default confidence of 50%. The source confidence can be modified by clicking the edit icon beside the source name.

Expiry

A default expiry of 2 weeks is applied to indicators, unless ***Valid Until*** is specified for the indicator in the information received by Darktrace. The default expiry can be modified to a value between 1 and 90 days, and set to override expiry information sent by the provider if desired. To do so, click the edit icon beside the source name.

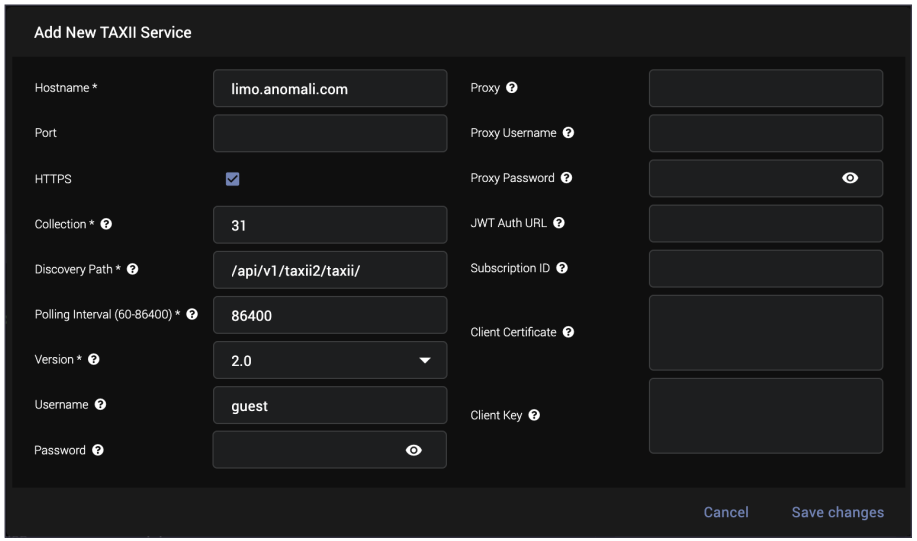
TAXII Services

TAXII threat intelligence feeds can be added on the **Configure TAXII Servers** tab.



- When a feed is initially added, Darktrace will request STIX packages that are a maximum of 2 weeks old. Subsequent queries will ask for STIX packages which have arrived on the TAXII server since the last known threat indicator received from that service.
- When polling is turned on, Darktrace will poll the TAXII server for new IoCs at the stated interval.
- Indicators which are downloaded in a STIX package from a TAXII service are then checked against data within Darktrace.
- Each IoC will be checked against historic connection data up to 2 weeks in the past (dependent on data stored).
- Each IoC will be added to the Watched Domains/IPs list with a default expiry date of two weeks in the future.
- If the STIX package includes a ‘valid until’ date this will be used as an expiry time for the IoC in the Watched Domains/IPs list.

Adding a TAXII Feed



1. Navigate to the **Configure TAXII Servers** tab and click "**Add New TAXII Service**".

2. Complete the required details of the TAXII feed Hostname, Collection, Discovery Path. Fields indicated with a * are mandatory.

3. Select a polling time in seconds to refresh the feed. There is a maximum setting of 86400 seconds (1 day) and a minimum of 60 seconds. This is the time between polling the individual TAXII services.

4. From the dropdown, select the STIX/TAXII version sent by the TAXII server.

5. Optionally configure a proxy and any authentication fields required by the third-party feed. Multiple methods are supported including Basic, Client Certificate and JWT authentication.

6. Finally, add the service by clicking **Save Changes**.

Ensure that TAXII polling is enabled or the feed will not be refreshed.

Uploading a STIX file

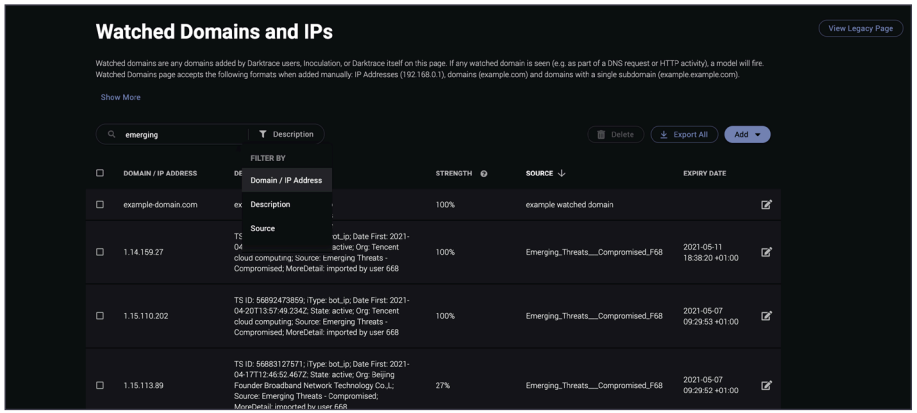
On the **Upload** tab, STIX files can be uploaded. Darktrace supports the following indicators: Domains, Hostnames and IPv4 addresses. URIs are not supported. Each indicator will be searched for historically and added to the watchlist in case it appears in the near future.

Navigate to the Upload tab and select the STIX-formatted XML or JSON file for upload. STIX versions 1.1.1, 1.2, 2.0 and 2.1 are supported. On successful upload, the number of observables and metadata fields derived from the file will be displayed. These can be reviewed on the **Inbound Sources** tab.

WATCHED DOMAINS

Watched Domains

Watched domains are endpoints of concern - whether due to known compromise, anomalous behavior or compliance - which will trigger a model breach and optional Darktrace RESPOND/Network response if observed in ingested traffic.



Domains can be added manually, through the API (`/intelfeed`), or via configured STIX/TAXII intelligence feeds. No default list exists for watched domains. The **Watched Domains** page accepts the following formats when added manually:

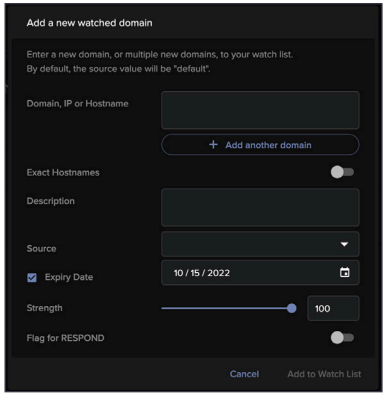
- IP addresses (192.168.0.1)
- domains (example.com) and domains with a single subdomain (example.example.com)
- IP address ranges (maximum subnet range is /8, previously /16)

An advanced configuration setting is available to allow a greater number of subdomains.

Access the **Watched Domains** page from the Threat Visualizer main menu under **Intel**.

In the **Watched Domains** workspace, all hosts, IP addresses, IP address ranges and endpoints on the watch list are listed and searchable. At the top-left corner of the workspace, you can use a free text search bar to locate a specific entry, and results can be filtered by clicking the filter icon. The default option is to search by **Domain/IP address**, which can be altered to **Description** and **Source** from the drop-down menu.

Adding New Entries



At the top-right corner of the **Trusted Domains** workspace, click **Add** to add new domains, hostname, IP addresses or IP address ranges to the Watched Domains list. If adding a range of IP addresses, the maximum subnet range is /8 (previously /16).

Choose between **Import CSV** or **Manual Input** from the drop-down menu.

- Choose **Import CSV** to upload an appropriately-formatted CSV format file.
- Choose **Manual Input** to add an IP address, domain or IP address range. The endpoint entered will be automatically detected as a **Domain, IP or Hostname**.

The **Exact Hostnames** flag indicates the value should be treated as an exact string only.

A free text **Description**, a selection of **Sources**, an **Expiry Date**, and a **Strength** from 1 to 100 can be set.

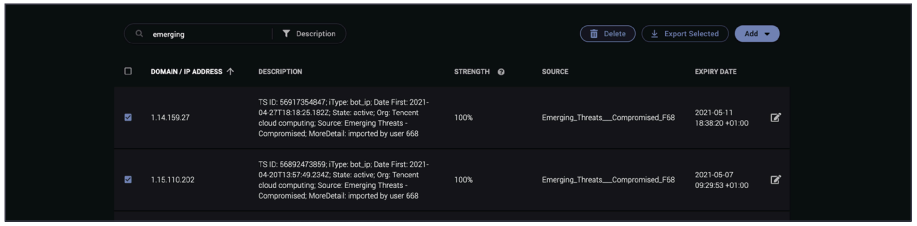
Finally, **Flag for RESPOND** will trigger an automatic Darktrace RESPOND/Network action if the entry is seen. It is not possible to trigger a RESPOND action against IP ranges, only discrete entries.

To add multiple endpoints concurrently, while at the **Manual Input** panel, click **+Add another domain** and a further free-text field will appear. This entry will share the parameters of the previous entry.

When all the information has been input, click **Add to Watched Domains**.

At the bottom of the screen, a notification will confirm the new addition.

Deleting and Exporting Endpoints



Delete an Endpoint

To delete a specific trusted endpoint from the list, click the square ■ icon on the left side of an entry.

The icon will become ticked ☑ and the **Delete** button on the top right will illuminate. Click **Delete** to remove the selected entry.

To delete multiple entries, select all applicable entries and click **Delete**.

To delete all entries, select the checkbox at the top of the list to select all entries, and then click **Delete**.

Export an Endpoint

To export a specific trusted endpoint, select the square ■ icon on the left side of an entry.

The icon will become ticked ☑ and the **Export Selected** button on the top right will illuminate. Click **Export Selected** to export the selected entry in CSV format.

To export multiple entries, select all applicable entries and then click **Export Selected**.

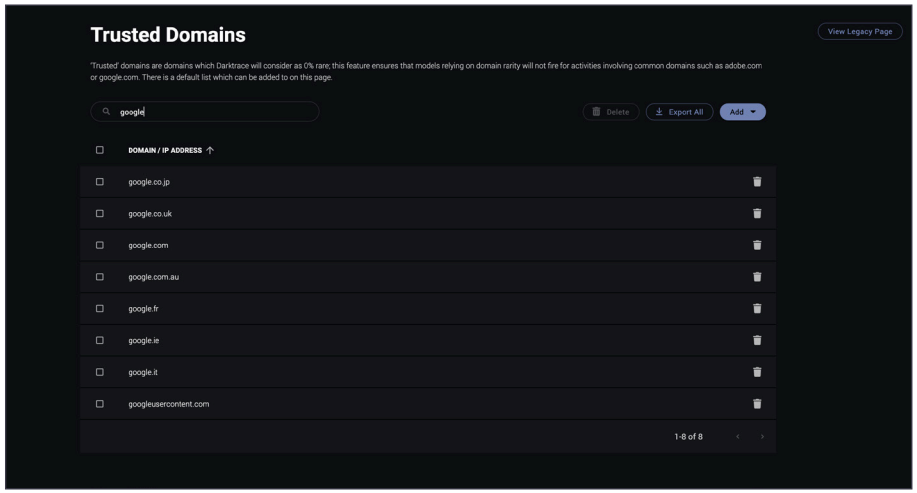
To export all entries, select the checkbox at the top of the list to select all entries, and then click **Export Selected**, or without selecting any entries click **Export All**.

Both actions will exporting all trusted domains as a CSV file.

Trusted Domains

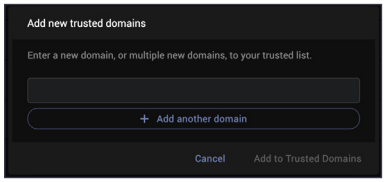
Trusted domains are domains which Darktrace considers as 0% rare. This feature ensures that models relying on domain rarity will not fire for activities involving common domains such as “google.com”. The list includes **Domains**, **IP Addresses** and **IP Address Ranges**. There is a default list, to which entries can be added on the **Trusted Domains** page.

Access the **Trusted Domains** page from the Threat Visualizer main menu under **Intel**.



In the **Trusted Domains** workspace, all endpoints that have been deemed as non-rare, or added by a user, are listed and searchable. At the top-left corner of the workspace, you can use a free text search bar to locate a specific entry.

Adding New Entries



Trusted endpoints are excluded from rarity calculations and models which use domain rarity - be careful when selecting new endpoints to be added.

At the top-right corner of the **Trusted Domains** workspace, click **Add** to add new domains or IP addresses.

Choose between **Import CSV** or **Manual Input** from the drop-down menu.

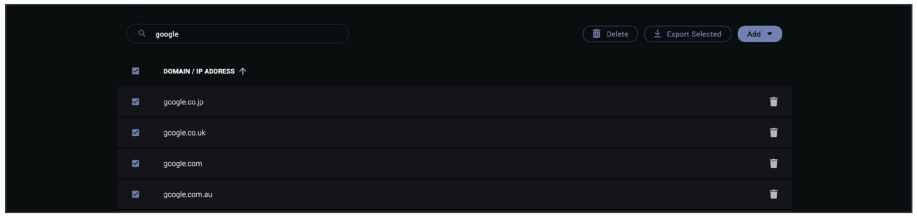
- Choose **Import CSV** to upload an appropriately-formatted CSV format file.
- Choose **Manual Input** to add an IP address, domain or IP address range as free text.

    To add multiple endpoints concurrently, while at the **Manual Input** panel, click **+Add another domain** and a further free-text field will appear.

    When all the information has been input, click **Add to Trusted Domains**.

At the bottom of the screen, a notification will confirm the new addition.

Deleting and Exporting Endpoints



Delete an Endpoint

To delete a specific trusted endpoint from the list, click the square ■ icon on the left side of an entry.

The icon will become ticked ☑ and the **Delete** button on the top right will illuminate. Click **Delete** to remove the selected entry.

To delete multiple entries, select all applicable entries and click **Delete**.

To delete all entries, select the checkbox at the top of the list to select all entries, and then click **Delete**.

Export an Endpoint

To export a specific trusted endpoint, select the square ■ icon on the left side of an entry.

The icon will become ticked ☑ and the **Export Selected** button on the top right will illuminate. Click **Export Selected** to export the selected entry in CSV format.

To export multiple entries, select all applicable entries and then click **Export Selected**.

To export all entries, select the checkbox at the top of the list to select all entries, and then click **Export Selected**, or without selecting any entries click **Export All**.

Both actions will exporting all trusted domains as a CSV file.

CHAPTER 8 - MODEL EDITOR

USING THE MODEL EDITOR

THE MODEL EDITOR 124
MODEL EDITOR OVERVIEW 125
FILTERING THE MODEL LIST 126
BULK ACTIONS IN THE MODEL EDITOR 127

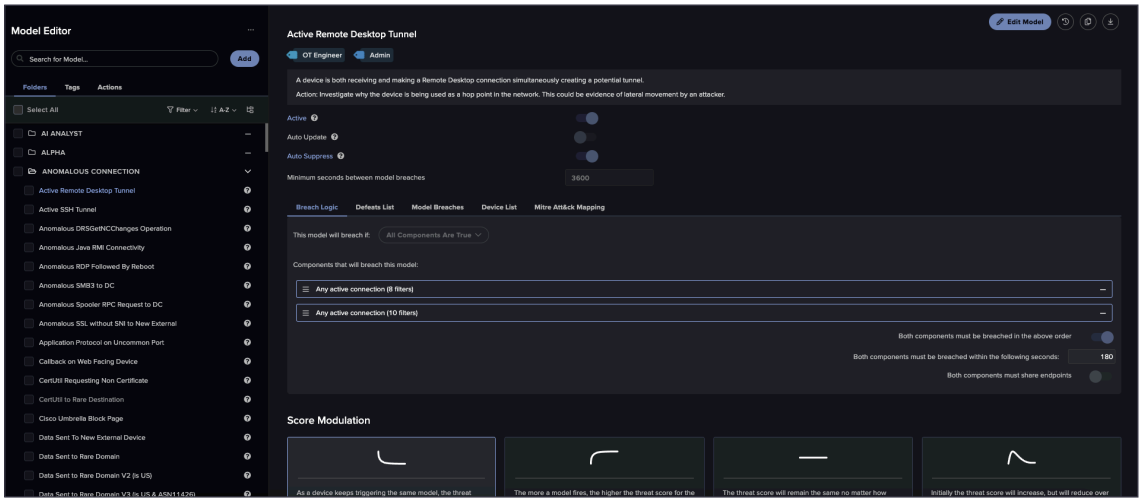
UNDERSTANDING A MODEL

LOOKING AT A MODEL 128
COMPONENTS AND FILTERS 129
MODEL BREACH LOGIC 130
MODEL CONDITIONS AND SCORING 131
MODEL ACTIONS 132
MODEL DEFEATS 134
MITRE ATT&CK MAPPINGS IN THE MODEL EDITOR 135
OTHER MODEL TABS 136

MAKING A NEW MODEL

CREATING A NEW MODEL 137
ADDING COMPONENTS TO A MODEL 138
ADDING FILTERS TO A MODEL COMPONENT 139
DEFINING MODEL COMPONENT BREACH CONDITIONS 141
TUNING MODEL SCORING 143
ADDING ACTIONS TO A NEW MODEL 143
TIPS FOR NEW MODEL MAKERS 145

THE MODEL EDITOR



Introduction to the Model Editor

Darktrace models are a series of logical statements and conditions which, if met, trigger an alert or action; models are primarily used to identify and alert on anomalous and potentially malicious behavior. The Threat Visualizer platform provides a curated set of default models as standard, designed and constantly tuned by the specialized Darktrace analyst team. Darktrace-maintained models are frequently updated in parallel with the evolving threat landscape.

The models framework leverages both the underlying 'pattern of life' detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules. Output from the complex anomaly framework is available in accessible, building block format and can be combined with simple conditions and logical expressions to create tailored activity detection.

Models can also be used to profile a device, assign tags, highlight misconfigurations or trigger autonomous responses from the Darktrace RESPOND framework. Creating custom models can be a straightforward and powerful way to integrate Darktrace into your existing security processes or replicate a SOC playbook.

Model Updates

Darktrace will periodically update the standard supplied models – either automatically for customers with Call-Home, or manually through software updates. If changes are made to a model logic outside of the defeats system, it will no longer be eligible for automatic updates.

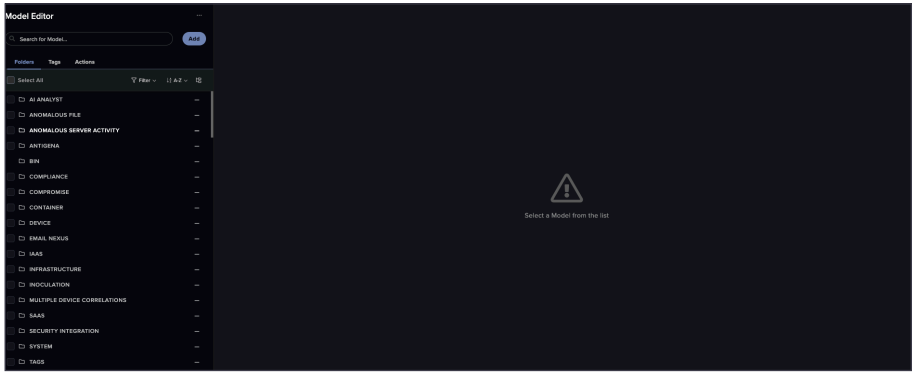
Please see Upgrading Darktrace Models or the System Administration guide for more information.

MODEL EDITOR OVERVIEW

The Model Editor

In the Threat Visualizer interface, the Model Editor is accessible from the main menu under **Models** > **Model Editor** or from any breach log with the " **Click to view model**" button. From the SaaS Console interface, the Model Editor is available from the sidebar ("⚠ Model Editor").

The editor is comprised of a list of current models (left) and a main workspace where model logic can be reviewed and edited. If the editor is opened from an event log - such as a model breach log - the main workspace will be focused on the model.



The 🏠 home button can be used to return to the main Threat Visualizer interface and the log out button to exit the Threat Visualizer platform. The ••• more options button in the top right provides access to the legacy model editor, bulk model import and bulk model export functionality.

The "Add" button opens a new, blank model for configuration. New models will be covered in more detail in Creating a New Model.

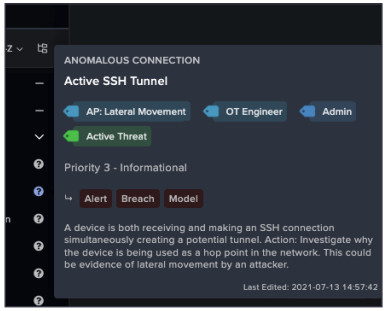
Model List

The default view is the "Folder" tab. Models are sorted into folders which describe the general purpose and detection area of the model. Models can also be tagged to identify the purpose or the potential attack phase for the activity it detects. Using tags can be particularly useful when categorizing custom models or when constructing filters to send email alerts to specific teams or users.

Clicking the name of a folder or the line beside a folder will expand to show the models and subfolders it contains. Clicking again will collapse the folder. The directory icon collapses all open folders.

Deactivated models in the model list are indicated by darker text.

Individual Models



Beside each model in the list is an ❓ info icon. Clicking this icon will open a summary tooltip containing:

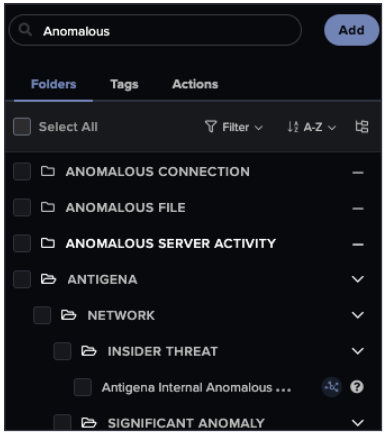
- The Folder(s) the model is in.
- The model title.
- Tags applied to the model, if applicable.
- The model priority.
- Actions in response to a breach.
- A Model description, where available.
- Last edited date.

The tooltip can be useful to quickly review a list of models in a folder before selecting one to duplicate for editing purposes.

Click on a model to open a detailed overview in the right-hand pane.

FILTERING THE MODEL LIST

Filtering the Model List

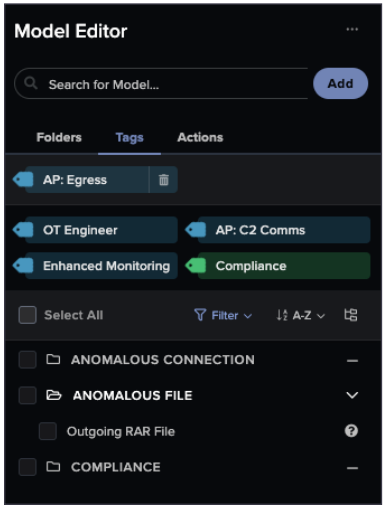


There are multiple ways to filter the model list to display only relevant models. The list of models can be sorted alphabetically, by priority, by creation date or by last edited date. Searching for a keyword or model title will filter the list of models and folders to only those with a title or description that includes the search terms.

Use the "Filter" option to show/hide models based upon their active state (enabled/disabled), whether they have any model defeats applied, and whether they are in the "Compliance" threat behavior category.

Bulk actions can be performed on multiple models at once; this is covered in Bulk Actions in the Model Editor.

Model Tags Tab



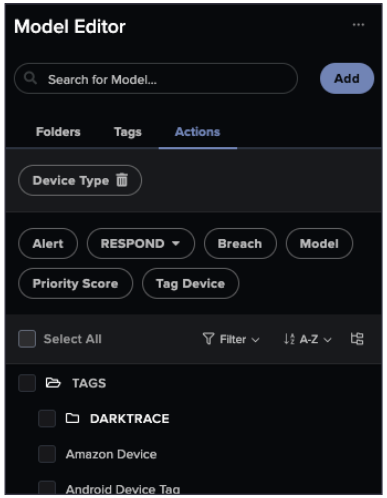
The tag view is a way to filter the list of models to those with specific tags applied.

When the view is set to tags, all tags applied to models are shown in the sidebar. Clicking on a tag will list all models it is applied to.

Above the model list is a selection of suggested tags to filter by - clicking these suggested tags will add them to the filter.

Tag filters can be removed with the trash icon.

Model Actions Tab



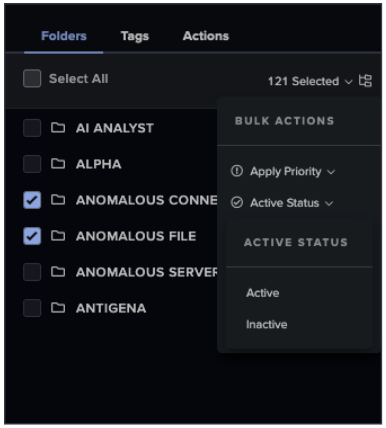
Model actions are system actions which can be triggered in response to a model breaching. Models can also be sorted by the action(s) they perform. This can be a useful way to identify all Darktrace RESPOND models ("Antigena"), or all models which apply tags.

When the view is set to action, all model actions are shown above the list of models. Clicking on an action will list all models that perform it. Additional actions can be added, or filters can be removed with the trash icon.

For Darktrace RESPOND actions, the ***state*** of the RESPOND action - e.g., "Force Human Confirmation" - can be further filtered by. This is a quick way to identify models which are permitted to take autonomous action at all times, or those forced to always require permission.

BULK ACTIONS IN THE MODEL EDITOR

Certain settings can be changed across multiple models at the same time using bulk actions. When one or more model is selected, a dropdown will appear where the sort option was previously.



- Clicking the ✔ tick icon that appears to the left of the model name will select that model.
- Clicking the ✔ tick icon that appears to the left of a folder name will select all models contained within.
- Clicking the ● circle icon at the top left of the list below 'Folders' will select all models.

Bulk Actions

The following options are available for bulk edit:

EDITABLE SETTINGS	DESCRIPTION
Model Priority	Model priority is a modifier applied to the final score assigned to the model - higher priority models will have a high model breach score. It also controls which threat behavior category a model is applied to (0-3 Compliance, 4 Suspicious, 5 Critical)
Active status	Whether the model is turned on or not. De-activating a model can be a useful way of preserving the logic for future tuning without removing it entirely.
Auto Updates	Not applicable to custom models. Whether the model can be automatically updated (if permitted by other factors).
Auto Suppression	Controls whether the model should be automatically suppressed by the modeling engine if triggered an excessive amount of times in a short period.
RESPOND State	The models "RESPOND" state impacts whether autonomous RESPOND actions can be taken as a result of the model triggering. Please refer to Model Actions for more information.
Model Tags	Add one or more tags to the model(s). Tags can be used as a filter within the model editor or for external alerts.
Compliance	When "true", includes the model in the "Compliance" threat behavior category. This category takes precedence over the priority setting. Refer to Filtering Alerts with Behavior Categories for more information.

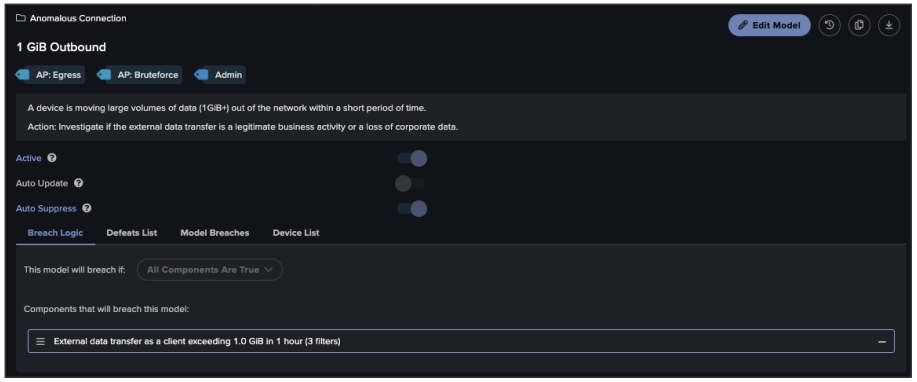
LOOKING AT A MODEL

For this example, we will look at a version of the "Anomalous Connection / 1 GiB Outbound" model. Clicking on the name of any model will open the model overview on the right.

Please note, as Darktrace models are tuned and updated with high frequency by a specialized analyst team, the versions of the model here may not correspond with the versions in your environment. These models are only used for example purposes and the concepts apply to any model - custom or default - which you may encounter.

Main Overview

At the top of the overview is the folder, the model name and any tags applied. Both the top-level folder and tags can be used as filters in the threat tray. Underneath is a description of the behavior the model is targeting and a potential action to take in response to a model breach.

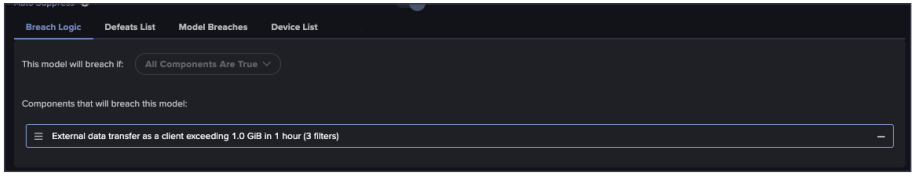


There are three toggles:

- **Active**: Determines whether the model is active or not. Allows the user to turn a model off - rather than deleting it and having to recreate it at a later date - so it will not breach during that time.

- **Auto Update**: Where models are set to update automatically in the global config setting, individual models can have auto-update disabled so that updates must still be manually approved.

- **Auto Suppress**: Determines whether this model will be automatically suppressed if it breaches repeatedly in a short timeframe for the same device and with the same set of breach conditions. If a repeatedly breaching model is suppressed, from that point onward it will generate at most one breach per week for any given device and the same breach conditions. Recommended for most models.



Breach Logic Tab

The first tab of the model overview contains a summary of the model logic. Models are constructed from components - a series of logical statements evaluated against a connection, event or device.

A model may have more than one component. The dropdown indicates the relationship between those components that is required to produce a model breach. There are two types of component relationship: "*All Components Are True*" and "*A Target Score is Reached*".

All Components Are True

In this mode, all components in the model must evaluate as "true" for the model to breach. This is known as a 'checklist' model.

A Target Score is Reached

In this mode, components are assigned a positive or negative weighting and a target score is set for the whole model. This is known as a 'weighted' model. If a component evaluates as "true", it then contributes the 'points' it has towards the target score. A negative weighting can be used as a control. If enough components evaluate as "true" to meet the target score in the timeframe specified, the model will breach regardless of the status of any remaining components.

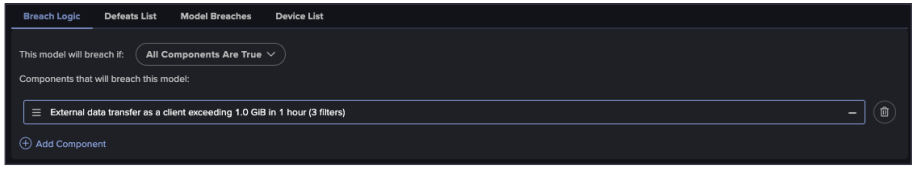
COMPONENTS AND FILTERS

Components

Components are made up of metrics and filters. Each component looks at exactly one metric, filtered with zero or more filters. The metric is the main behavior, event, or condition that defines the component, and filters are restrictions or stipulations that matching events or activity must meet for the whole component to be satisfied. The list of filters changes according to the selected metric - only filters that are relevant to a metric are available.

Metrics can cover many different types of activity or event, for example:

- Specific activity performed by a device, such as number of connections or volume of data transfer.
- Distinct events, such as a failed login to a SaaS platform or the use of a new credential.
- Activity deemed unusual in comparison to the 'pattern of life' and other outputs from Darktrace Cyber AI.
- A reference to components in another model or Darktrace RESPOND (*formerly Antigena*) actions that have taken place.

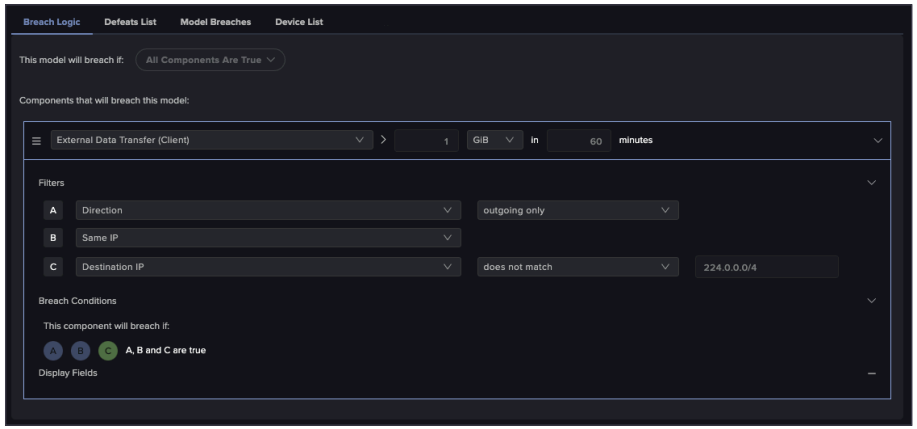


In this example, the metric is "External Data Transfer (Client)" and the activity must meet the condition "> 1 GiB" within the timeframe of 1 hour. The component is therefore targeting large data transfers from standard, non-server devices to an external location. Each component is summarized in this way to give a quick understanding of the type of activity that it is looking for.

Filters

The number of filters applied to the metric is also listed - click the component line to expand the filters.

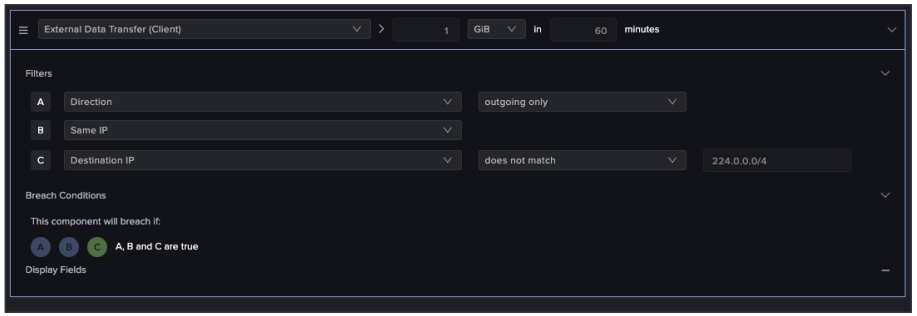
Here, there are three filters applied to the External Data Transfer: the data transfer must be outbound, it must be to the same IP and it must not be to an address in the Multicast network range.



MODEL BREACH LOGIC

Breach Conditions

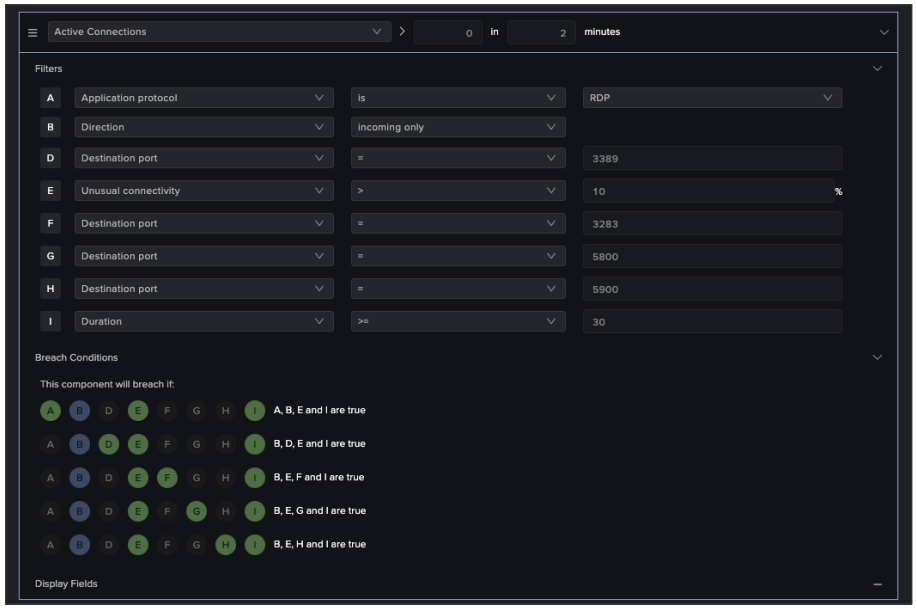
Each filter is assigned an alphabetical reference which is used to define the breach conditions. Breach conditions for a component are a list of possible filter combinations that should trigger a breach. In this example model, the relationship between the three filters is very simple - all three filter conditions must be met for the component to evaluate as true - and therefore the breach conditions are represented as “A & D & E”.



Frequently a component will have many filters which cover different scenarios or act as controls, where the filters are not intended to all be used at once. In this case, multiple lines of logic will be defined in the Breach Conditions.

For example, the “Active Remote Desktop Tunnel” model has “Active Connections” components where the filters correspond to a selection of ports. Each filter in this case reflects a slightly different connection scenario and consequently, multiple combinations of these filters can trigger a breach.

Where a component has multiple lines in the breach conditions, these should be treated as an ‘OR’ relationship. In this example above, the logic looks like “A & B & E & I ***OR*** B & D & E & I ***OR*** B & E & H & I”. For a component to evaluate as true overall, one of these lines of breach condition logic must be satisfied.



Display Fields

Display fields are extra filters which do not have an impact on the component logic. Instead, they are used to include additional information in the model breach alert or breach log entry when a component breaches. For example, adding a display field of “Source Port” will result in the port number that the breaching connection originated from being included in the log.

Display fields make it easier to triage model breaches in the threat tray, as you can quickly see the interesting features of the activity Darktrace has detected as anomalous.



MODEL CONDITIONS AND SCORING

Component Settings

Where a model has more than one component, component settings can be defined which limit the timeframe and order in which components must be triggered for a model breach. Component settings available differ between the two model types: checklist and weighted models.

Checklist Model



All components must be breached in the above order

When true, components must be triggered in order from top to bottom. Components can be rearranged by clicking and dragging.

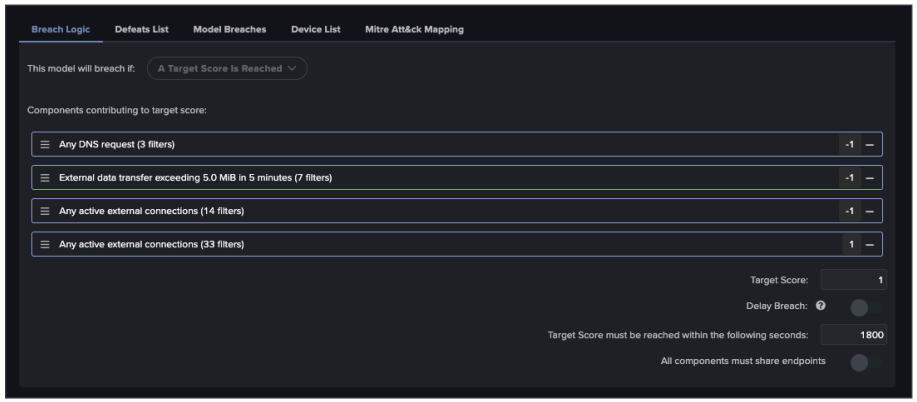
All components must be breached within the following seconds

A limit on the time within which all components must trigger for the model to breach overall.

All components must share endpoints

Requires components to be triggered by the same destination endpoint for the overall model to breach. This option will not be displayed if there is only one component in the model.

Weighted Model



Target Score

The total score that must be reached when adding the points assigned to triggered components to create a model breach.

Delay Breach

If the model has a negative component, this option will appear. If turned on, adds the "Delay Breach By (seconds)" field to specify a delay to allow activity to happen that meets the negative component's criteria.

Delay Breach By (seconds)

Minimum delay in seconds after a positive-scoring component has fired before the overall model score is calculated. This option is useful to prevent a model breach where the negative component activity may occur shortly after.

Target Score must be reached within the following seconds

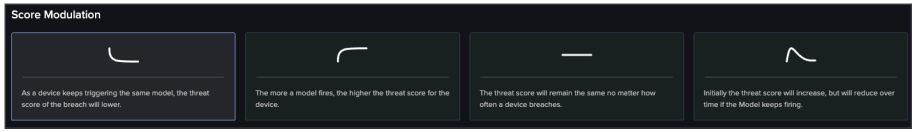
A limit on the time within which enough components must trigger to reach the target score, so that the model may breach overall.

All components must share endpoints

Requires components to be triggered by the same destination endpoint for the overall model to breach. This option will not be displayed if there is only one component in the model.

Score Modulation

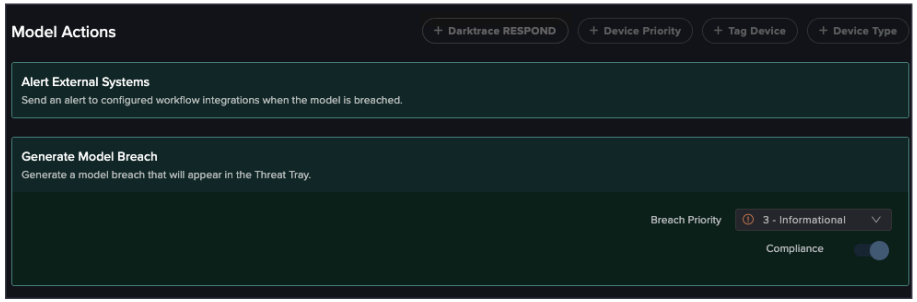
Determines the modulation curve for this model – when the same device repeatedly breaches a model, the score of subsequent breaches will be adjusted. For some behavior types, repetition indicates that the activity is likely less worthy of attention from the security team.



For other behavior types repetition makes the activity potentially more serious. Selecting a modulation curve will determine how the scoring of repeated breaches for a given device should be adjusted.

MODEL ACTIONS

Model actions are system actions which can be triggered in response to a model breaching.



Generate Model Breach (+ Breach)

The most basic response is breach, which triggers the system to generate a model breach alert in the Threat Visualizer threat tray. In the majority of cases, triggering a threat tray entry is desirable to alert Threat Visualizer users to the activity detected. Where a model is used for identification purposes - for example to tag a device that exhibits specific, non-malicious behavior - producing a threat tray entry is not desirable and the breach action can be removed.

The model's priority will affect the strength with which it breaches, so that if particular types of behaviors are of greater interest to you, these behaviors will register as more strongly relevant and will be more obvious in the threat tray than if the priority was lower. Breach priority can also be used to control alerts to external systems where a model is also set to alert.

Alert External Systems (+ Alert)

The next step from generating a threat tray entry is the alert action. A model with this action will trigger the system to send an alert to all configured alerting outputs. Only models with this action will trigger external alerts.

Darktrace RESPOND (+ Darktrace RESPOND)

Darktrace RESPOND (***formerly "Antigena"***) is the framework that takes autonomous actions in response to potentially malicious activity. Default Darktrace RESPOND/Network actions are triggered with meta-models which look for the presence of one or more Darktrace RESPOND (Antigena) tags on a model breaching device. Darktrace RESPOND capability can also be added directly to a model by adding this 'Antigena' action section. Inhibitors - response actions - can be selected for each Darktrace RESPOND type enabled on the deployment. Inhibitors can be automatic, where Darktrace RESPOND selects the most appropriate action at the time, or pre-defined from the list of options. By default, inhibitors will be active for one hour (3600 seconds).

RESPOND State controls when the action is taken without human confirmation. If set to "Force human confirmation", the model will always ask for confirmation before taking an action - regardless of the global state. In "Take autonomous action" mode, a response will be taken automatically unless confirmation for actions is required globally. If "Force autonomous action" is chosen, the model will override the global schedule and always take action without confirmation.

Optionally, a threshold can be defined which restricts Darktrace RESPOND response to breaches of the model that exceed the limit. For models where the threat score reduces over time, it can be desirable to take autonomous actions against the first one or two model breaches but not after that point. The Darktrace RESPOND Threshold therefore sets the lowest breach score above which Darktrace RESPOND will take the specified response.

Assign Priority Score (+ Priority Score)

Device priority is a scoring system which marks device importance - higher priority devices will produce higher model breach scores. Device priority can be increased and decreased as a model breach action. For example, a model which profiles devices may increase the priority of a device if it displays activity consistent with a server. Similarly, a model breach can be used to raise the priority of a device so that future alerts - which may indicate a developing security event - are more prominent and higher scoring.

Tag Device (+ Tag Device)

A model can add a tag to a device on breach. This can be particularly useful where models are used to profile a device or to mark a device as high risk due to its behavior. Tags can be given an optional expiry date; a device can therefore be added to a 'risk pool' for a set amount of time in response to a model breach, and automatically leave it after a set period such as two weeks.

Device Type (+ Device Type)

Darktrace detects device types automatically and assigns a most likely type based upon device activity and traffic analysis. Device types can be overridden on the Device Admin page or via the API. Models can also reassign device types as an action. Instead of setting device types manually to match asset registers, a model can be configured to match a hostname convention and assign device types accordingly. The additional benefit of a model is that new devices will be assigned automatically as they appear and trigger a breach, reducing the amount of repeated manual configuration required.

Model

The Model action is a default action applied to all models.

MODEL DEFEATS

What Are Defeats?

Defeats are different from negative filters. Filters in the model logic limit the connections, events or activity eligible to trigger a breach down to only those that are relevant. These are integral parts of the model logic. Defeats are specific conditions which prevent a model from breaching - they are distinct from the logic. Because of this distinction, defeats are specific to each environment and do not prevent the main model logic from being updated. Adding a defeat to a Darktrace-maintained model **does not make it ineligible for model bundle updates**.

Defeat conditions offer more specificity than device-based exclusions or restrictions and can also target a large range of devices by characteristics like device type or hostname. Defeats can be a useful tool to tune existing models to suit your environment as your deployment develops, whilst also leveraging the expertise of the Darktrace analyst team as they update default models to match the ever-changing threat landscape.

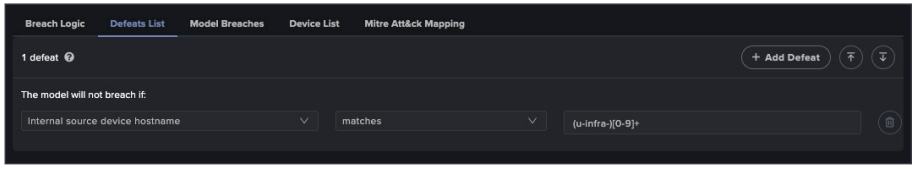
Defeat Logic

In contrast to components, no defeats can evaluate as true for the model to breach - each defeat should be considered as a series of "AND NOT" logic lines against the whole model. Where the model is a checklist, the overall logic can be thought of like:

[component 1] AND [component 2] AND NOT [defeat 1] AND NOT [defeat 2]

For example, an organization has an update infrastructure that regularly sends large files to external locations - all internal devices in that infrastructure use a hostname convention "u-infra-##". These devices can be excluded from the scope of the model with the defeat:

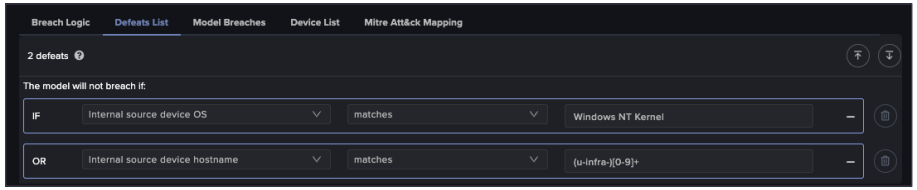
[Internal source device hostname] [matches regular expression] [(u-infra-)[0-9]+]



In this example, the model is prevented from breaching if the device that sends data outbound has a hostname that matches the regular expression `(u-infra-)[0-9]+`.When the comparator is positive - e.g. "matches", "is", "contains" - the defeat essentially excludes a subset of devices or activity from triggering a breach.

Conditional Defeats

From Darktrace Threat Visualizer 6, defeats can be ***conditional*** - one defeat can contain multiple "AND" conditions that limit the scope. Where the model is a checklist, the overall logic can be thought of like:

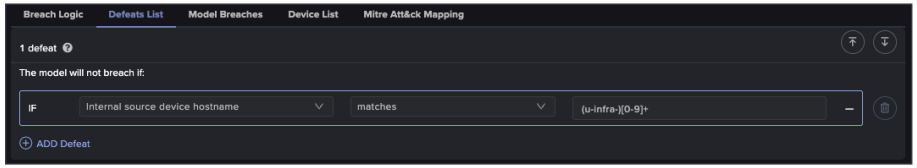


[component 1 must be true] AND [component 2 must be true] AND NOT [IF ([A]) is true] AND NOT [IF ([A] AND [B]) is true]

Continuing the example above, we can restrict the defeat to just those which match the regular expression and are detected as **server** device types.

IF ([Internal source device hostname] [matches regular expression] [(u-infra-)[0-9]+]) AND ([Internal source device type] [is] [Server - Any])

To add a condition, click the — line icon on the right of the defeat line.



Negative Defeats

Defeats can also be used to limit the scope of a model - this is a very powerful tool to tailor models to your environment but must be carefully executed to avoid over-limiting. For example, adding the following defeat would restrict the model to **only breach when the device has the File Server tag**:

[Tagged internal source] [does not have tag] [File Server]

If a defeat condition evaluates as true, it prevents the model from breaching entirely. In the above case, the defeat will evaluate as true for all devices not tagged with "File Server" - essentially limiting the model to **only** breach when the device is tagged.

Defeats with negative comparators can therefore constrain a model to very specific breach criteria.

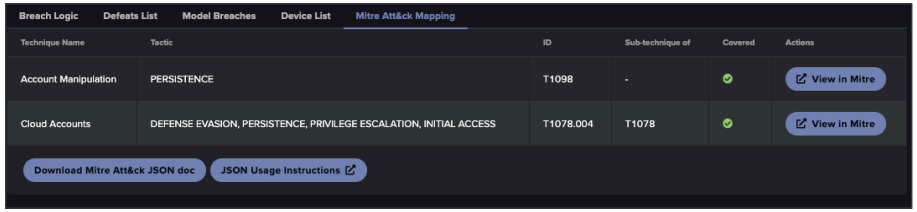
Bulk Editing Defeats

Defeats can be downloaded and edited as a .csv file, then re-uploaded to the model. This can be a quick way to add the same defeat conditions across multiple models.

MITRE ATT&CK MAPPINGS

Models curated and maintained by the Darktrace analyst team are mapped to the MITRE ATT&CK Framework, where applicable. This mapping is a valuable tool to understand coverage and for teams with internal playbooks for how to address each technique.

For each model, the MITRE ATT&CK Mapping tab displays the tactics and techniques the model covers are displayed on this tab.



For each technique, click " View in MITRE" to open detailed information about the technique on the MITRE website.

A JSON file of all mappings can be downloaded from any model. This mapping file can be uploaded to the MITRE ATT&CK navigator to browse the mapped techniques. If you are unsure how to use this file, refer to the instructions provided in "JSON Usage Instructions".

OTHER MODEL TABS

Model Breaches

The **Model Breaches** tab displays a graph of breaches for the selected model over the given timeframe - by default, two weeks are displayed. The time window covered by the graph can be adjusted in the bottom right and can return up one year of model breach data (where available). Acknowledged breaches can be displayed or hidden with the "**Show Acknowledged Breaches**" toggle.



The graph displays each model breach as a node. A tooltip with the device, breach score and exact time of breach appears when hovering over a node. The average breach score across the timeframe is displayed as a red line. Grey, vertical lines indicate a that the model was edited, and the logic potentially altered.

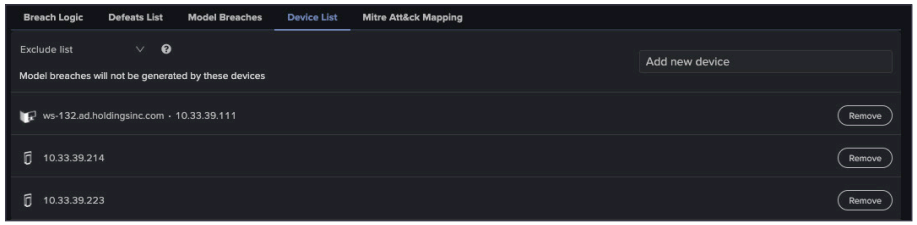
Under the graph is a list of devices that breached the model in the given timeframe, where each row corresponds to a specific device. The type of devices breaching the model are also summarized in the top left. Each device is identified by the hostname, IP and/or MAC address where available. On the right, the number of breaches and the time of the last breach are displayed. Clicking the number of breaches opens a new window or tab with a Breach Log for each event.

The ignore device button is located beside the number of breaches - ignoring a device is in addition to the defeats system. When a device is ignored, it will never trigger a model breach for the model in question, even if it performs behavior that matches the model criteria in future. If the icon is greyed-out, the device is already ignored and must be added to or removed from the device list, depending on mode, to become eligible to create breaches again.

To download the list of devices that have triggered the model over the time period, click the "Export Model" button beside the time range.

Device List

The device list displays the devices current eligible or ineligible to generate model breaches. A model can either be applied to all relevant devices - those not already removed by defeats or model logic conditions - or a select subset of devices.



- In the first mode, ignoring any device will place it on a list of devices ineligible for breaches. This is an **"Exclude"** list.
- In the second mode, all devices apart from those explicitly listed will be treated as ignored. This is an "**Include**" list.

Only one mode can be selected.

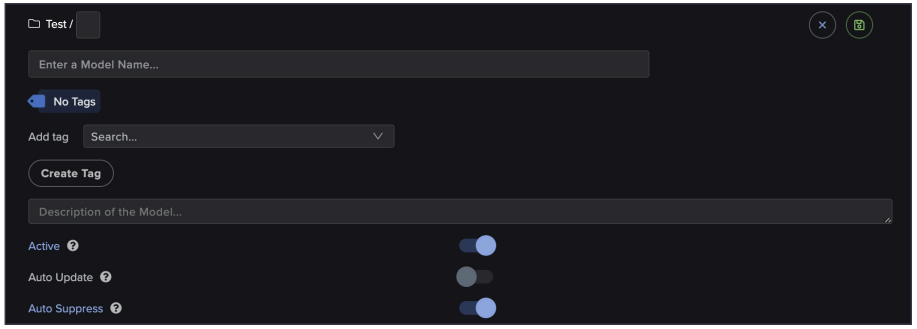
When the model is edited, devices can be added via the search bar and remove with the "**Remove**" button.

CREATING A NEW MODEL

Creating a Custom Model

To create a new model, click the “Add” button beside the search bar.

Models are sorted into folders to categorize their purpose or by the type of activity they are looking for. If a model contained within a folder is open when the “Add” button is clicked, the new model will be automatically placed within that folder. Models can also be dragged-and-dropped into folders using the model list after creation.



Basic Information

Model Name

First, select a name for the model. The model name can also be used to control alerts to external outputs so should be clear and concise.

Description

Adding a description will help other users of the Threat Visualizer understand the purpose of the model when they observe it in the main Threat Tray. A description is optional but highly recommended.

Active

Now set whether the model should be active when it is created - only active models can breach. Toggling a model to inactive allows the user to turn a model off for a time (rather than deleting it and having to recreate it at a later date).

Auto update

This setting determines whether a model will be automatically updated or not when a new version is available from Darktrace. This is not applicable for custom models.

ADDING COMPONENTS TO A MODEL

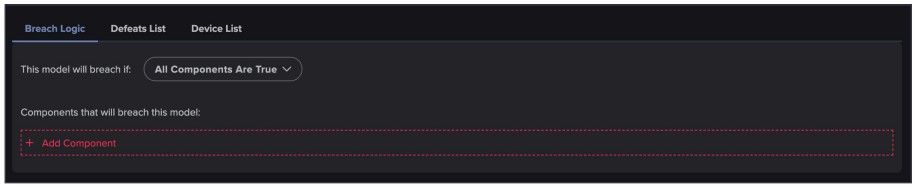
Breach Logic

The breach logic is the most important part of the model. The components and filters that make up the breach logic define the activity the model is looking for. Each model has one or more components - types of activity to detect - which are controlled with filters to limit the eligible connections or activity that could trigger a model breach.

There are two types of model: checklist and weighted. These were discussed in more detail in Looking at a Model. Simply, the type of model defines the relationship between those components that is required to produce a model breach:

- Checklist = "*All Components Are True*"
- Weighted = "*A Target Score is Reached*".

Select the type of model before proceeding.



Building Components

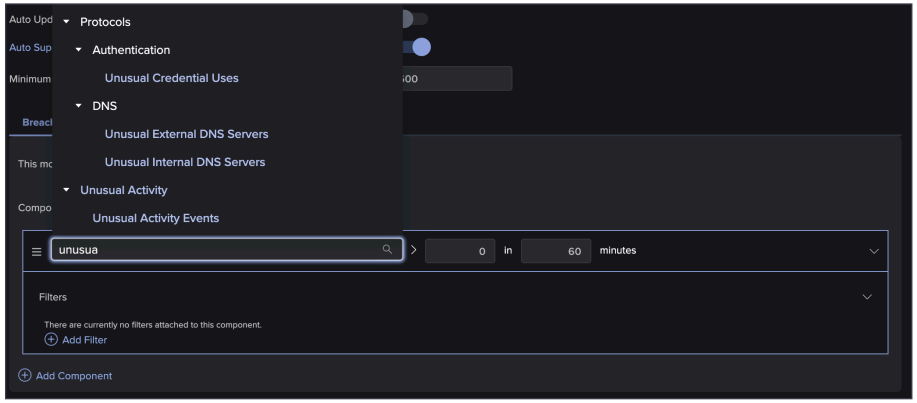
Components are created from metrics - discrete events or activities that are observed by the Threat Visualizer. To start building the model, click "+ Add Component" to add the first component.

Select the Metric the Model should include. The available Metrics are listed in Model Editor Metrics but may differ depending on your environment and any additional Darktrace modules and integrations providing data.

At the right of the '>' symbol, enter the number of times the filtered metric should be seen before the component will trigger. A zero in this field means any number of times. The default component identifies at least one connection in 60 minutes.

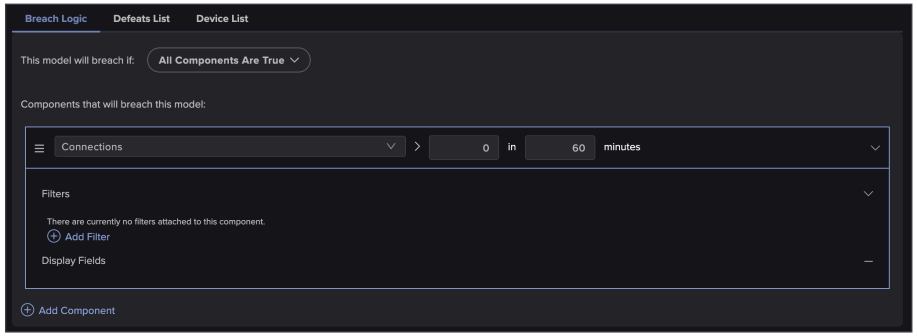
The frequency and the time period can be altered to identify activity over a shorter or longer period. Other metrics can be selected for evaluation - click "Connections" to open a dropdown of metrics categorized by their type or by the models they are relevant to.

Either select a new metric from the dropdown or proceed with the default "Connections". *The relationship between components and metrics is covered in more detail in Components and Filters.*



Filters

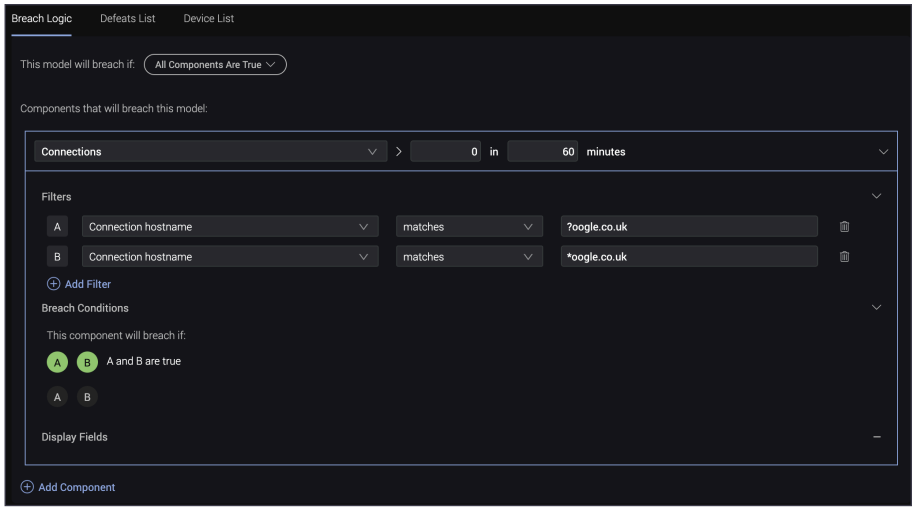
It is not desirable to alert on *every* connection that occurs, just some interesting connections; filters therefore limit the number of events that can trigger the component. Some metrics are very precise and do not need qualifying with filters. We will now proceed to add filters to the new model.



ADDING FILTERS TO A MODEL COMPONENT

Model Filters

Each metric has a list of applicable filters that can be used to limit its scope. For example, connection metrics can be filtered by hostname rarity, ports and connection details such as length and data transfer. SaaS activity, representing discrete events, can be filtered by the user involved, the SaaS platform, the geographic location of the action, etc.



Filters are combined with ***comparators*** and ***values*** to create breach conditions.

Filter Comparators

Matches / does not match

Attempts to match the string provided against the value in the field. If * is used, performs a simple wildcard match of any length. If ? is used, performs a wildcard match against a single character.

Example Breach Condition

```
[Connection Hostname] [contains] [*oogle.co.uk]
```

Example Match:

```
mail.google.co.uk
google.co.uk
```

Example Breach Condition

```
[Connection Hostname] [contains] [?oogle.co.uk]
```

Example Match:

```
google.co.uk
```

Contains / does not contain

Tries to locate the exact string provided anywhere within the filtered field.

Example Breach Condition

```
[Connection Hostname] [contains] [`darktrace`]
```

Example Match:

```
darktrace.com
```

Matches regular expression / does not match regular expression

Attempts to match a value in the field against a regular expression. Does not support lookups.

The comparator is case-sensitive but can be made insensitive by surrounding the expression with forward-slash and appending an ‘i’.

Example Breach Condition

```
[Message] [matches regular expression]
[(Anomalous|Compromise|Device|Unusual Activity|User|Infrastructure).*]
```

Example Match:

```
Device / New Failed External Connections
```

Is longer than / is shorter than

This comparator accepts a number, representing the number of bytes.

Example Breach Condition

```
[DNS Hostname Lookup] [is longer than] [0]
```

Is / Is not

Allows for values to be selected from a list of predefined values. Also includes booleans.

Example Breach Conditions

```
[Day of the Week] [is] [Sunday]
[Proxied Connection] [is] [False]
```

Has tag / does not have tag

For tagged internal sources or destinations, allows the tag to be selected from a list.

Example Breach Conditions

```
[Tagged Internal Source] [does not have tag] [Admin]
```

Numeric Comparators

The following numeric comparators are supported. Depending on the filter, not all comparators may be available.

- < (Less than)
- <= (Less than or equal to)
- = (Equal to)
- != (Not equal to)
- >= (Greater than or equal to)
- > (Greater than)

Example Breach Conditions

```
[Rare external IP] [>=] [90]
```

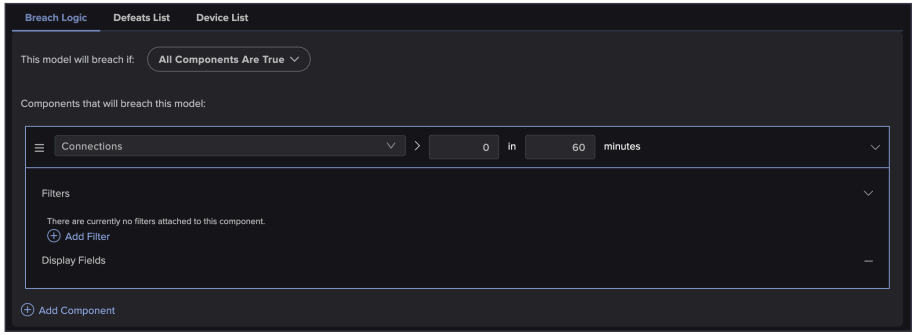
Other Filters

Some very specific filters do not have comparators or values. For example, the “Direction” filter offers the values “Incoming only” and “Outgoing only” instead of defining a comparator. The “Same IP” and “Unique IPs” filters do not have any comparators or values at all. The majority of filters behave exactly as described above. Where filters differ, the logic should still be clear.

DEFINING MODEL COMPONENT BREACH CONDITIONS

Filter Logic

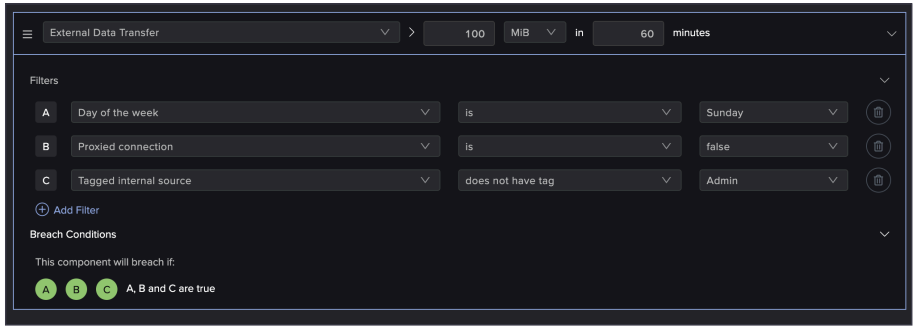
When more than one filter is defined for a component, the **Breach Conditions** section will appear - this is where the relationship between filters is defined.



Breach conditions connect filters together in logical sequences. Each filter has an alphabetical reference which corresponds to a 'bubble' in the conditions. When a bubble is clicked on, it becomes highlighted and is now a required condition for the component to breach.

Multiple required components should be thought of in an 'AND' relationship.

Worked Example



Metric: [External Data Transfer] [>] [100 MB] in [60 Minutes]

Filters:

A - [Day of the Week] [is] [Sunday]

B - [Proxied Connection] [is] [False]

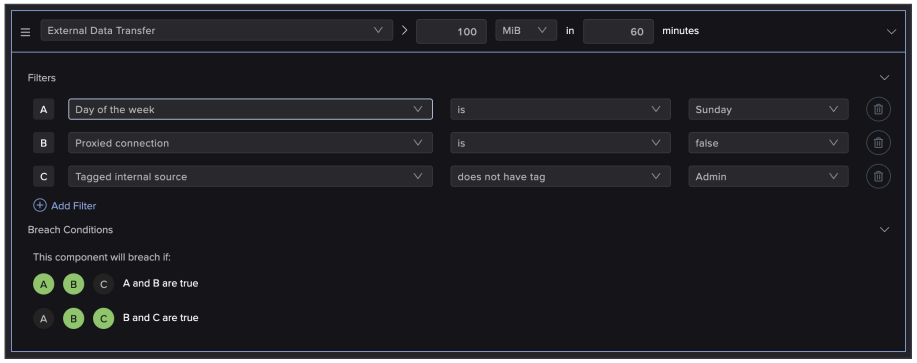
C - [Tagged Internal Source] [does not have tag] [Admin]

Breach Conditions: A & B & C

This limits the component to only breach when all filters - A and B and C are satisfied. In this case, the component would only breach when a non-admin device transfer more than 100mb externally within an hour on a Sunday without using a proxy.

There is often more than one set of criteria we wish to apply to a component, for example when an event can be identified differently over UDP or TCP. Components are therefore not limited to one set of breach conditions - multiple sets can be built in an OR relationship with one another.

Worked Example



Metric: [External Data Transfer] [>] [100 MB] in [60 Minutes]

Filters:

A - [Day of the Week] [is] [Sunday]

B - [Proxied Connection] [is] [False]

C - [Tagged Internal Source] [does not have tag] [Admin]

Breach Conditions:

A & B

B & C

Here, the same filters are in use, but component can breach in two different scenarios - either A and B are satisfied, or B and C are satisfied. In this case, the component would breach when any device transfers more than 100mb externally within an hour on a Sunday without using a proxy, OR if a non-admin device transfers more than 100mb externally within an hour on any day without using a proxy.

At the end of each breach condition, the logic is summarized for review. These conditions can be chained together to create complex and carefully targeted model components.

"Always Required"

Some filters are very restrictive, such as connection direction or restrictions on IPs, and therefore must be applied across all breach condition lines regardless. These filters have blue 'bubbles' and will display "Always Required" in the tooltip on hover.

Adding Display Fields

The final step of defining a component is selecting display fields. Display fields control the at-a-glance information displayed to the user when a breach is triggered by the component, they do not impact the logic of the model itself. Display fields should include the most important information an analyst would need to triage a model breach.

For example, in a connection-based model an analyst would need to know the IP, protocol and port involved in the breach event. In a SaaS model, it would be useful to know the resource (like file or cloud infrastructure element) involved in the activity or the ASN of the source IP to identify the geographic location.

In the component you are working on, expand the **Display Fields** header and click '**+ Add Field**'. The default display field for the main component metric will be added. For example, for the "Connections" metric, the default display field is "Connection Hostname". To choose a different field, click on the name and a selection of relevant filters will appear in the dropdown.

TUNING MODEL SCORING

Score Modulation

Score modulation controls how the model score should change as the model breaches over time for a given device. There are four configuration options.

- **Decreasing**: As a device keeps triggering the same model, the threat score of the breach will lower over time.
- **Increasing**: The more a model fires, the higher the threat score for the device.
- **Stable**: The Threat score will remain the same no matter how often a device triggers a breach of this model.
- **Increase then Decrease**: Initially the threat score will increase but will reduce over time if the model keeps firing.

Some activity becomes less anomalous over time. For example, connections to a new internal server become less concerning as they become part of the 'pattern of life'.

Contrastingly, some activity becomes more anomalous over time. For example, the more failed connections a device makes to internal devices in a short space of time, the higher the likelihood of a malicious network scanning incident.

Other activity remains concerning regardless of frequency. This is particularly true of compliance issues where an activity is always in contravention of a policy regardless of how often it occurs.

Finally, some activity is potentially malicious when it happens frequently in quick succession but can be less concerning if remains frequent over a long period. For example, the scanning activity above is concerning at first but if repeatedly occurring may indicate a network configuration error or a DNS issue rather than a malicious incident.

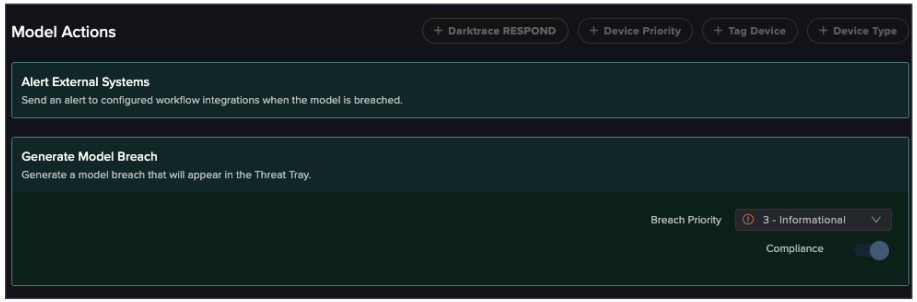
Selecting a score modulation is important for controlling how your new model behaves over a longer period and ensures the alerts it produces stay relevant and useful to other Threat Visualizer users.

ADDING ACTIONS TO A MODEL

Model Actions

The purpose of each model action was covered in further detail in Model Actions. A model can trigger more than one action in response to a breach. Here, it is important to consider the most relevant set of actions for the behavior targeted by the model.

If the model is looking for malicious activity that requires analyst intervention - triggering an alert output is a reasonable response. However, for a model identifying all Office 365 admin users, it is not the most appropriate action to trigger. Instead, an increase in device priority is a logical response to ensure any potential compromise of an admin account has a suitably high threat score.



The **Model** action is default for all models. This creates an event in the device's event log without creating an alert or model breach in the threat tray.

The **Generate Model Breach** action causes the system to generate a model breach that will appear in the threat tray. Within this action, a priority can be set for the overall model breach:

Breach priority

Assign this model a priority score of 0-5. The model's priority will affect the strength with which it breaches, so that if particular types of behaviors are of greater interest to you, these behaviors will register as more strongly relevant and will be more obvious in the threat tray than if the priority was lower.

- **0** System Event
- **1** Low Impact
- **2** Interesting Behavior
- **3** Medium Impact
- **4** Significant Behavior
- **5** High Impact

Other Actions

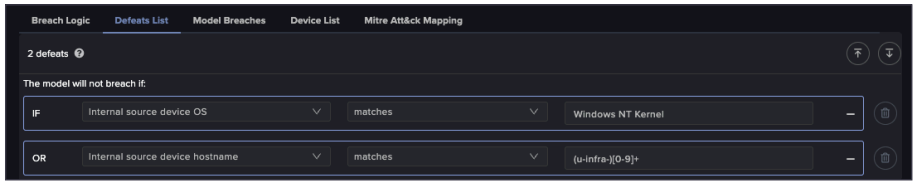
The other available actions are:

- **Alert External Systems**: models with alert turned on will be pushed out to external systems if conditions for such alerting are met.
- **Darktrace RESPOND**: take a Darktrace RESPOND action. See Darktrace RESPOND in the Model Editor for more information.
- **Tag Device**: automatically tag the device or user that meets the conditions.
- **Assign Priority Score**: assign a priority score to this device or user to increase or decrease the scoring of model breaches associated with it.
- **Device Type**: change the device type to one of the preset values (e.g., server, IP phone). Only valid for network devices, not users.

Select the most appropriate actions for the model purpose and proceed.

Defeats List

The final element to configure before saving the new model is the defeats list. Defeats are conditions which prevent the model from breaching but are separated from the model logic. How defeats can help tune the models in your environment was covered in Model Defeats.



For new models, there may not be any defeats you wish to add at this time. Once the model is live, exclusions may become obvious as specific devices trigger false-positive breaches or you wish to exclude certain subnets from the criteria due to the devices within them.

Tuning a model with defeats keeps it relevant and producing valuable alerts without requiring changes to the underlying logic.

Finalizing the Model

Once all the above steps are completed - save the model using the save icon in the top right of the workspace.

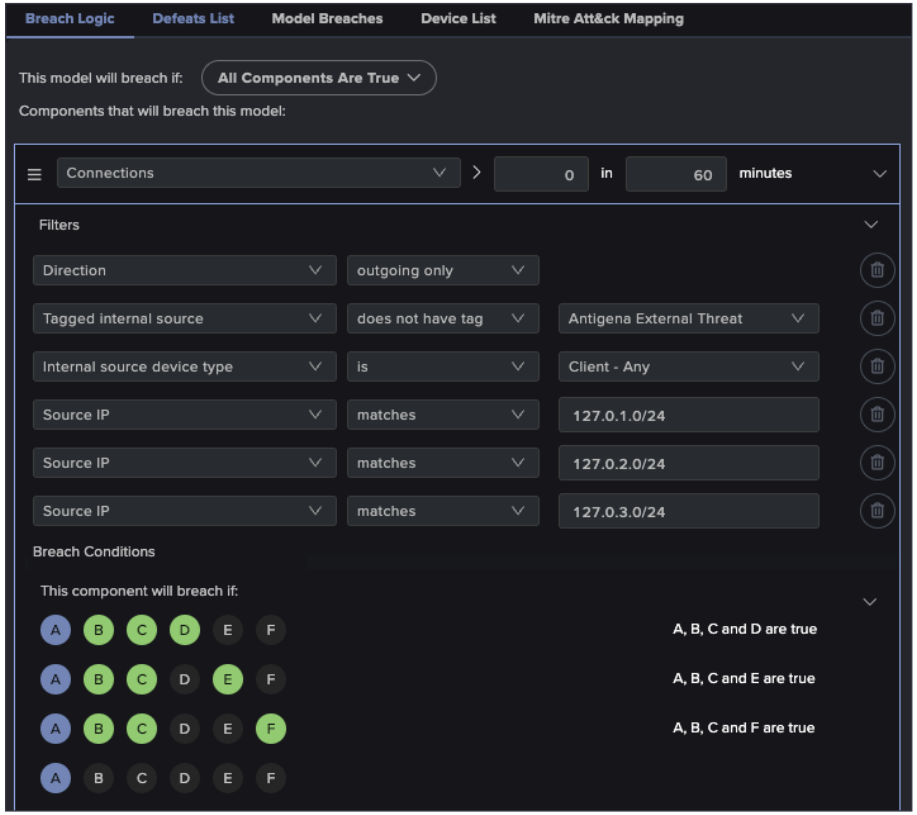
Enter a **commit message** that describes the purpose of the model creation, or of any changes made to the model, and click **Save** once again.

TIPS FOR NEW MODEL MAKERS

The Darktrace model development team create and maintain the extensive Darktrace model deck to detect a wide range of potentially malicious indicators, unusual activity and unwanted network behavior. These experts have provided some top tips and points to consider for new model makers.

Top Tips from Darktrace Experts

- Think about the connection direction for the model you are creating. Are you only interested in inbound or outbound data transfer?
- When creating a tagging model - a model that adds a tag to a device based on given behavior - make sure to add a condition to exclude devices that already have the tag you want to apply. This will keep your model breaches controlled and relevant.
- Don't neglect low-level indicators - it can be useful to set your unusual activity thresholds lower at first and tune them later to catch all the events you are interested in.
- If you are worried that a model may fire too often, check in on it after 15 minutes, then an hour - are you happy with how often it is breaching? If it seems too busy, increase the minimum seconds between model breaches or increase thresholds in your anomaly scores.



CHAPTER 9 - ADMIN INTERFACES

DEVICE & SUBNET ADMIN

DEVICE ADMIN 147
SUBNET ADMIN 148

PERMISSIONS ADMIN

PERMISSIONS ADMIN 148
GROUPS IN PERMISSIONS ADMIN 150

SYSTEM STATUS

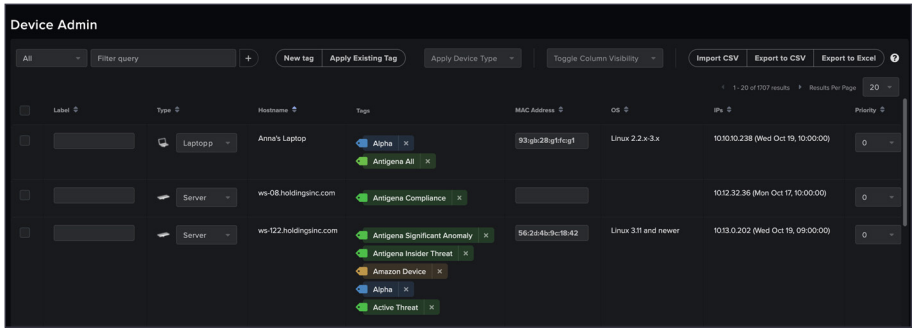
THE SYSTEM STATUS PAGE 152
ALERTS ON THE STATUS PAGE 153
SYSTEM TOPOLOGY 154

DEVICE ADMIN

Viewing Device Admin

The device admin is an up-to-date and historical catalog of all devices ever seen on your network. The devices are by default sorted by when they were last seen. The device admin allows for both the bulk application of tags and the ability to change device types. Some examples of these include adding the tag “Microsoft Windows” to any device that meets certain conditions, or changing all devices with hostnames starting with SEP to IP Phones.

Device Admin can be accessed from the Main Menu under Admin.



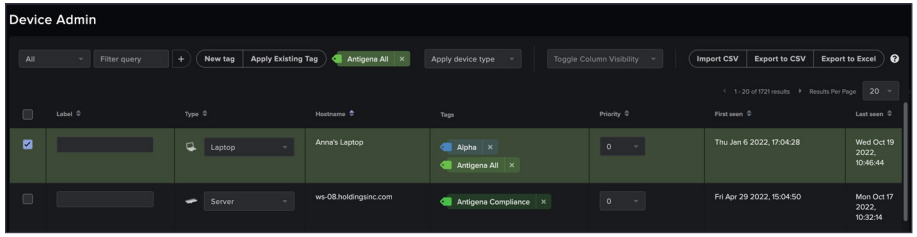
Filtering and Modifying Devices

The search bar allows users to narrow down the listing of devices and can be used in conjunction with previous searches.

Users can search for Hostname: SEP and press + to add that query, and then add additional queries to narrow the results down based on any criteria including: Labels, Tags, Device Priority, Device Types, Hostnames, IPs, Mac Addresses, Vendor (based on OUI of Mac Address), Operating System, First Seen, and Last Seen.

To apply tags, simply click “Apply Tag” which will display all available Tags (for more information on how to create tags, see Adding and Reviewing Tags) as well as a checkbox next to each device.

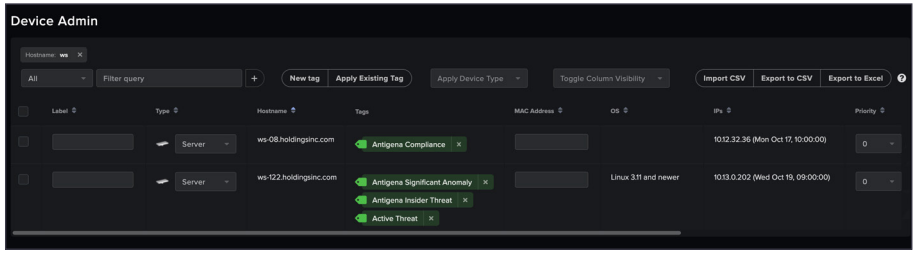
To apply the tag to all devices on the page, you can check the top checkbox; otherwise, you can simply check the checkboxes one by one. To remove tags, you can click the tag you would like and uncheck the box next to a box that is tagged.



Click the 🔗 icon to open the Threat Visualizer centered on the specific device.

Bulk Modifications

At the top of the **Device Admin** page there are buttons to **Import CSV**, **Export to CSV** and Export to Excel. These buttons may be used for audit purposes, or to modify a large number of devices.



To modify a large number of devices, export the **Device Admin** table, edit the fields that need to be changed; save the table as a CSV file; and import the file with new changes. Only the ***Label***, ***Type***, ***MacAddress***, and ***Priority*** fields may be edited in this way. It is not possible to bulk-modify tags using CSV.

It is strongly recommended to keep a copy of the original table in case a backup is needed. Please contact your Darktrace representative if you require assistance in making these modifications.

SUBNET ADMIN

Viewing Subnet Admin

Subnet Admin can be accessed from the main menu under Admin.

The subnet admin provides a catalog of the current subnets being modeled on your environment. Every field is customizable, which allows for enrichment of investigations and usage of the tool.

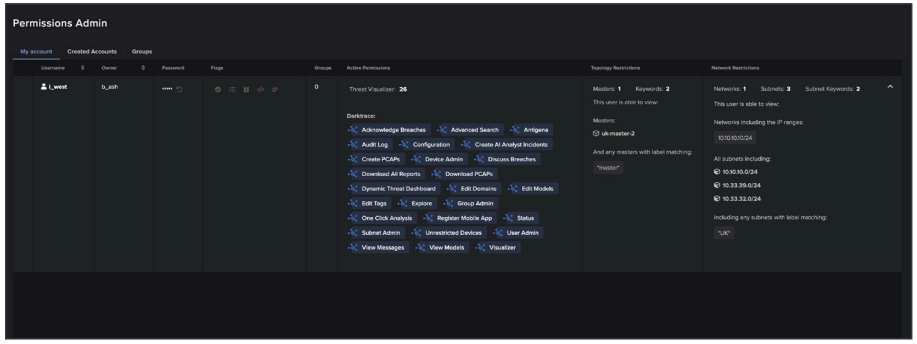
Modifying Subnets

- The label of a subnet can be used to provide a nickname to a device, this will show up throughout Darktrace's Threat Visualizer. For more information about labelling subnets, please see Labelling Subnets and Devices or the System Administration guide.
- The network is where a user can change the subnet size, for example from a /24 to a /25. This may change if DHCP traffic sees a more accurate subnet mask or if there is a successful TCP connection to the broadcast IP of the subnet.
- The location of the subnet can be changed by providing the latitude and longitude, altering where the subnet is displayed on the home screen of the Threat Visualizer.
- The DHCP field is used to specify if the Darktrace system should expect to see DHCP.
- Hostname and Credentials control the type of device tracking assigned to that subnet. Please see Hostname Tracking and Credentials Tracking or the System Administration guide for more details.

PERMISSIONS ADMIN

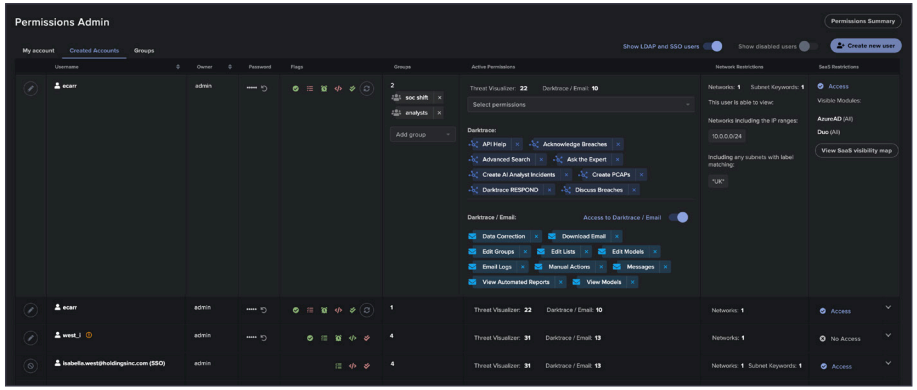
Permissions Admin is the primary location for managing user permissions and visibility for the Darktrace Threat Visualizer, SaaS Console and Darktrace/Email Console. it is accessible from the main menu of the Threat Visualizer (**Admin** > **Permissions Admin**), from the sidebar of the SaaS Console (**User Admin** or **User Groups**), and from Darktrace/Email (**User Admin** or **User Groups**).

By default, the page will open on "My account" which summarizes the settings, group membership, permissions and visibility of the current user. This data is read-only.

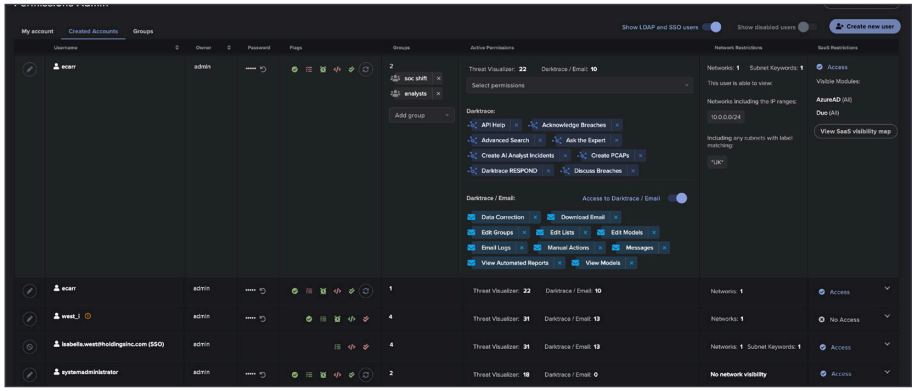


Created Accounts

Permissions in the Darktrace Threat Visualizer, SaaS Console and Darktrace/Email interfaces are handled in a hierarchical structure. The "Created Accounts" tab lets the user review users they have created (and therefore "own"), or that have been created by users they have created.



Click on a row to expand and show more details about the user.



Users are granted permissions directly, or through group membership. The Permissions column shows the effective permissions the user has.

- If the user is a member of one or more groups, these are the permissions granted by group membership.

- If the user is not a member of the group, these are the permissions directly applied to the user. Permissions can be added and removed from the column view.

The state of flags - simple settings associated with the user - can be modified from the column view. Users can also be added to and removed from groups. For more complex edits, click the ✎ pen icon in the first column. This will open a step-by-step process to edit the users' permissions, group membership and flags.

Values in the ***Visibility*** and ***Restriction*** columns (for example, "Topology Restrictions", "cSensor Visibility") are controlled through group membership. By default, all users will have full visibility unless "Manual Data Visibility" is set on the "Groups" tab. For more information, please refer to Controlling Data Visibility and How Data Visibility is Defined.

Some users may appear with a warning icon. This indicates the user has been modified by a higher-permission user and granted permissions the current user does not possess - whether through group membership or direct grant - and the current user is no longer able to modify it.

LDAP & SSO Users

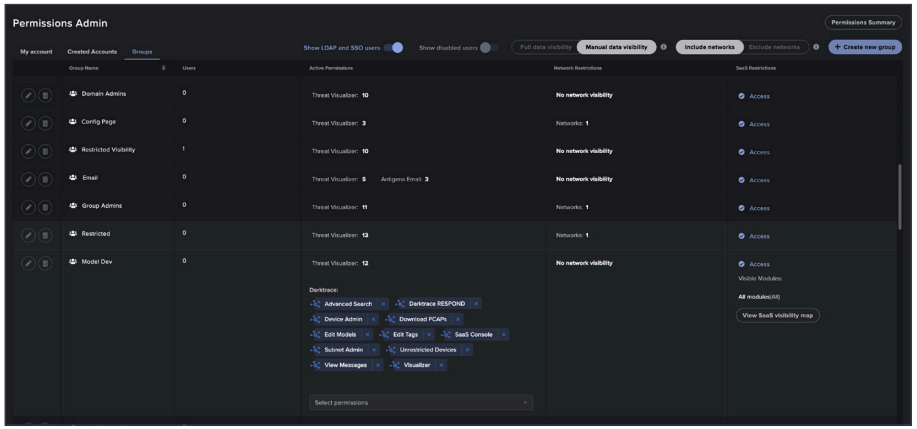
The `admin` user can also review the permissions of individual SSO and LDAP users on this tab; these users will display `(LDAP)` or `(SSO)` after the username.

Permissions and visibility for these non-local users are configured in groups and cannot be edited. Only a limited selection of flags are available as of Darktrace Threat Visualizer 6.

Instead of the user edit ✎ pen icon, the edit column contains a button to revoke mobile app access for these users ⊘. This is a necessary step after the "Register Mobile App" permission has been removed from their corresponding groups.

GROUPS IN PERMISSIONS ADMIN

The “Groups” tab of Permissions admin is used to manage groups of local, LDAP and SSO users. Threat Visualizer users can be added to groups to manage permissions in bulk. Once a user is added to at least one group, the permissions of that group take precedence.



SAML SSO and LDAP users are part of groups in external AD or IdP environments - information about these groups are sent as part of the authentication process. The permissions of these users is controlled by modifying the groups in Darktrace that correspond to those in the AD/IDP environment.

The Threat Visualizer also allows data visibility to controlled at a group level. To create manual data restrictions, users must be placed in groups.

Visibility Options

The “Groups” tab has two additional options which are related to data visibility: “Manual data visibility / Full data visibility” and “Include networks / Exclude Networks”. These settings impact the visibility of all groups and therefore can only be modified by a Darktrace representative or the default `admin` user.

The data visibility setting chosen (“Manual data visibility” or “Full data visibility”) controls whether users are able to see all data, or whether restrictions should be placed at a group level. The default value is “Full data visibility”. For more information, please refer to Controlling Data Visibility.

The networks setting chosen (“Include networks” or “Exclude Networks”) is only relevant if “Manual data visibility” is set above and is applicable to network visibility restrictions only. “Include Networks” is the default state. In this mode, users will only see data for network ranges they have been specifically granted access to by group membership. If set to “Exclude Networks”, users will see all network data **except** for those ranges specified for groups they are members of.

Please refer to How Data Visibility is Defined for more information on controlling network visibility.

Groups

Users will only be shown groups in the “Groups” tab that have equivalent or lower permissions than they possess. Group members will also only be displayed if they are part of the user’s hierarchy.

Click on a row to expand and show more details about the group.

Groups can grant permissions to their members; the Permissions column shows the permissions the group grants. Group permissions are additive - if a user is a member of more than on group, the permissions of all groups are aggregated together.

The permissions assigned to a group can be modified from the column view. Users can also be added to and removed from the group from this view. For more complex edits, click the ✎ pen icon in the first column. This will open a step-by-step process to edit the group permissions, members and visibility.

LDAP & SSO Users

LDAP and SSO users are now included in the membership information but will only be visible to the `admin` user. These users cannot be added or removed from groups manually; membership is controlled by the groups the user is part of in the AD or IdP environment.

Visibility

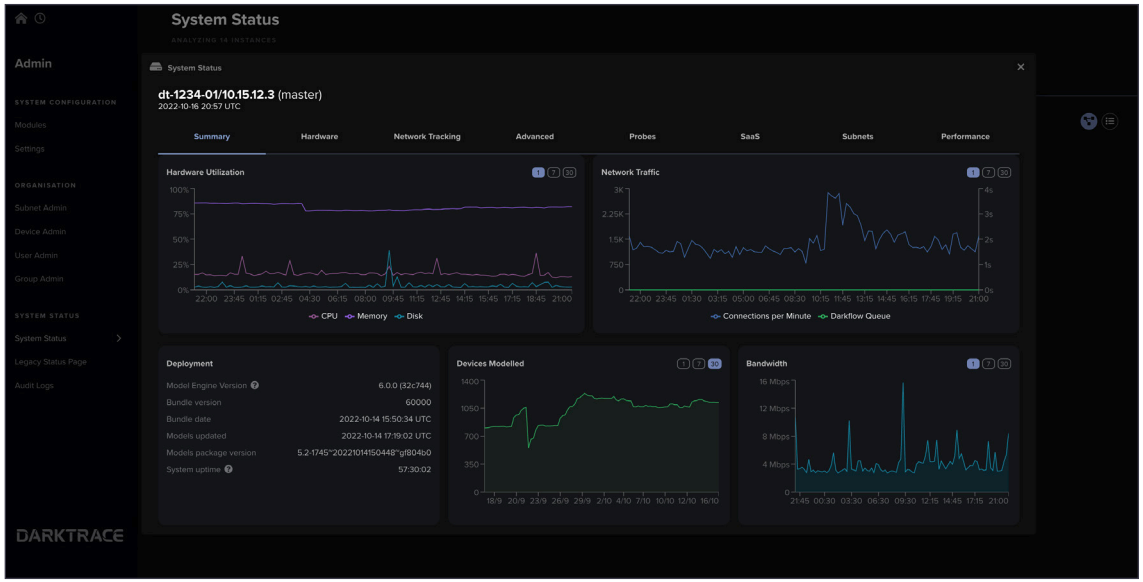
By default, visibility is set to "Full Data Visibility". Users and groups are able to see all data monitored by Darktrace that their permission set allows. Users who are not part of a group are also able to see all data.

When a deployment is moved from "Full Data Visibility" to "Manual Data Visibility", the system applies the effective equivalent of full data visibility to all current groups in the form of visibility expressions. Users not within a group (***excludes default users***) will lose data visibility unless added to a group. New groups will also possess no visibility unless granted during creation.

The visibility granted by each group is displayed in the ***Visibility*** and ***Restriction*** columns (for example, "Topology Restrictions", "cSensor Visibility").

For more information on controlling data visibility, please refer to Controlling Data Visibility and How Data Visibility is Defined.

THE SYSTEM STATUS PAGE



The System Status page contains detailed information about the health and scope of your Darktrace deployment. Here, metrics on hardware utilization, throughput, software bundle versions, component health, and modeled devices can be monitored.

It is important to ensure every part of your deployment is running successfully and within specification, particularly when a deployment is new, or the scope has been increased significantly. Unhealthy deployments, such as those which are overloaded, observing far too many connections for their specification or with unreachable probes or components, will not experience the full benefits of network visibility and consistent monitoring.

In the Threat Visualizer interface, the Status page is accessible from the main menu under **Admin** > **System Status**.

From the SaaS Console interface, the Status page is available from the sidebar (“System Status”).

System Alerts

System alerts keep operators informed of the health of the Darktrace instance and any changes in modeling, data or the health of connected integrations and modules.

When changes or errors occur, a “System Alerts” notification will also appear in the bottom right of the Threat Visualizer workspace, above the Threat Tray. This notification will open the System Alerts tab directly.

ALERTS ON THE STATUS PAGE

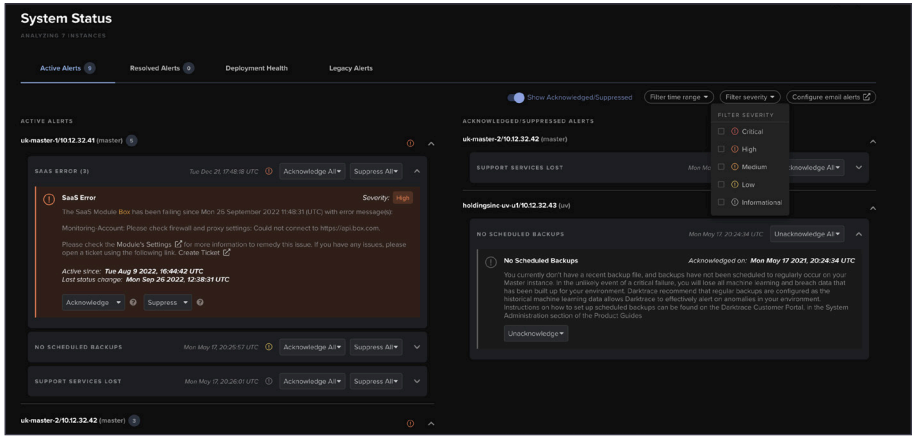
System alerts are notifications about changes to the scope or health of the overall Darktrace deployment. For example, system alerts inform operators when new subnets are detected, when a probe or component loses traffic or becomes unreachable, or when a SaaS Module is experiencing issues.

Historic system alerts were created from model breaches from the System model folder; previously these models would appear in the Threat Tray and generate repeated alerts whilst active. Darktrace Threat Visualizer 5.1 introduces smart, high fidelity alerting with NOC-level capabilities. Alerts will maintain an “active” state whilst relevant, becoming “resolved” once the issue is rectified.

These new system status alerts are available for issuing to external systems - alerts will only be issued when a system event becomes active or resolved. Each alert will also maintain a consistent unique identifier, allowing recurring issues to be identified.

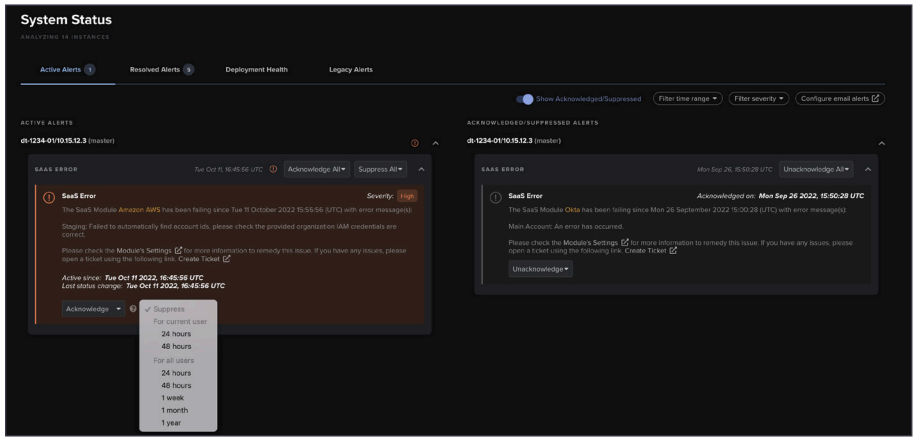
Active Alerts

The **Active Alerts** tab contains current system alerts that are occurring on the Darktrace instance or any child instances and probes associated with it. Alerts are categorized into severity levels - low, medium, high and critical - and can be filtered by severity or by creation time. The level is indicated by the color of the left-hand bar and the ❗ alert icon.



Notifications of the same kind or on the same instance are grouped together - to expand nested alerts, click the instance or alert name. For nested alerts, the highest severity level of all nested alerts will be shown by the ❗ icon color .

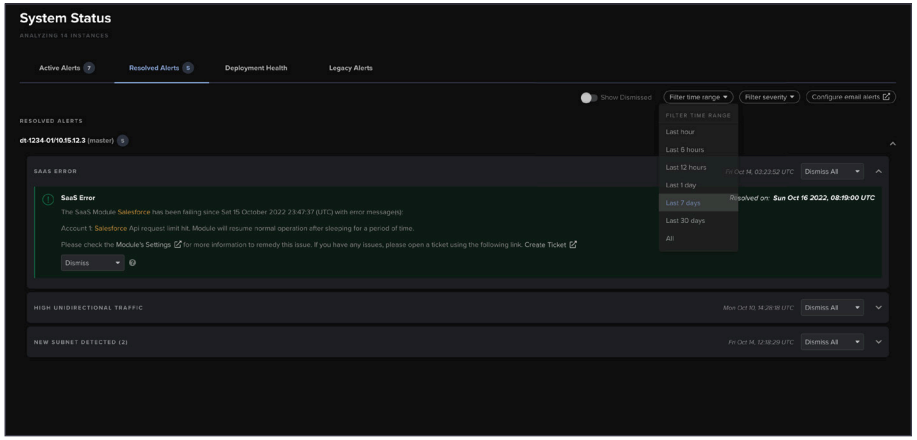
The details section contains information about the event such as the subnet detected, the error experienced by the component or the time that packet loss was detected. Where relevant, details will also include further investigation or remediation steps. Each event includes a link to the Customer Portal to raise a support ticket if further assistance is needed.



Active alerts can be acknowledged to hide them whilst they remain active, or suppressed for a specific duration. Both actions can be performed for the current user, or for all users. Alerts in this state are still viewable by enabling the **Show Acknowledged/Suppressed** toggle and can be unacknowledged or unsuppressed at any time.

Resolved Alerts

The **Resolved Alerts** tab contains system alerts that have been resolved, such as where a SaaS module is no longer experiencing errors, where a drop in traffic has been corrected, or an informational subnet alert has reached expiry.



Alerts on this page can be dismissed for the current user or for all users, removing the resolved alert permanently. To review dismissed alerts, enable the "Show Dismissed" toggle.

If a resolved system event recurs, it will become active again.

Legacy Alerts

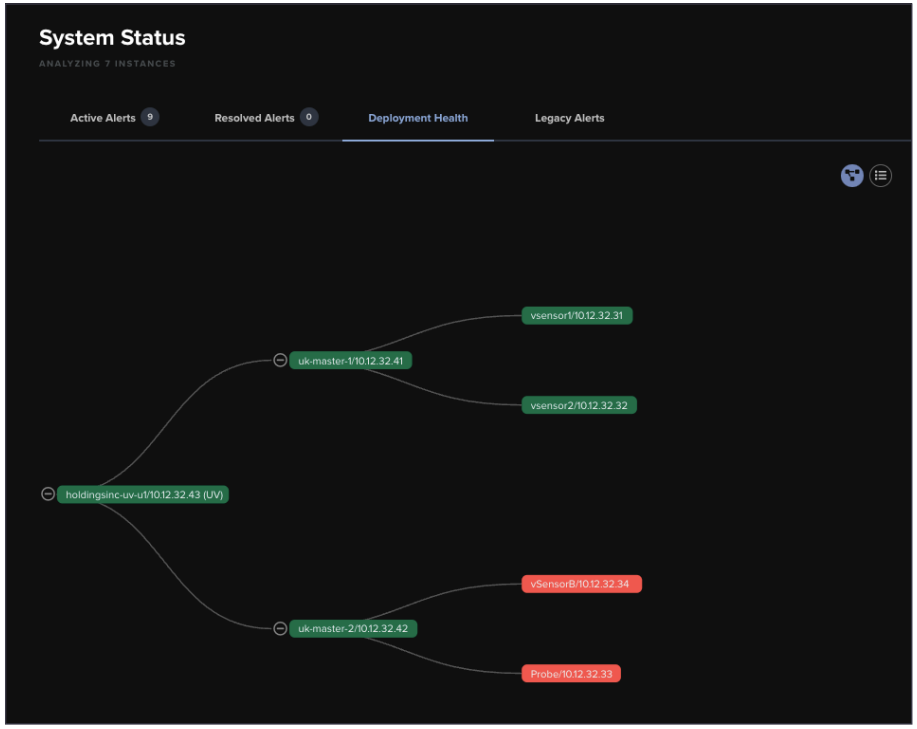
Legacy system model breaches can still be reviewed on the **Legacy Alerts** tab.

Deployment Health

This tab displays an overview of all masters, hardware probes and virtual probes connected to the deployment. There are two views - topology view and list view - which can be toggled in the top left of the workspace. By default, the workspace will open in topology view unless there are a large number of probes and master instances associated with the deployment.

Topology View

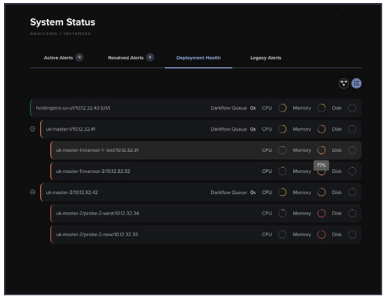
The master appliance accessed is always displayed at the head on the left with any subordinate masters or probes to the right. Where an instance such as a sub master has multiple probes, these can be hidden ⊖ or expanded again ⊕ with the relevant icon.



The color of the instance label represents the overall health. For example, a vSensor with a green label is operating well within its scope, and a master with a red label may be overloaded on one or more key metrics. Hovering over the label will provide more information.

Please note, this view may not be available in deployments with a substantial number of instances.

List View



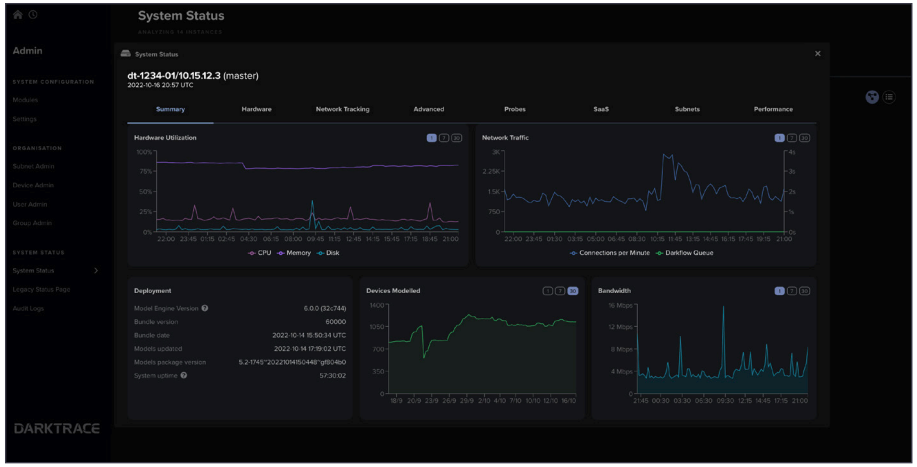
In list view, instance health is shown by the colored line on the left of each row. Current utilization of key metrics - CPU, Load, Memory, Disk and Darkflow Queue - are displayed on the right. Where an instance such as a sub master has multiple probes, these can be hidden ⊖ or expanded again ⊕ with the relevant icon.

Health Metrics

In either view, clicking the name of a master instance will open a "System Status" overview with detailed graphs covering deployment health. For probes, a simple deployment overview is displayed. The instance identifier, the IP and the type of instance are displayed in the top left. Below this, the time the information was retrieved at is shown.

Each graph can show 1, 7 or 30 days to identify trends or a specific date/time where something changed significantly. Hovering over a point on the graph will give details on the specific values at that time.

- The **Summary** tab contains graphs for all key metrics and details of the current Threat Visualizer and Model Bundle versions.
- The **Hardware** tab covers disk utilization with greater granularity alongside other key metrics.
- The **Network Traffic** tab includes ***last seen*** dates for major protocols, greater detail on modeled devices, processed bandwidth and network interface load.
- The **Advanced** tab contains deployment version information, license expiry, and details of internal domains and servers observed by the appliance.



- The **Probes** tab lists all virtual and hardware probes associated with the instance under review.
- The **SaaS** tab lists the current status of Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust Modules.
- The **Subnets** tab lists all subnets currently modeled and their overall tracking quality. Where DHCP warnings are displayed, a link to **Subnet Admin** is provided to modify the tracking settings.

CHAPTER 10 - DARKTRACE RESPOND

DARKTRACE RESPOND ACTIONS

WHAT IS DARKTRACE RESPOND? 157
REVIEWING DARKTRACE RESPOND ACTIONS 159
FILTERING DARKTRACE RESPOND ACTIONS 160
UNDERSTANDING THE DARKTRACE RESPOND ACTIONS PAGE 161
MANUAL DARKTRACE RESPOND ACTIONS 163
WHERE CAN I TRIGGER A MANUAL DARKTRACE RESPOND ACTION? 165

DARKTRACE RESPOND MODELS

DARKTRACE RESPOND MODELS 167
DARKTRACE RESPOND IN THE MODEL EDITOR 168

DARKTRACE RESPOND STATE

HOW IS THE DARKTRACE RESPOND STATE DECIDED 170
THE DARKTRACE RESPOND SCHEDULE 172
USE THE SCHEDULE TO CONTROL AUTONOMOUS RESPOND ACTIONS 173
CONTROL AUTONOMOUS RESPOND ACTIONS ON MODELS 175
ADVANCED RESPOND MODE INFORMATION 177

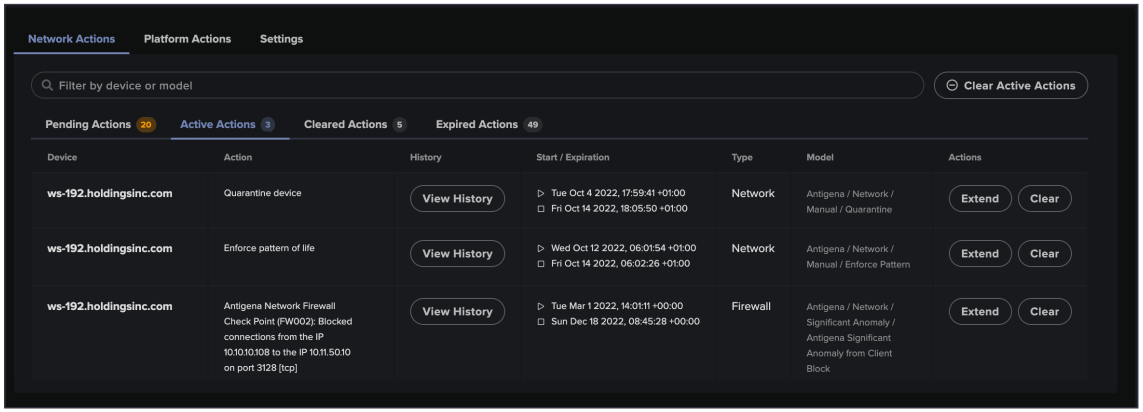
DARKTRACE RESPOND/APPS, RESPOND/CLOUD AND RESPOND/ZERO TRUST

ENABLING DARKTRACE RESPOND FOR DARKTRACE/APPS 178
DARKTRACE RESPOND INHIBITORS FOR THIRD-PARTY SAAS PLATFORMS 181
MAKING USERS IMMUNE TO DARKTRACE RESPOND SAAS ACTIONS 182

DARKTRACE RESPOND/NETWORK

DARKTRACE RESPOND/NETWORK MODELS 183
DARKTRACE RESPOND/NETWORK TAGS 184
DARKTRACE RESPOND/NETWORK INHIBITORS 186
WORKFLOW FOR REVIEWING DARKTRACE RESPOND ACTIONS 187
DARKTRACE RESPOND/NETWORK DEPLOYMENT SCENARIOS 188
DARKTRACE RESPOND/NETWORK AND SERVER DEVICES 190

WHAT IS DARKTRACE RESPOND?



Darktrace RESPOND Framework

The Darktrace RESPOND framework leverages the power of the 'pattern of life' developed across the platform to respond, contain and neutralize emerging threats across the entire digital estate. With its unique understanding of "normal" for each user, device, and component, each autonomous action is targeted and proportionate - disruption to user workflow is minimal.

Darktrace RESPOND is available across multiple coverage areas including Network, Email, Endpoint, Zero Trust and SaaS, extending autonomous response across the enterprise. Where Darktrace DETECT provides the 'pattern of life' to intelligently identify threats, the Darktrace RESPOND framework places the tools in your hands to action and remediate those threats - autonomously, or with human oversight.

Each component offers a combination of autonomous, semi-autonomous, manual and triggered actions. The Darktrace RESPOND framework works with models – when a model breach occurs, the system can be configured to take a range of automatic actions in response or recommend actions for human confirmation or later review. A range of options exist within the platform to configure the operation of Darktrace RESPOND and tailor it to individual requirements.

Darktrace RESPOND/Network

Formerly Antigena Network

The Darktrace RESPOND/Network component applies Darktrace RESPOND capabilities to physical and cloud network devices by controlling connectivity. This control ranges from interrupting communications between distinct endpoint/port combinations up to complete quarantine - actions are proportional to threat and may be escalated if granular blocks are not sufficient. Response can be performed through integration with a number of popular firewalls or by Darktrace virtual sensors such as vSensors and osSensors (both agent and dockerized), ensuring it extends to the farthest reaches of the cloud environment.

Darktrace RESPOND/Apps

Formerly Antigena SaaS

Darktrace RESPOND/Apps extends autonomous response to enterprise software environments. Where anomalous behavior in a third-party SaaS platform begins to escalate, Darktrace RESPOND/Apps can step in and utilize access policies and administrative tools to control the account and sever the malicious actor's access. Eligible modules and the range of SaaS actions available are continually expanded over future software updates.

Darktrace RESPOND/Zero Trust

Formerly Antigena SaaS

Darktrace RESPOND/Zero Trust offers two approaches to securing Zero Trust environments with Darktrace autonomous response; per-user SaaS actions at the ID-provider level and connection (traffic) actions for Zero Trust VPN solutions. By integrating with Zero Trust solutions including Okta, Duo and Zscaler, Darktrace can ensure machine-speed response at the critical pivot-point of many digital businesses. Eligible modules and the range of SaaS actions available are continually expanded over future software updates.

Darktrace DETECT & RESPOND/Email (Darktrace/Email)

Formerly Antigena Email

Darktrace/Email represents a powerful expansion of Darktrace DETECT and RESPOND autonomous, responsive capabilities into the email domain. Complex patterns of life for individual senders, recipients, groups and correspondents are developed and used to detect even the subtlest threats entering via the inbox. The Darktrace/Email component is available for Gmail, Microsoft 365, Hybrid and On-Premises Microsoft Exchange environments. It can be deployed standalone, with one or more Darktrace/Apps modules or integrated with an existing Darktrace DETECT deployment.

A fully comprehensive user guide for the Darktrace/Email system is made available separately.

Darktrace DETECT & RESPOND/Endpoint (Darktrace/Endpoint)

Formerly Darktrace for Endpoint and Antigena Endpoint

Darktrace RESPOND/Endpoint brings award-winning autonomous response capability to the endpoint, enabling AI to take targeted, autonomous actions through Darktrace cSensor agents. Darktrace RESPOND can control network traffic to restrict anomalous connectivity at the system-level, even on remote devices. Devices monitored by cSensors are eligible for Darktrace RESPOND/Endpoint actions if they are licensed, in a Darktrace RESPOND/Endpoint-enabled group (*5.2+*), and possess one or more of the Darktrace RESPOND tags.

As Darktrace RESPOND/Endpoint takes action at the network level on remote endpoint devices, it shares many detection frameworks, inhibitors, settings and configuration processes with Darktrace RESPOND/Network.

Advanced Integrations

Darktrace RESPOND/Cloud for AWS Lambda (Custom)

Where Darktrace RESPOND/Network and Darktrace RESPOND/Apps offer carefully curated inhibitors selected by Darktrace, Darktrace RESPOND/Cloud for AWS Lambda (Custom) instead opens the RESPOND framework to allow the creation of custom actions through invoked AWS Lambda functions. Through this powerful, flexible tool, Darktrace anomaly detection can be used to drive any action and response desired.

REVIEWING DARKTRACE RESPOND ACTIONS

The Darktrace RESPOND Actions Page

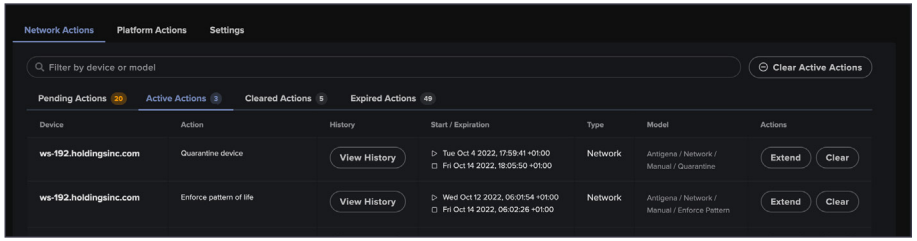
The Darktrace RESPOND Actions window lists all current and expired Darktrace RESPOND Actions. Actions from Darktrace RESPOND components such as Darktrace RESPOND/Network, Darktrace RESPOND/Apps and Darktrace RESPOND/Endpoint will appear in this window. Darktrace RESPOND/Network (including firewall integrations) and Darktrace RESPOND/Endpoint devices appear under the "Network" tab as they impact network devices. Actions against users or entities in third-party platforms will appear under the "Platform" tab.

The types of action shown will differ depending on where the Darktrace RESPOND Actions page was accessed from and which Darktrace RESPOND components are enabled in the environment.

Please note, Darktrace DETECT & RESPOND/Email is displayed separately in the dedicated Darktrace/Email interface.

Threat Visualizer

In the Threat Visualizer interface, Darktrace RESPOND Actions are accessible from the main menu (**Darktrace RESPOND Actions**) or filtered on a specific device from the **Darktrace RESPOND Actions** icon in the Omnisearch bar.

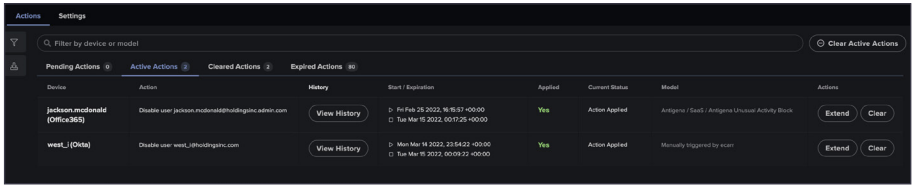


A new window will open on the dashboard.

SaaS Console

From the SaaS Console interface, Darktrace RESPOND Actions is available from the sidebar ("**Darktrace RESPOND Actions** icon"). In this interface, **only platform actions will be visible**.

The Darktrace RESPOND Actions page will open in the main workspace.



Overview

Actions on the Network and Platform tabs are split into pending, active, cleared and expired. Active is the default view.

Pending Actions

Indicates Darktrace RESPOND Actions that have not yet been approved by a human operator.

When there are pending Darktrace RESPOND actions, a notification will appear above the threat tray in the Threat Visualizer. A notification will also appear in the SaaS console threat tray for pending platform actions only.

Active Actions

Displays current Darktrace RESPOND Actions which are in place and have not yet expired.

All currently active actions can be cleared at once using the "**Clear Active Actions**" button in the top right.

Cleared Actions

Lists all Darktrace RESPOND actions that were manually cleared by an operator. Clear will inform Darktrace RESPOND to stop controlling the device or user and suppress the combination of Darktrace RESPOND Action and model breach conditions for 24 hours.

Expired Actions

Includes all previous Darktrace RESPOND activity for devices and users in your environment. Actions can be reactivated and manually extended if desired.

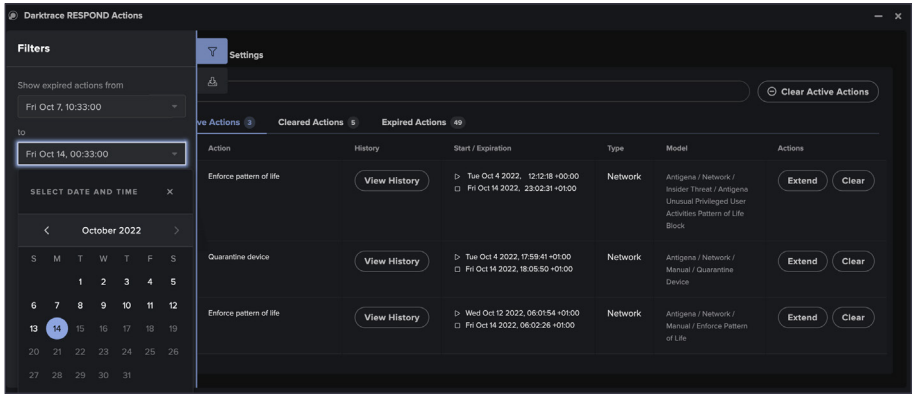
FILTERING DARKTRACE RESPOND ACTIONS

Filtering can be done by action type, by device, by model, or by status.

Each tab can be filtered with a free text search against the model that triggered the Darktrace RESPOND action, or the device or user impacted by the Darktrace RESPOND action. This filter is applied across all action states (active, cleared, etc.)

Filter Panel

Click the filter icon on the left to open the filter panel.



For Network actions, expired actions can be filtered by their expiry time using the time selectors. By default, currently active actions, or those with an expiry in the last 7 days, are returned.

Actions can also be filtered by the exact action Darktrace RESPOND applied, for example, “Block Outgoing Connections” or “Disable User”. By default, all actions are returned.

Manual actions - actions triggered by Threat Visualizer users - are included by default. These actions can be filtered out by turning off the “Manual Actions” option.

Type-specific Filters

Platform Actions

Platform actions can be further filtered by whether the action was successfully applied in the third-party platform.Where “**Applied**” : “**Yes**”, this means Darktrace RESPOND successfully performed the action - such as a forced logout or an IP block. If the status is “**No**”, the **Current Status** column will give more information about why it was not possible to apply the action.

Actions may be prevented for many reasons, such as insufficient permissions for the module or because the user is marked as immune at the configuration level. Setting the “**Applied**” filter to “**Yes**” will only display actions where Darktrace RESPOND successfully actioned the user.

Network Actions

For Darktrace RESPOND/Network actions (***not applicable for firewall integrations or Darktrace RESPOND/Endpoint***), the “blocked” status of the action can be optionally displayed. “**Blocked**” : “**Yes**”, indicates Darktrace RESPOND/Network blocked one or more connections that matched the criteria.

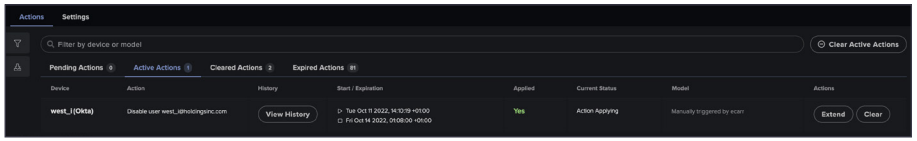
Blocked can refer to the original connection that triggered the action if still active when Darktrace RESPOND responded, or it can refer to subsequent connections that matched the criteria for blocking. In some cases, Darktrace RESPOND/Network could identify a suspicious connection, block any future connections, but no subsequent connections actually occur - here, “**Blocked**” would be “**No**”.

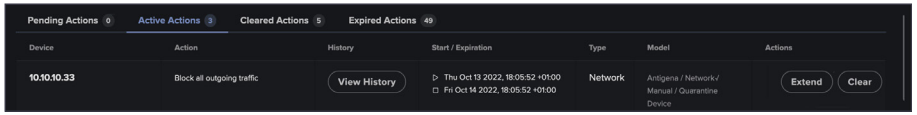
UNDERSTANDING THE DARKTRACE RESPOND ACTIONS PAGE

Now, a single active Darktrace RESPOND action will be focused upon.

Reviewing a Single Action

From left to right, each entry on the Darktrace RESPOND Actions page contains the following:

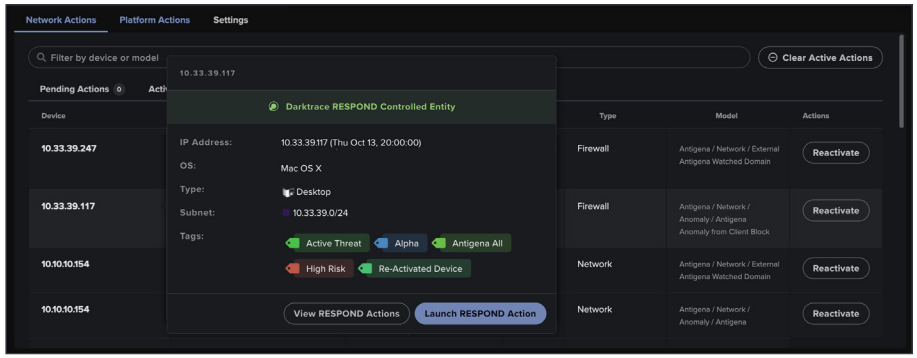




- A device or user under Darktrace RESPOND control
- A summary of the action taking place
- An option to review the history of the action (creation/activation/expiry)
- The start and expiration time of the action.
- The action status, if applicable.
- The model that triggered the action, if applicable.
- Controls to extend, activate, reactive or clear the action.

Device

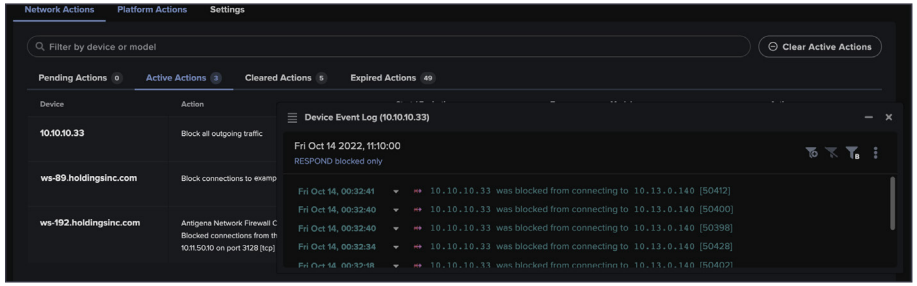
In the Threat Visualizer, hovering over the name of the device or user will show a summary.



Clicking the name will center the visualizer on the device. Clicking the name of a user from the SaaS Console will open their profile.

Action Summary

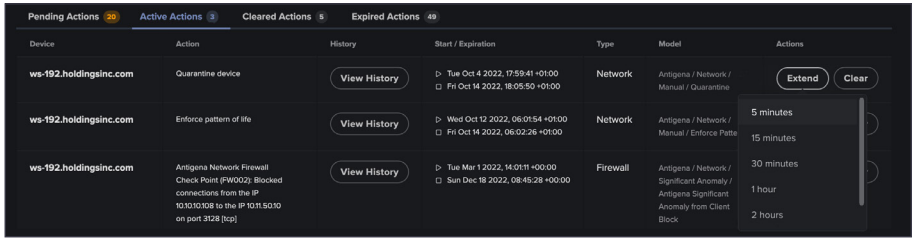
The summary states the type of action and the conditions, such as the IP or port involved. For example, blocking outgoing connections to a specific external IP on port 443, invoking a block through a firewall or blocking a user from logging in from a specific IP.



In the Threat Visualizer, clicking on the action summary for Darktrace RESPOND/Network actions will open an event log of blocked connections.

History

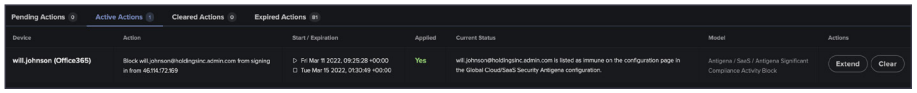
New for Threat Visualizer 6, the "View History" button opens a new window which describes the history of the action state. Actions can be pending (***created (requires confirmation)***), active, cleared, extended, expired or reactivated; the history window allows changes to this status to be tracked over time and linked back to users or system processes. For example, did an action expire before the security team activated it? Did a user manually clear an action?



If configured, the mandatory "reason" users provide when changing the state are also displayed here.

Duration (Start/End)

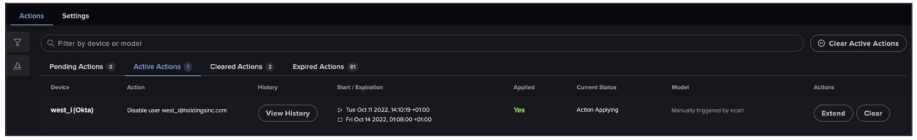
The default action duration is set in the model that triggered the action. For manual actions, the duration can also be set when performing a manual action. Actions can be extended for the desired length of time with the **Extend** button.



Where Darktrace RESPOND performs a one-off, discrete action such as revoking all current logins on a platform, the action will be repeated for the duration of the action. The interval at which it is repeated is defined in the configuration settings.

Action Status (*SaaS/Platform-type actions only*)

The status represents whether the action could be successfully applied in the third-party platform and its last known state.



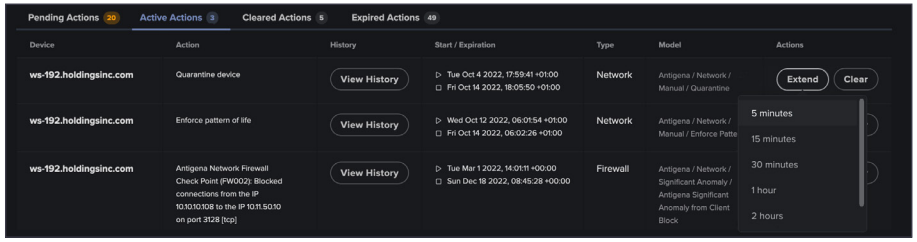
Model

In the Threat Visualizer, clicking the name of the model that triggered the breach will open a Breach log for the model. In the SaaS Console, clicking the model will open the breach investigation overlay centered on the model breach.

Actions triggered manually will be indicated as such in this column.

Extend, Activate, Reactivate, Clear

These buttons allow an operator to change the state of current actions. These states were also covered in Reviewing Darktrace RESPOND Actions.



Extend causes the block to last longer. **Activate** starts a pending action and **Reactivate** restores an expired one. To end the block, click the **Clear** button. **Clear** will inform Darktrace RESPOND to stop controlling the device and suppress the combination of Darktrace RESPOND Action and Breach Condition for 24hrs.

Any changes to the action state - and the user who triggered the change - are recorded in the action history window (see above).

MANUAL DARKTRACE RESPOND ACTIONS

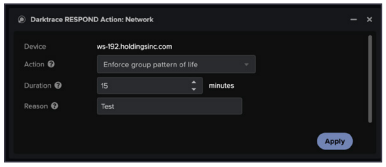
Darktrace RESPOND actions can be triggered manually for Darktrace RESPOND/Network, for eligible devices through the Darktrace Microsoft Graph Security API Integration and for valid modules that support Darktrace RESPOND/Apps, Darktrace RESPOND/Cloud or Darktrace RESPOND/Zero Trust.

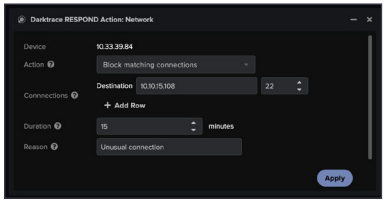
Manual actions let users take the power of autonomous response into their hands, building upon Darktrace AI-powered 'pattern of life' to control device behavior. This can be particularly valuable where external factors, or third-party systems not yet integrated with Darktrace, indicate a device or user may behave strangely or maliciously in the near future. Through Darktrace RESPOND, trigger a preemptive control for a departing employee's device, or put a 'pattern of life' restriction on a device known to be vulnerable from an emerging zero day - buying your team time to mitigate with minimal user disruption.

Manual actions are also a key part of testing your RESPOND configuration - whether to test Darktrace RESPOND/Network reachability after network changes or test upgraded permissions for a newly deployed Darktrace RESPOND/Apps module.

Network Actions

Manual Darktrace RESPOND actions against network devices can be triggered from a number of locations (described in Where Can I Trigger a Manual Darktrace RESPOND Action). For these actions, there are two settings that must be set at a minimum - the desired action and the duration for that action.





Action is the type of response taken against the device. The list of actions available to launch for Darktrace RESPOND/Network is available in Darktrace RESPOND/Network Inhibitors. Most inhibitors (actions) are focused on the AI-powered 'pattern of life' for the device and do not require any connection or endpoint information to be populated.

Depending on the location the action is launched from, some inhibitors may appear prepopulated. For example, a manual action launched from a connection event log (see Investigating a Connection in the Device Event Log will prepopulate with the "block matching connections" inhibitor and details of the connection.

The **Duration** is the initial time the action should be applied for.

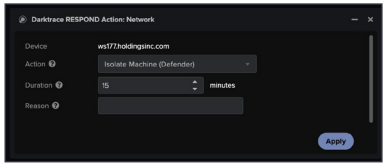
Some actions may also have additional fields, such as defining an IP to be blocked, which are also required. To launch a "Block matching connections" action, the destination IP/hostname and optional port must be specified. Multiple entries to block as part of the action can be added (➕ Add Row).

Complete any additional fields, then click **Apply** to trigger the action.

Security Integrations

The reach of Darktrace RESPOND/Network can also be extended through Darktrace Threat Intelligence ***Security Integrations*** for Microsoft Graph Security API (via Microsoft Defender for Endpoint) and CrowdStrike. These integrations invoke capabilities on host-based Defender or CrowdStrike agents to quarantine, isolate or contain devices that may not be reachable by Darktrace RESPOND/Network at the time.

These actions can only be taken manually and are not available as RESPOND inhibitors for Darktrace models.

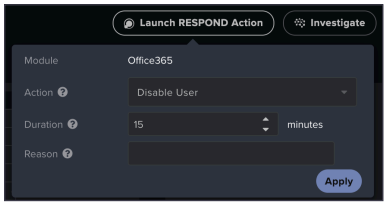


For these actions, a special inhibitor will appear when launched against a valid device. The inhibitor will reference the security integration platform as part of its name.

Proceed to trigger these actions as above.

Platform Actions - Users

Manual actions against users detected by Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules can be manually performed on a user from the SaaS Console or the Threat Visualizer (refer to Where Can I Trigger a Manual Darktrace RESPOND Action for exact details). For these actions, there are three settings that must be set at a minimum - the module intended to take the action, the desired action and the duration for that action.



The **Actions** available to take will differ between different third-party platforms - a list is available in Darktrace RESPOND SaaS Inhibitors. The **Duration** is the initial time the action should be applied for - one off actions will be repeated for the length of the duration.

- If the user has been seen on multiple platforms where Darktrace RESPOND actions can be taken, the **Module** field can be used to search for the SaaS module intended to take the action.

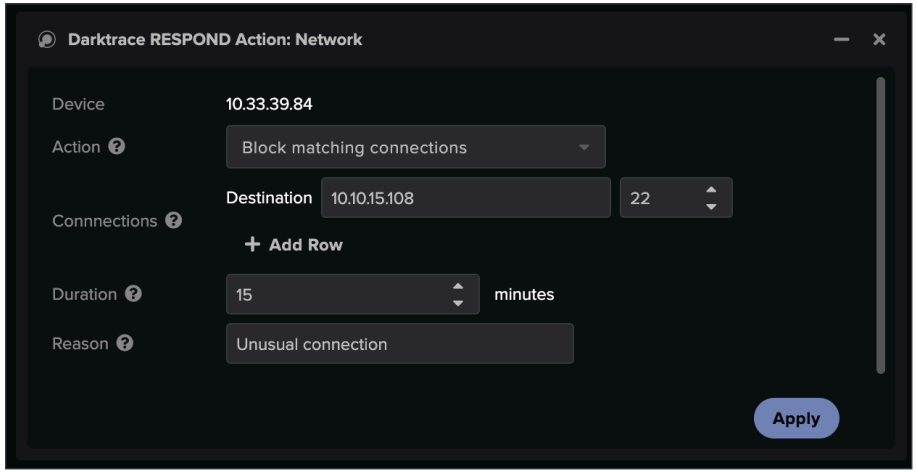
- If the user has been seen on only one platform where Darktrace RESPOND actions can be taken, the **Module** field will be pre-populated with the Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust module that can take an action.

  From Darktrace Threat Visualizer 6, for users detected on more than one platform with Darktrace RESPOND capabilities, the "Disable User" action can be triggered on all modules at once. To do so, select "All Eligible Modules" rather than an individual module. This functionality is only supported for actions launched in the SaaS Console interface.

- If the user has not been seen on a platform where Darktrace RESPOND actions can be taken, the button will not appear and manual actions cannot be performed.

Some actions may have additional fields, such as defining an IP to be blocked, which are also required. Complete any additional fields, then click **Apply** to trigger the action.

Action Reason



From Darktrace Threat Visualizer 6, users can be required to provide a mandatory "reason" (text) when they create a manual action or change the state of an action (activate/extend/clear). This is controlled by the System Config setting "Require Antigena Audit". If enabled, the reason field will also appear in the dialogue for creating manual Darktrace RESPOND actions.

WHERE CAN I TRIGGER A MANUAL DARKTRACE RESPOND ACTION?

The following locations describe where a manual action can be triggered from.

SaaS Console (Platform Actions Only)

Darktrace RESPOND SaaS actions for supported Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules can be manually performed on a user from the SaaS Console

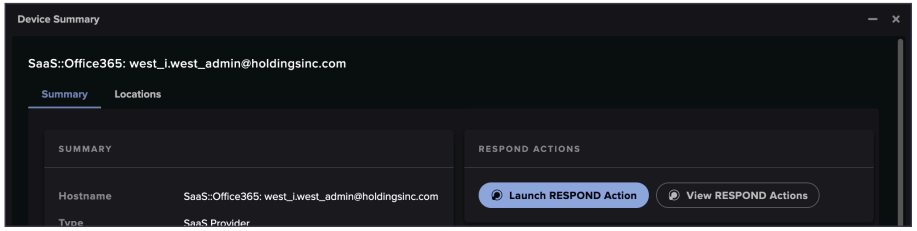
To perform an action on a user from the **SaaS Console**, navigate to the user's **Profile** page from the main menu, the omnisearch or from an event clickthrough. In the top right, the **Launch RESPOND action** button should be available. Click on this button to open the manual action dialog.

Users currently being actioned by Darktrace RESPOND have a green bar beside their profile tile. More information can be found in User Profiles as part of the SaaS console guide.

Threat Visualizer

There are multiple ways to launch a manual action against a user or device from the Threat Visualizer. If the Darktrace Defender Advanced Hunting integration is present, this method can also be used to trigger "Isolate" actions against eligible devices.

Device Summary (Network & Platform)

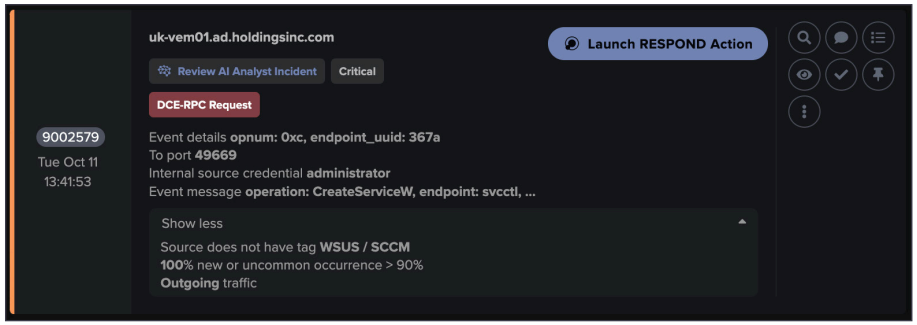


Manual actions can be triggered from Device Summary. To open this window, click the device icon from the Threat Visualizer omnisearch bar. If the user or device is eligible for Darktrace RESPOND actions, the RESPOND Actions section will appear. Click "Launch RESPOND Action" to trigger a manual Darktrace RESPOND action on the device.

For more information on Device Summary, see Device Summary. For other omnisearch options, refer to Device Investigation Options.

Model Breach Logs (Network & Platform)

Model breach logs contain a list of model breaches for a model, a device or a subnet over a chosen time window. Breach logs can be launched from AI Analyst incidents (related model breaches), from the homepage summary, from the threat tray, from the omnisearch for devices or subnets ("View Breaches") and from many other locations in the Threat Visualizer.



On the body of the entry, the "Launch RESPOND Action" button appears for eligible device or user types. Click to launch a manual action against the device that triggered the model breach.

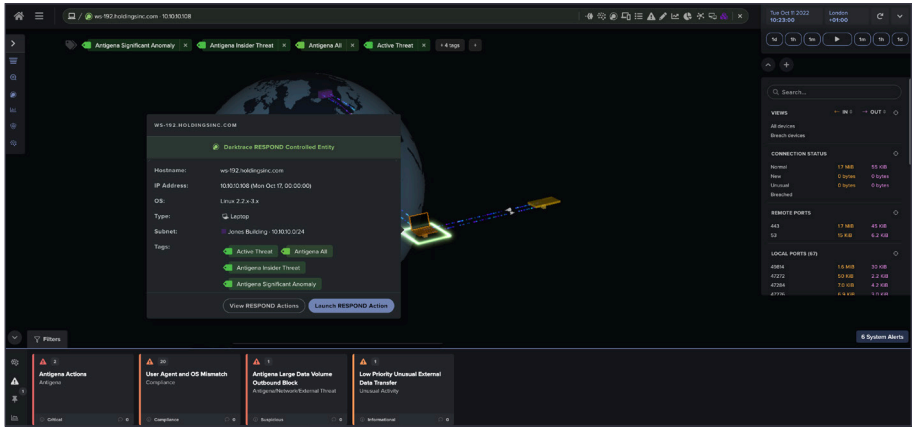
For more information on model breach logs, see Opening the Model Breach Log.

Device Tooltips (Network & Platform)

In almost any location where a device appears in Threat Visualizer, hovering over the text or node displays a short summary tooltip. For example, this tooltip can be observed when hovering over:

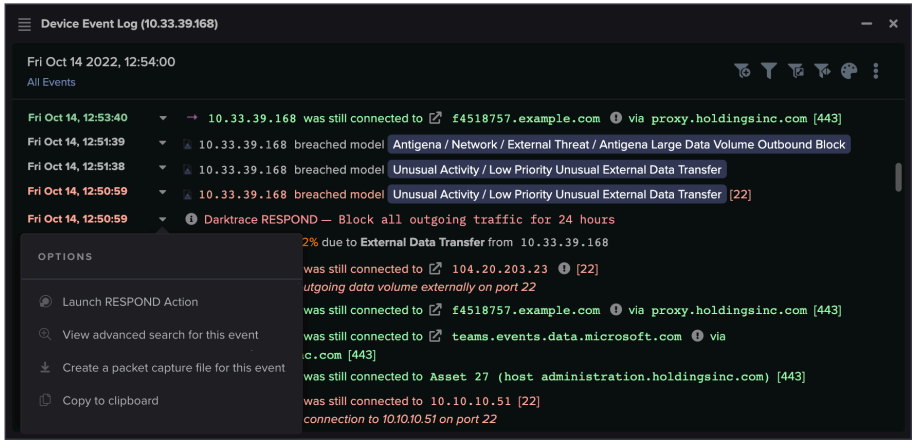
- Device nodes in the 3D Threat Visualizer subnet or device visualization
- Device labels, hostnames and IPs that appear in event logs (network only)
- Device details in AI Analyst incidents
- Devices that have triggered models in the model breach log

From these tooltips, click "Launch RESPOND Action" to trigger a manual Darktrace RESPOND action on the device.



For more information on device visualization, see Focusing on a Device. To learn more about event logs, please refer to Using the Device Event Log.

Device Event Logs (Network Only)



In event logs, manual Darktrace RESPOND actions can be triggered from specific connection entries or from general device tooltips. If you wish to launch a RESPOND action against network entities in a connection, there are two ways to do so:

- From the hover-over summary, click "Launch RESPOND Action" (as above)
- From the ▾ down arrow, chose "Launch RESPOND Action" to create a "Block Matching Connections" action from the connection itself.

For more information on device event logs, see Investigating a Connection in the Device Event Log.

Darktrace RESPOND Actions Window (Platform Only)

Actions against users from supported Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules can also be launched from a filtered version of the Darktrace RESPOND Actions Window. To open this window, click the Darktrace RESPOND icon from the Threat Visualizer omnisearch bar when the user is selected.

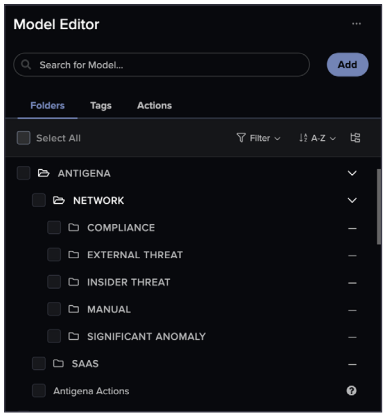
If the user is eligible for Darktrace RESPOND actions, the **Trigger Action** button will appear in the top right.

DARKTRACE RESPOND MODELS

Darktrace models are a series of logical statements and conditions which, if met, trigger an alert or action; models are primarily used to identify and alert on anomalous and potentially malicious behavior. The models framework leverages both the underlying 'pattern of life' detection and outputs from Darktrace Deep Packet Inspection, telemetry inputs, Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules.

Darktrace RESPOND takes action by monitoring model breaches and connectivity for indications of potentially malicious activity or significant anomalies that deviate from the 'pattern of life'. When Darktrace RESPOND components (*excludes Darktrace/Email*) are enabled in your environment, the models used by Darktrace RESPOND (Antigena) will become visible in the Model Editor.

Finding Darktrace RESPOND Models in the Model Editor



Darktrace RESPOND Models are found in the **Antigena** folder of the Model Editor. When Darktrace RESPOND components (*excludes Darktrace/Email*) are enabled in your environment, this folder will become visible in the Model Editor. There are subfolders which correspond to different types of Darktrace RESPOND coverage - for example, **Antigena** > **Network** for Darktrace RESPOND/Network.

Each model detects a specific type of activity or looks for unusual activity above a certain threshold, then triggers an autonomous response if found/met.

To look at all possible models that can take a RESPOND action, navigate to the **Model Editor**. Select the "Actions" tab, then add a filter for "RESPOND - All". Models which can trigger an autonomous response are returned. The state - whether the model will act autonomously or ask for human confirmation - is depicted by one of three icons:

ICON	DESCRIPTION
	Take autonomous
	Force autonomou
	Force human con

More information about what these states mean is covered in Darktrace RESPOND in the Model Editor and How is the Darktrace RESPOND State Decided.

Darktrace RESPOND (Antigena) Models

Some Darktrace RESPOND models are what are known as meta-models - models which use other models breaching as part of their logic. For example, the model "*Antigena Large Data Volume Outbound Block*" (**Antigena** > **Network** > **Insider Threat**) looks for model breaches with a score over 20% that refer to large outbound data transfers. When one of these models breaches - such as the *"Anomalous Connection / Data Sent to Rare Domain"* model - Darktrace RESPOND/Network will observe the breach, trigger the "*Antigena Large Data Volume Outbound Block*" model and take an automatic action against the device transferring the data.

Other Darktrace RESPOND models take action directly in response to specific criteria. For example, the model "*Antigena FTP Block*" (**Antigena** > **Network** > **Compliance**) looks for outbound FTP transfers over 1 MB. If a device is observed making these insecure transfers, Darktrace RESPOND/Network will block matching connections for 1 hour.

If the "Generate Model Breach" action is set on a Darktrace RESPOND (Antigena) model then, like a standard model, it will create model breaches that can be seen in the Threat Tray of the Threat Visualizer or the SaaS Console.

DARKTRACE RESPOND IN THE MODEL EDITOR

Darktrace RESPOND in the Model Editor

Models can take a series of actions in response to breaching such as triggering an alert, raising the priority of a device or adding a tag. When Darktrace RESPOND is enabled on a deployment, the additional **Darktrace RESPOND** action becomes available. This **Darktrace RESPOND** action is already used by the Darktrace RESPOND models described above but can also be added directly to custom models.

Darktrace RESPOND responses are described as 'inhibitors' as they inhibit anomalous behavior. Full details of the inhibitors available are included in the documentation for each Darktrace RESPOND component.

Open a model with Darktrace RESPOND capabilities or add the '**Darktrace RESPOND**' action to an existing model to review the following settings. Some settings may not be visible if the Darktrace RESPOND component is not present in your environment.

Network Inhibitors

This dropdown selects the response that Darktrace RESPOND/Network or Darktrace RESPOND/Endpoint should take if this model breaches and the **Darktrace RESPOND Threshold** value is met.

Full details of Darktrace RESPOND/Network inhibitors are available in Darktrace RESPOND/Network Inhibitors.

Platform Inhibitors

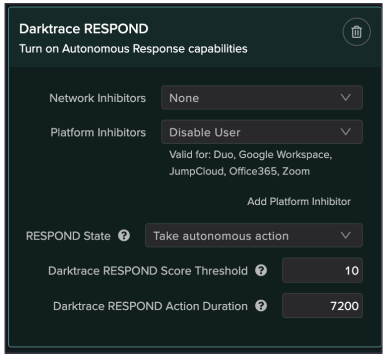
Sets one or more responses that Darktrace RESPOND should take in third-party platforms if this model breaches and the **Darktrace RESPOND Threshold** value is met. **Multiple dropdowns can be added**.

SaaS inhibitors cover Darktrace RESPOND/Apps, Darktrace RESPOND/Zero Trust and Darktrace RESPOND/Cloud modules. As each SaaS and enterprise platform is different; the same capabilities may not be available from one platform to the next. To ensure an action is taken on every applicable platform, more than one can be specified. Once selected, the inhibitor will display the platforms it is applicable to.

Full details of Darktrace RESPOND inhibitors for third-party platforms are available in Darktrace RESPOND Inhibitors for Third-Party SaaS Platforms.

RESPOND State

This setting controls whether the model can take autonomous actions.



- If the model breaches when the Darktrace RESPOND schedule (See How is the Darktrace RESPOND State Decided) is in the "Use Model Setting" state and "Take autonomous action" is selected, an action will be taken.
- If "Force human confirmation" is chosen, or the Darktrace RESPOND Schedule is in the "Require Confirmation" state, human approval will be required for the action to proceed.
- The final option - "Force autonomous action" - will override the global schedule and always take action without confirmation.

Darktrace RESPOND Score Threshold

This field sets the model breach score that must be achieved to trigger the action.

Darktrace RESPOND Action Duration

This setting specifies how long in seconds the action should last for. The default is 3600 (1 hour).

RESPOND STATE: AUTONOMOUS ACTIONS VS HUMAN CONFIRMATION

Deployment Modes

When the criteria for a Darktrace RESPOND action are met, Darktrace RESPOND can either take action autonomously or wait for human approval. These two states are referred to as:

- *Human Confirmation Mode*
- *Active Mode* or *Fully Autonomous Mode*.

In ***Human Confirmation Mode***, Darktrace RESPOND will request approval from a human operator before taking any action. In ***Active Mode***, autonomous actions are taken without human intervention. For many organizations, the end goal of a Darktrace RESPOND deployment is some level of Active Mode, whether applied to select models, select times of the day, or across all devices and use cases.

Human Confirmation Mode

In Human Confirmation mode, Darktrace RESPOND will create recommended actions but it will not take them unless approved by a user. Actions created but waiting for human confirmation are sometimes referred to as "pending" actions, as they are pending human approval.

New deployments of Darktrace RESPOND/Network, for example, are kept in human confirmation mode for a short period to allow Darktrace RESPOND to demonstrate the type of recommended actions it would take across your extended network.

You can review the average time it takes for your team to approve human confirmation actions - and how many actions expired before approval - on the Darktrace RESPOND/Network Summary.

Active Mode / Autonomous Mode

In the second mode, Darktrace RESPOND will create and take actions automatically, without the need for human intervention. These actions can still be cleared or extended at any time by a human operator.

This is the ideal end state for Darktrace RESPOND - actions can be taken at machine speed, without delays waiting for user approval.

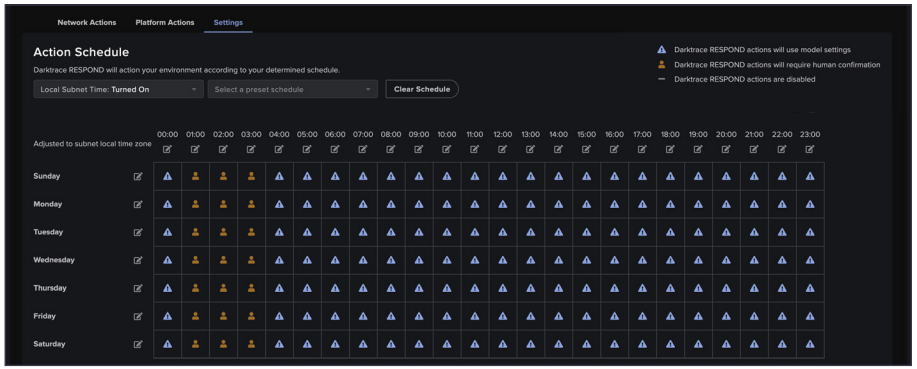
HOW IS THE DARKTRACE RESPOND STATE DECIDED

There are two ways to control whether Darktrace RESPOND will action autonomously: the Darktrace RESPOND schedule, and the Darktrace RESPOND state on the model that triggers the RESPOND action. **The Darktrace RESPOND schedule is the simplest, recommended way to control the mode that Darktrace RESPOND operates in.**

Modifying the model state is straightforward, but should only be attempted when you are familiar with Darktrace models and the Darktrace RESPOND system. Later articles will describe how to change the state globally using the schedule, or on a per-model basis.

Darktrace RESPOND Schedule

The Darktrace RESPOND schedule is the simplest, recommended way to control the mode that Darktrace RESPOND operates in. The schedule allows periods of time to be allocated for autonomous actions and other periods of time allocated to human confirmation mode. A common configuration is that Darktrace RESPOND is permitted by the schedule to operate in active mode outside working hours, where operators may not be available to approve pending actions.



There are three possible states:

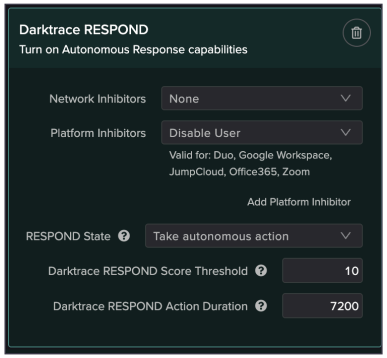
- Always Require Confirmation
- Use Model Settings
- Darktrace RESPOND Disabled

When criteria for a Darktrace RESPOND model are met, Darktrace RESPOND creates an action. The schedule state then controls the actions that result from these Darktrace RESPOND models; "Always Require Confirmation" will make models that would usually ***take autonomous action*** ask for human confirmation before proceeding.[1] Those configured to always request (***force***) human confirmation will continue to do so.

Similarly, "Use Model Settings" will allow those that are configured to ***take autonomous action*** to do so. If the model is set to always request (***force***) human confirmation, this is also respected. "Darktrace RESPOND Disabled" will prevent those models configured to ***take autonomous action*** or request (***force***) human confirmation from creating any actions during its allocated time period.

[1] *Force autonomous action" excluded, please refer to Darktrace RESPOND in the Model Editor or further information below.*

Darktrace RESPOND Models



Darktrace RESPOND models also carry an individual RESPOND state; there are three possible states:

- Take autonomous action
- Force autonomous action
- Force human confirmation

By default, all models are in "Take autonomous action" state - this means they will take an autonomous action **if permitted** by the system schedule. This is the preferred state as it means activity can be controlled on a weekly schedule, rather than model-by-model.

Models in "Take autonomous action" state defer to this system schedule - they take autonomous action ***when permitted***. Models in a "Force" RESPOND state will ignore this global schedule and follow their own setting. This is why it is highly recommended to leave all Darktrace RESPOND models in this Take autonomous action" state, at least during early deployment, so that they can be controlled centrally from the schedule.

So is Darktrace RESPOND in Human Confirmation Mode or Active Mode?

When deciding whether a created action should be autonomous or not, Darktrace RESPOND evaluates the model state ***and*** the schedule state.

By default, all models are set to "Take Autonomous Action", therefore:

- Setting the schedule to "Always Require Confirmation" will result in global *Human Confirmation mode*
- Setting the schedule to "Use Model Settings" will result in global *Active mode*

Changing the state of individual models can impact this mode. This relationship can be illustrated in the following table of outcomes, where a "pending" action is one that requires human confirmation:

(SCHEDULE STATE VS MODEL STATE)	FORCE HUMAN CONFIRMATION	TAKE AUTONOMOUS ACTION	FORCE AUTONOMOUS ACTION
Always Require Confirmation	Action created (pending)	Action created (pending)	Action created (autonomous)
Use Model Settings	Action created (pending)	Action created (autonomous)	Action created (autonomous)
Darktrace RESPOND Disabled	No Action	No Action	Action created (autonomous)

The following articles will explain how to change the global state and - for advanced configuration - individual model state.

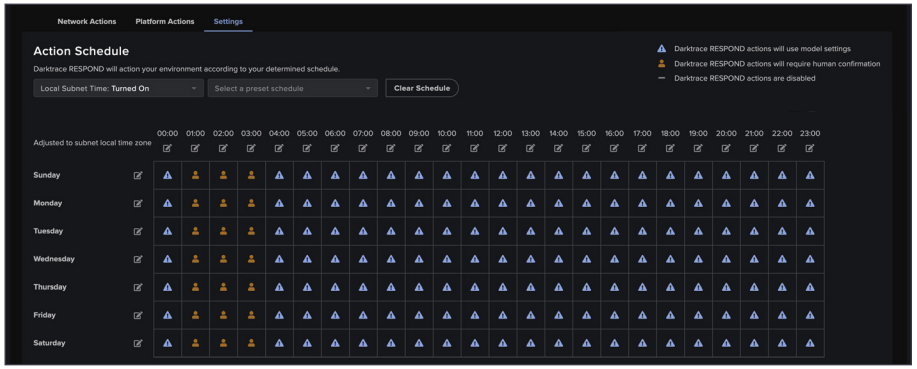
- Use the Darktrace RESPOND Schedule to Control Autonomous RESPOND Actions
- Control Autonomous RESPOND Actions on Individual RESPOND Models.

THE DARKTRACE RESPOND SCHEDULE

Darktrace RESPOND Schedule

Darktrace RESPOND actions can be configured to fit a weekly/hourly schedule. Human Confirmation Mode can be activated during business hours and deactivated when personnel are no longer in the office, ensuring autonomous actions can be taken when an operator is not available to approve them. Response mode can be tailored by hour and day so weekend periods are protected. Similarly, Darktrace RESPOND actions can be disabled entirely during business hours if desired.

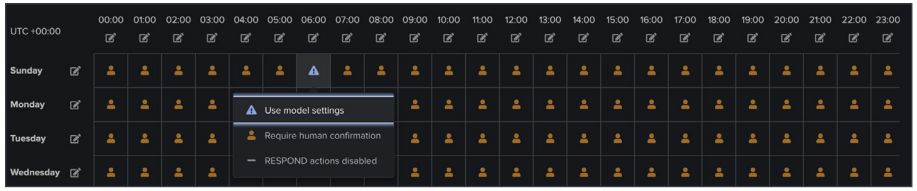
The **Darktrace RESPOND Actions** weekly schedule is located on the **Darktrace RESPOND Actions** page in the **Settings** tab.



The grid represents seven days, comprised of 24 hours. Each block defines when Darktrace RESPOND will seek operator approval for actions, when it will take them automatically (unless the individual model indicates otherwise) and when it will take no action at all. There are three settings: **Use Model Setting**, **Require Confirmation** and **RESPOND Actions Disabled**.

Configuration is set on one, shared master layout. The schedule replaces 'Always Require Confirmation', previously set on the System Config page. The default setting for new deployments is all hourly blocks in active mode.

Configuring a Mode



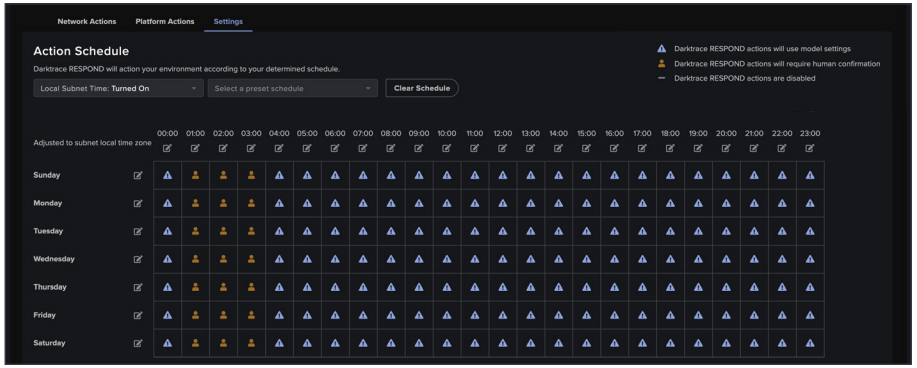
How to configure a mode - or a variation of the two modes - will now be covered in Use the Darktrace RESPOND Schedule to Control Autonomous RESPOND Actions and Control Autonomous RESPOND Actions on Individual RESPOND Models.

For specific scenarios where Darktrace RESPOND may operate outside these configured modes, please refer to Advanced RESPOND Mode Information

USE THE SCHEDULE TO CONTROL AUTONOMOUS RESPOND ACTIONS

The Darktrace RESPOND schedule allows you to control when the system will request human confirmation for actions and when it will take them autonomously. The schedule works in 1 hour blocks over a 7 day period and can be localized to the timezone where your subnets are located. The simplest way to control Darktrace RESPOND is to utilize this schedule.

For deployments which have never used Darktrace RESPOND before, for example, setting the schedule entirely to "Require Human Confirmation" can be a quick away to allow Darktrace RESPOND to demonstrate the actions it would take. Actions created during a block of "Require Human Confirmation" will be left "pending" unless a human chooses to approve the recommendation.



Over time, the schedule can be used to permit blocks of autonomous action ("Use model settings") outside working hours. The final goal is to reach fully autonomous operation.

Human Confirmation Mode (Global)

On initial deployment, it is recommended that Darktrace RESPOND be configured in **Human Confirmation Mode**. In this mode, Darktrace RESPOND will recommend actions but will not perform them unless a human operator gives confirmation. Reviewing the actions Darktrace RESPOND recommends in Human Confirmation Mode will build familiarity and confidence in the kind of autonomous responses Darktrace RESPOND will take and when they are appropriate.

In this mode, Darktrace RESPOND will require an operator to confirm before taking any actions, regardless of whether the individual model breaching has autonomous actions enabled.

1. To enable global Human Confirmation Mode, navigate to the **Darktrace RESPOND Actions** page.

   In the **Settings** tab, add periods of **Require Human Confirmation** to the Antigena (Darktrace RESPOND) weekly schedule.

2. Click in one of the hourly blocks and select **Human Confirmation**. Repeat for all desired hours.

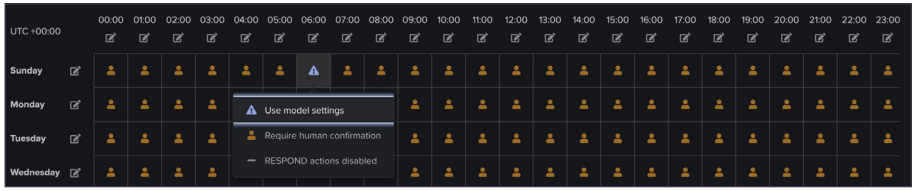
   Alternatively, select the ***Require Confirmation*** preset from the dropdown or click the icon beside each day to fill all boxes with the setting, then refine as desired.

   If you wish Darktrace RESPOND to be disabled outside of these hours, select **RESPOND Actions Disabled** for these time periods.

If Darktrace RESPOND is in passive mode, grid settings will have no impact on Darktrace RESPOND behavior.

Active Mode (Global)

The final goal of Darktrace RESPOND deployments is a fully autonomous deployment state. When you are happy with the actions that Darktrace RESPOND recommends, enabling Active Mode for at least some periods of the week should be strongly considered.



1. To enable global Active Mode, navigate to the **Darktrace RESPOND Actions** page.

   In the **Settings** tab, add periods of **'Use Model Setting'** mode to the Darktrace RESPOND schedule.

2. Click in one of the hourly blocks and select **'Use Model Setting'**. Repeat for all desired hours.

   Alternatively select the **Model Setting** preset from the dropdown or click the icon beside each day to fill all boxes with the setting, then refine as desired.

   If you wish Darktrace RESPOND to be disabled outside of these hours, select **'RESPOND Actions Disabled'** for these time periods.

If one or more models are set to "Force human confirmation", Darktrace RESPOND will **still ask for confirmation** before taking the action. These models must be set to "Take autonomous action" to be in active mode.

If Darktrace RESPOND is in passive mode, grid settings will have no impact on Darktrace RESPOND behavior.

Darktrace RESPOND Models vs the Darktrace RESPOND Schedule

Models in "Take autonomous action" state defer to this system schedule - they take autonomous action ***when permitted***. Models in a "Force" RESPOND state will ignore this global schedule and follow their own setting. This is why it is highly recommended to leave all Darktrace RESPOND models in the Take autonomous action" state, so that they can be controlled centrally from the schedule.

How to create a hybrid deployment - where some models can ignore the schedule - is described in Control Autonomous RESPOND Actions on Individual RESPOND Models.

CONTROL AUTONOMOUS RESPOND ACTIONS ON INDIVIDUAL MODELS

For a more granular approach to Darktrace RESPOND actions, whether Darktrace RESPOND is permitted to act autonomously can be controlled on a per-model basis. This means certain activity can always be actioned without human approval, whilst other activity requires a human to confirm the action Darktrace RESPOND recommends. Conversely, organizations who are moving towards a fully autonomous deployment may still wish to keep some models in human confirmation mode, or roll out autonomous mode slowly.

This hybrid approach allows organizations to tailor Darktrace RESPOND to their daily schedule, compliance policies and risk appetite.

How is Per-Model State Configured?

Individual models have a setting - "Darktrace RESPOND state" (formerly ***Automatic Antigena***) - which controls their individual autonomy. When deciding whether a created action should be autonomous or not, Darktrace RESPOND evaluates the model state ***and*** the schedule state. For example, consider the following scenarios:

SCENARIO	SCHEDULE	MODEL	RESULT
1	Always Require Confirmation	Take Autonomous Action	Action created in human confirmation mode
2	Use Model Settings	Take Autonomous Action	Action taken autonomously
3	Always Require Confirmation	Force Autonomous Action	Action taken autonomously despite schedule state
4	Use Model Settings	Force Human Confirmation	Action created in human confirmation mode

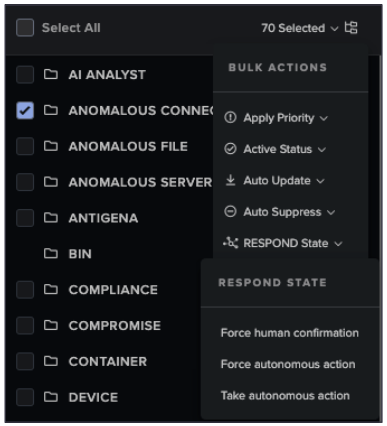
Models in "Take autonomous action" state defer to the system schedule - they take autonomous action ***when permitted***. Models in a "Force" RESPOND state will ignore this global schedule and follow their own setting. Therefore, changing the "Darktrace RESPOND state" field on a select number of models can be a useful way to control individual responses without changing global behavior.

The full interaction between the states is covered in more detail in How is the Darktrace RESPOND State Decided.

Example 1: Mostly Autonomous, Some Models Require Human Confirmation

Human Confirmation can be configured on a model-by-model basis, where Darktrace RESPOND takes some actions autonomously but seeks operator confirmation for specific cases. Models intended for human oversight must be modified in the Model Editor to ensure confirmation will be requested.

Set Specific Models to Human Confirmation Mode



1. To enable Human Confirmation Mode on select models, access the model editor and locate the models you wish to edit.

2. Click on the left of the model name in the list to select it, or select all desired models within a folder.

3. At the top of the list, click the '# Selected' dropdown and choose **Darktrace RESPOND state** > **Force human confirmation**.

   This will place all selected models into confirmation mode.

Put All Other Models in Active Mode

After editing the Darktrace RESPOND state setting for all required models, navigate to the Darktrace RESPOND Actions page.

1. In the Settings tab, add periods of 'Use Model Setting' mode to the Darktrace RESPOND schedule.

2. Click in one of the hourly blocks and select the **Use Model Setting** option. Repeat for all desired hours.

   Alternatively select the ***Model Setting*** preset from the dropdown or click the icon beside each day to fill all boxes with the setting, then refine as desired.

   If you wish Darktrace RESPOND to be disabled outside of these hours, select **RESPOND Actions Disabled** for these time periods. If Darktrace RESPOND is in passive mode, grid settings will have no impact on Darktrace RESPOND behavior.

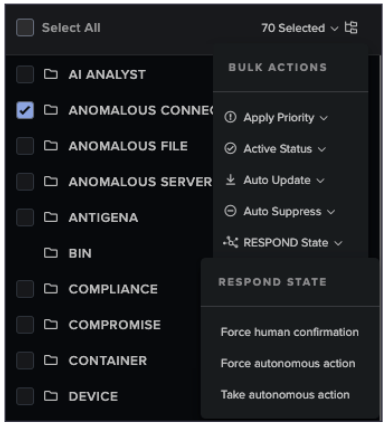
Models which are in RESPOND state "force human confirmation" will still require confirmation, regardless of this configuration.

Example 2: Mostly Human Confirmation, Some Models Force Autonomous Action

The following example demonstrates how to achieve a mostly human confirmation mode deployment but still permit certain models to take autonomous actions in highly anomalous scenarios. Recommendations for new Darktrace RESPOND/Network deployments, for example, strongly recommend that the model "Antigena Ransomware Block" (found in the **Antigena** > **Network** > **External Threat** folder) is permitted to act autonomously at all times. This is regardless of the general schedule state.

Set Specific Models to Forced Autonomous Mode

1. To force specific models to always take autonomous actions, access the model editor and locate the models you wish to edit.



2. Click on the left of the model name in the list to select it, or select all desired models within a folder.

3. At the top of the list, click the '# Selected' dropdown and choose **Darktrace RESPOND state** > **Force autonomous action**.

   This will place all selected models into forced autonomous mode.

The model will now ignore scheduled periods of "Require Human Confirmation" **and** periods of "RESPOND Actions Disabled" in the global schedule. For more information, please refer to How is the Darktrace RESPOND State Decided.

Leave All Other Models in Human Confirmation

After editing this setting for all required models, navigate to the Darktrace RESPOND Actions page to ensure human confirmation is still globally enabled for all other models.

1. In the Settings tab, ensure the schedule shows periods of 'Require Human Confirmation' mode to the Darktrace RESPOND schedule.

2. If not, click in one of the hourly blocks and select the **Require Human Confirmation** option. Repeat for all desired hours.

   Alternatively select the ***Require Human Confirmation*** preset from the dropdown or click the icon beside each day to fill all boxes with the setting, then refine as desired.

   If you wish Darktrace RESPOND to be disabled outside of these hours, select **RESPOND Actions Disabled** for these time periods. This will only impact models which are in "Take Autonomous Action" or "Force Human Confirmation".

If Darktrace RESPOND is in passive mode, grid settings will have no impact on Darktrace RESPOND behavior.

ADVANCED RESPOND MODE INFORMATION

Exceptions

- Manually created actions, such as Darktrace RESPOND/Apps actions, are not subject to the schedule as they have been deliberately taken by a human operator. These actions will not require confirmation and will occur even during periods of 'RESPOND Actions Disabled'.

- Actions reactivated or extended manually will also not be subject to the schedule as they have been deliberately taken by a human operator.

- Models which use the Automatic Darktrace RESPOND option "Force autonomous action" will override the global schedule and always take action without confirmation.

- Actions created by Darktrace RESPOND-enabled endpoints on the Watched Domains will be automatically taken regardless of mode or device tags.

Localizing the Schedule

The setting ***Local Subnet Time*** can be useful for global networks. If enabled, for Darktrace RESPOND/Network actions, the setting applies schedule timings locally based upon the latitude and longitude information for the subnet.

For example, a period of human confirmation mode is set between 9am and 5pm. From a UTC perspective, this is then applied as a period of human confirmation mode between 4pm and 12am for devices in the Los Angeles subnet (9am - 5pm localized) and between 7am and 3pm for devices in the Frankfurt subnet (9am - 5pm localized).

For Darktrace RESPOND/Endpoint actions or Darktrace RESPOND for Apps, Cloud or Zero Trust module actions, the schedule is applied as UTC in this mode.

ENABLING DARKTRACE RESPOND For Darktrace/Apps, Darktrace/Zero Trust and Darktrace/Cloud

Darktrace RESPOND can take a range of proactive, measured, automated actions in the face of confirmed cyber-threats detected in real time. Where anomalous behavior in a third-party SaaS platform begins to escalate, Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules can step in and utilize access policies and administrative tools to control the account and sever the malicious actor's access.

In contrast to ***network actions*** - taken against connectivity between endpoint, cloud and network devices - Darktrace RESPOND SaaS actions are those taken in third-party environments, primarily against user accounts. The suite of actions differs between each SaaS, Zero Trust and Cloud platform - full details are available in Darktrace RESPOND Inhibitors for Third-Party SaaS Platforms.

Platform-based autonomous response is currently available for the following modules:

- Darktrace RESPOND/Apps for Microsoft 365
- Darktrace RESPOND/Apps for Zoom
- Darktrace RESPOND/Apps for Google Workspace
- Darktrace RESPOND/Apps for Salesforce
- Darktrace RESPOND/Zero Trust for Okta
- Darktrace RESPOND/Zero Trust for JumpCloud
- Darktrace RESPOND/Zero Trust for Duo

This provision will be continually expanded over future software updates.

Enable RESPOND capabilities for a Darktrace/Apps or Darktrace/ Zero Trust module

The following steps are required to enable Darktrace RESPOND actions in Darktrace/Apps or Darktrace/Zero Trust modules:

1. Ensure that a Darktrace RESPOND license has been added to your Threat Visualizer environment for the relevant module.

2. Check the level of permissions granted to the Darktrace/Apps or Darktrace/Zero Trust module that has been licensed for actions.

   To do so, navigate to the System Config page and select the appropriate module. On the "settings" tab, for each account, check the state of the "Account Permissions" field (see below). This will show "Monitoring" (DETECT) or "Antigena" (RESPOND).

   To progress, the account must display "Antigena" in this field. The permissions level is granted when the app is authorized. Modules authorized before Darktrace Threat Visualizer 5 or authorized solely for monitoring may require reauthorization to begin taking actions.

   If you are not sure if reauthorization is necessary, consult the Darktrace RESPOND section of the documentation for the module. More information is also available below in "Account Permissions".

3. Optionally configure immunity from actions for certain users and IP addresses. more information can be found in Making Users Immune to Darktrace RESPOND SaaS actions.

4. Finally, Darktrace recommends that new users of Darktrace RESPOND enable autonomous actions at a later date, when they are more comfortable with Darktrace RESPOND functionality.

   For this reason, we recommend navigating to the Darktrace RESPOND Actions page and configuring the Schedule. See Darktrace RESPOND State: Autonomous Actions vs Human Confirmation and The Darktrace RESPOND Schedule for more information.

Account Permissions

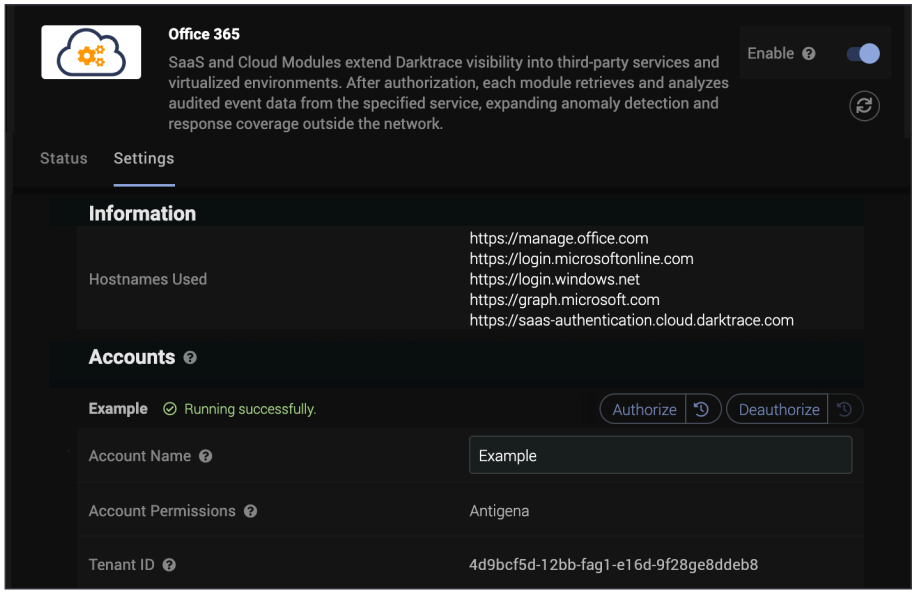
Darktrace/Apps or Darktrace/Zero Trust modules can be authorized for monitoring (DETECT), or for monitoring and Darktrace RESPOND actions. For some modules, Darktrace RESPOND capabilities require more permissions than initially granted for DETECT coverage in order to take these actions.

The type of authorization is listed for each account in the "Account Permissions" field.

Where the module needs to be reauthorized to add Darktrace RESPOND permissions, the reauthorization process will include additional prompts for Darktrace RESPOND permissions or provide alternative authorization links.

Modules operating with monitoring (DETECT) permissions may need to be reauthorized with Darktrace RESPOND permissions before actions can be taken.

Whether this is applicable for a module will be described in the detailed documentation for each module.



MANUAL RESPOND ACTIONS

For Darktrace/Apps, Darktrace/Zero Trust and Darktrace/Cloud

Darktrace RESPOND SaaS actions for supported Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust modules can be manually performed on a user from the SaaS console or the Threat Visualizer.

This is also covered in Where Can I Trigger a Manual Darktrace RESPOND Action?.

SaaS Console

To perform an action on a user from the **SaaS Console**, navigate to their **Profile** page from the main menu, the omnisearch or from an event clickthrough. In the top right, the **Darktrace RESPOND** button should be available. Click on this button to open the manual action dialog.

Users currently being actioned by Darktrace RESPOND have a green bar beside their profile tile. More information can be found in User Profiles as part of the SaaS console guide.

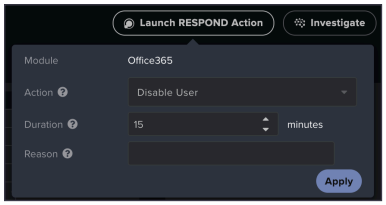
Threat Visualizer

To perform an action on a user from the **Threat Visualizer**, search for the user in the omnisearch bar and select them. From the device options in the omnisearch bar, click the Darktrace RESPOND icon. The Darktrace RESPOND Actions page will open focused on the user. In the top right, click the **Trigger Action** button to open the manual action dialog.

Triggering an Action

There are three settings that must be set at a minimum - the module intended to take the action, the desired action and the duration for that action.

The **Actions** available to take will differ between different third-party platforms - a list is available in Darktrace RESPOND SaaS Inhibitors. The **Duration** is the initial time the action should be applied for - one off actions will be repeated for the length of the duration.



- If the user has been seen on multiple platforms where Darktrace RESPOND actions can be taken, the **Module** field can be used to search for the SaaS module intended to take the action.

- If the user has been seen on only one platform where Darktrace RESPOND actions can be taken, the **Module** field will be pre-populated with the Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust module that can take an action.

- If the user has not been seen on a platform where Darktrace RESPOND actions can be taken, the button will not appear and manual actions cannot be performed.

Some actions may have additional fields, such as defining an IP to be blocked, which are also required. Complete any additional fields, then click **Apply** to trigger the action.

DARKTRACE RESPOND INHIBITORS FOR THIRD-PARTY SAAS PLATFORMS

Models can take a series of actions in response to breaching such as triggering an alert, raising the priority of a device or adding a tag. When Darktrace RESPOND is enabled on a deployment, the additional **Antigena** (Darktrace RESPOND) model action becomes available.

If a Darktrace/Apps, Darktrace/Zero Trust or Darktrace/Cloud module with RESPOND capabilities is enabled within your Darktrace environment, it comes with a category of **Darktrace RESPOND** models already using the **Darktrace RESPOND** action to perform responses.

This action allows an operator to set one or more 'inhibitors' - actions to ***inhibit*** the behavior that matches the model criteria. Darktrace RESPOND SaaS actions are those taken in third-party environments, primarily against user accounts. The suite of actions differs between each SaaS, Zero Trust and Cloud platform - models can therefore have multiple inhibitors set to ensure they can take an action in every possible environment.

Darktrace RESPOND SaaS inhibitors can also be triggered manually.

Users and IP addresses can be made 'immune' from specific inhibitors or from all inhibitors on a per-module basis. By default, all users known to a module licensed for Darktrace RESPOND are eligible for actions. Immunity from actions is controlled on the System Config page.

Please see Making Users Immune to Darktrace RESPOND SaaS actions for more detail.

SaaS Inhibitors

INHIBITOR	APPLICABLE MODULES	DESCRIPTION
Block IP(s)	Microsoft 365	Prevents access to the account from an IP or IP range for the duration set.
Disable User	Microsoft 365, Zoom, Okta, Duo, Google Workspace, JumpCloud, Salesforce	Disables a user account for the duration set.
Freeze User	Salesforce	Freezes a user account for the duration set.
Force Logout	Microsoft 365, Zoom, Google Workspace	Forces the user to log out from the platform. This action is a one-off action, so will be repeated at the configured interval for the duration set. The default interval is 15 seconds and can be altered by a member of Darktrace support if required.

For each module, any considerations for Darktrace RESPOND actions are covered in more detail in the module documentation:

- Darktrace RESPOND/Apps for Microsoft 365
- Darktrace RESPOND/Apps for Zoom
- Darktrace RESPOND/Apps for Google Workspace
- Darktrace RESPOND/Apps for Salesforce
- Darktrace RESPOND/Zero Trust for Okta
- Darktrace RESPOND/Zero Trust for JumpCloud
- Darktrace RESPOND/Zero Trust for Duo

MAKING USERS IMMUNE TO DARKTRACE RESPOND SAAS ACTIONS

All users are eligible for Darktrace RESPOND actions by default if the Darktrace/Apps, Darktrace/Cloud or Darktrace/Zero Trust module they were observed by has licensed RESPOND capabilities. Immunity from actions is controlled on the System Config page.

To make changes, access the **System Config** page from the main menu and select **Modules** from the left-hand side. Under Cloud and SaaS Security, select the module you wish to change and click on it to open the configuration window.

Action Eligibility

When a module has **Account Permissions** - **Darktrace RESPOND**, new settings will appear which control whether a user or an IP can be actioned by Darktrace RESPOND SaaS actions. These are **Immune Users** and **Immune IP Addresses**.

- **Immune Users** are users which will not be actioned by Darktrace RESPOND, whether by models or by manual actions. Actions blocked due to immunity will report this status in the Darktrace RESPOND Actions page. The field takes a comma-separated list.

  To make a user immune, enter their username as it appears in the Darktrace Threat Visualizer or SaaS console. Do not include the prefix. For example, **SaaS::Office365: john.ayre@holdingsinc.com** would be entered as just **john.ayre@holdingsinc.com**.

- **Immune IPs Addresses** are IP addresses that will not be blocked by actions that control access to the relevant third-party platform, such as "Block IP". This field will only appear if the module can take IP-based actions.

  The field takes a comma-separated list of IPs or valid CIDR IP ranges.

When "**Per Inhibitor Antigena**" is enabled, immunity can also be set for specific actions. For example, a user could be eligible for lower severity actions like "Force Logout", but immune from more stringent actions like "Disable User". These settings are in addition to the general **Immune Users** and **Immune IP Addresses** settings.

Inhibitors will also differ between modules as capabilities and appropriate actions vary between different platforms.

Example Configuration

Immunity and Account Permissions are set for each account within a Darktrace/Apps, Darktrace/Cloud or Darktrace/Zero Trust module.

For example, an organization may have two Office 365 tenancies - Business and Development. For the Darktrace/Apps Microsoft 365 module, they have two **accounts** which correspond to these two tenancies. These accounts have the **Account Permissions** level of **Monitoring**.

For actions to be taken in both tenancies, both of these accounts need to be reauthorized to add **Darktrace RESPOND Permissions** and upgrade to the **Account Permissions** level to **Antigena** (Darktrace RESPOND). Once this is complete, immunity settings will appear against each account.

The organization would like to ensure the IP `35.176.59.98` is not actioned in **either tenancy**. To do this, it must be added to **Immune IP Addresses** for **both** accounts.

DARKTRACE RESPOND/NETWORK MODELS

Darktrace RESPOND/Network and Darktrace RESPOND/Endpoint expand Cyber AI response to devices by severing network connections, restricting access and quarantining devices by limiting their outbound connectivity.

Actions can be taken when a device exhibits significantly anomalous behavior, when it contravenes a compliance policy, when a device attempts to access a specific watched endpoint, or any other custom criteria defined in a model. Unlike anomaly level, which is determined on behavior, risk profiles will vary across an organization's network according to business requirements.

You may wish some devices and users to be exempt from Darktrace RESPOND Actions, to restrict the actions to only match anomalies exhibiting certain behavior or apply actions automatically only at certain times of the day.

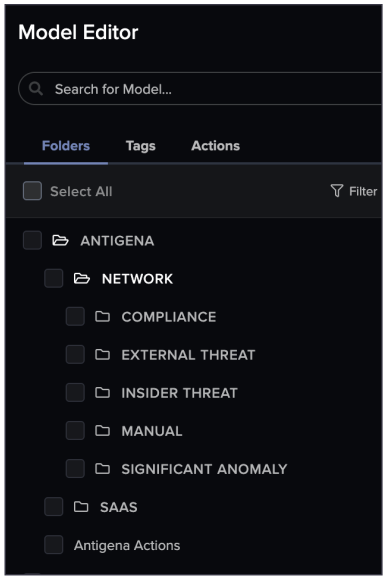
Model Categories

Darktrace RESPOND/Network and Endpoint responses are triggered by model breaches within the Threat Visualizer - Darktrace RESPOND models look for specific behavior or for indicators triggered by other model breaches. Darktrace models are a series of logical statements and conditions which, if met, trigger an alert or action; models are primarily used to identify and alert on anomalous and potentially malicious behavior. The models framework leverages both the underlying 'pattern of life' detection and outputs from Darktrace Deep Packet Inspection and security integrations and modules.

When Darktrace RESPOND/Network or Darktrace RESPOND/Endpoint are enabled within your Darktrace environment, a new collection of models will become available: **Antigena** > **Network**. This collection contains specific models for unusual device behavior which are set to trigger on specific types of connection or activity and perform different Darktrace RESPOND actions depending on the incident identified.

Reviewing the Darktrace RESPOND (Antigena Network) Models

Open the Model Editor. In the Threat Visualizer interface, the Model Editor is accessible from the main menu under **Models** > **Model Editor** or from any breach log with the " Click to view model" button.



From the left-hand model list, select the **Antigena** > **Network** folder.

In this folder, Darktrace RESPOND models for Network and Endpoint are grouped into four categories: **Compliance**, **External Threat**, **Insider Threat** and **Significant Anomaly**.

- Models inside the **Compliance** folder are linked to potential compliance breaches like unauthorized file storage and TOR usage.
- Models inside the **External Threat** folder contain behavioral elements suggestive of an external threat.
- Models inside the **Insider Threat** folder contain behavioral elements which suggest highly unusual internal or lateral behavior.
- The **Significant Anomaly** folder contains models which are independent of attributable aspects, but are consider significantly anomalous.

When a model breaches that would trigger a Darktrace RESPOND/Network action, the device is checked for the relevant tag before the action is performed. For example, a device performs unauthorized file transfer activity which triggers a compliance breach. If the **Antigena Compliance** models are active, they will check that the device has either the **Antigena Compliance** or **Antigena All** tags before the action is performed.

DARKTRACE RESPOND/NETWORK TAGS

Darktrace RESPOND/Network can take a range of proactive, measured, automated actions in the face of confirmed cyber-threats detected in real time. Darktrace understands the 'pattern of life' of users, devices, and networks and so Darktrace RESPOND can take action in a highly targeted manner, mitigating threats while avoiding over-reactions. Typically, a subset of all devices will be deemed appropriate to place under Darktrace RESPOND control.

Whether a device is eligible to be autonomously actioned by Darktrace RESPOND/Network is defined by the tags applied to that device. There are five Darktrace RESPOND/Network (Antigena) tags, corresponding to the model categories:

- Antigena Compliance
- Antigena External Threat
- Antigena Insider Threat
- Antigena Significant Anomaly
- Antigena All

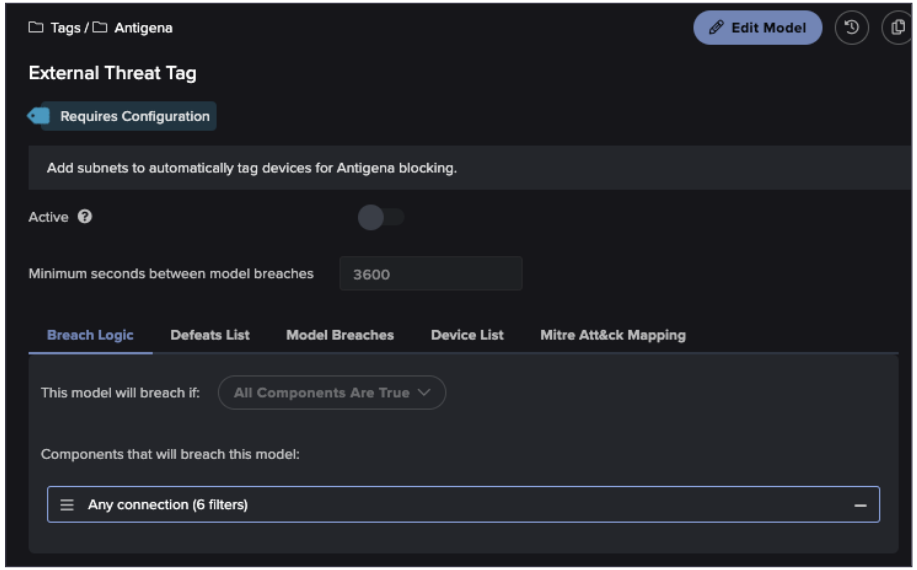
When a device is given the **Antigena Compliance** tag, for example, it becomes eligible to receive Darktrace RESPOND/Network responses triggered by models in the **Compliance** category. The final tag - **Antigena All** - makes the device eligible for responses from all categories.

By default, no devices are tagged. Tags can be applied in three ways: individually, by bulk tagging in **Device Admin**, or by enabling tagging models to provide automatic tagging. The latter - automatic tagging through tagging models - is often the easiest method in large environments.

Automatic Tagging

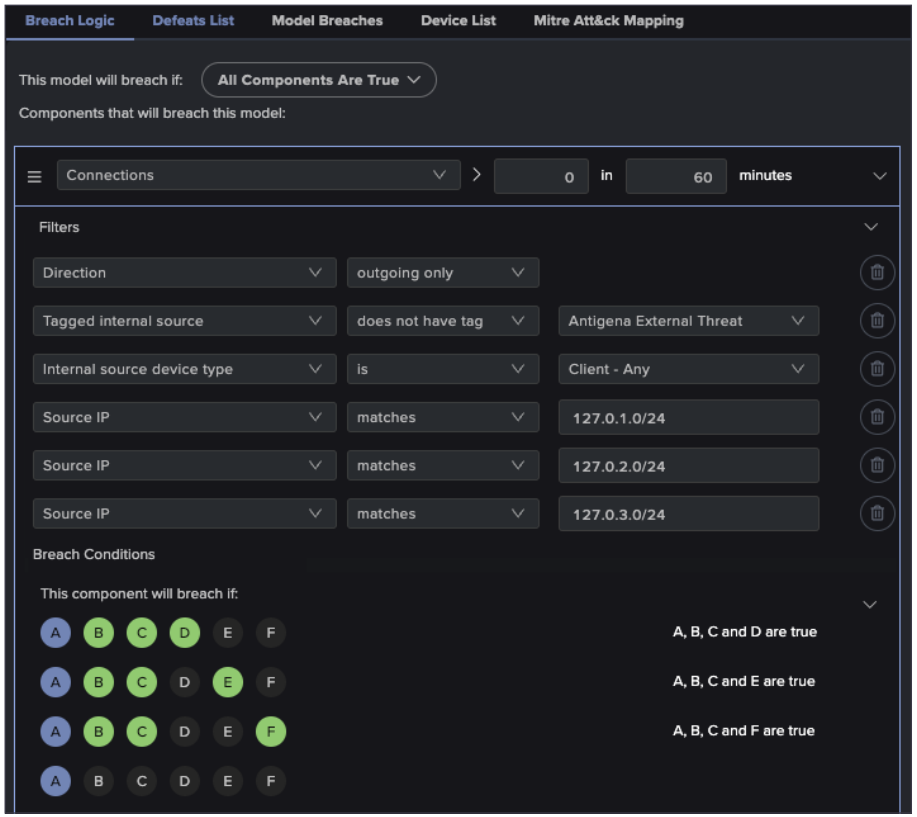
The Model Editor contains a series of models which apply tags automatically based on criteria.

Return to the model editor, but instead select the **Tags** > **Antigena** folder. Inside this folder are models which correspond to each Darktrace RESPOND/Network (Antigena) tag. These models are marked as "**Requires Configuration**", meaning they are templates and must be modified to suit your environment before being enabled. The **External Threat Tag** model will now be used as an example.



This model has one component which looks for connections - expand this component.

The component is comprised of three key filters and three example filters. The key filters identify outgoing connections from client devices which do not already have the **Antigena External Threat** tag. These filters are always required in the breach logic.



Below are some subnet ranges - these filters are templates which demonstrate how subnet ranges can be included for tagging. These should be substituted for the subnets you wish to be automatically enrolled in Darktrace RESPOND/Network (Antigena) External Threat actions, following the example given.

Once the model is enabled, any client devices in these subnets will automatically be tagged and the tag remain active for one week, at which point it will expire and either be immediately re-added or not, if the device no longer meets the criteria. This approach can be used with any Darktrace RESPOND/Network tagging model to slowly expand coverage across all subnets.

Darktrace Threat Visualizer 5.1 introduced targeted blocking for server device types. The client devices filter in this model can be altered or removed to widen the range of tagged device types. Servers tagged with Darktrace RESPOND/Network (Antigena) tags will, however, only be actioned if the relevant setting is enabled on the System Config page.

Please see Darktrace RESPOND/Network and Server Devices for more information.

DARKTRACE RESPOND/NETWORK INHIBITORS

Models can take a series of actions in response to breaching such as triggering an alert, raising the priority of a device or adding a tag. When Darktrace RESPOND is enabled on a deployment, the additional **Darktrace RESPOND** model action becomes available. When Darktrace RESPOND/Network or Darktrace RESPOND/Endpoint are enabled within your Darktrace environment, category of **Darktrace RESPOND (Antigena)** models already using the **Darktrace RESPOND** action are exposed.

This action allows an operator to set an 'inhibitor' - an action to ***inhibit*** the behavior that matches the model criteria. Darktrace RESPOND/Network and Darktrace RESPOND/Endpoint share a set of inhibitors that can be selected.

In some cases, Darktrace RESPOND may be unable to perform the specific inhibitor as set on the model - for example, where the target device is a server and the inhibitor is "quarantine". Similarly, an inhibitor may be upgraded under specific circumstances. For example, the "block matching connections" inhibitor may be upgraded by Darktrace RESPOND if the quantity of outgoing matching connections to block becomes too large, indicating escalating malicious activity that would be better targeted with "block outgoing connections".

Inhibitors

INHIBITOR	DESCRIPTION
Automatic	This option allows Darktrace RESPOND/Network to automatically choose the best option using information gathered from the alert. For example, if the alert concerns suspicious behavior to an SMB share on port 445, it may block connections just to that port. If a range of suspicious connections to various external endpoints was seen, it may instead select to block all external connections.
Block Matching connections	Darktrace RESPOND/Network will block the specific connection and all future connections that match the same criteria. For example, the FTP block prevents all future outbound connections to port 21 from the target device.
Enforce pattern of life	Enforce 'pattern of life' allows a device to make the connections that it usually makes. Therefore, it only allows connections and data transfers which Darktrace considers normal for that device. Any unusual traffic (marked with an unusual message in Darktrace) is blocked.
Enforce group pattern of life	Group 'pattern of life' allows a device to make any connections and data transfers that it or any of its peer group typically make. This action is less strict than the individual 'pattern of life' enforcement.
Quarantine device	All incoming and outgoing traffic from or to a device is blocked.
Block all outgoing traffic	All outgoing traffic from a device is blocked.
Block all incoming traffic	All incoming traffic to a device is blocked.

WORKFLOW FOR REVIEWING DARKTRACE RESPOND ACTIONS

Reviewing and Tuning Darktrace RESPOND/Network

When Darktrace RESPOND/Network is first enabled, reviewing each Darktrace RESPOND action can be a useful way of gaining confidence in the autonomous actions and locating any models which may require further tuning. It is recommended an operator should review any currently active actions first.

After reviewing a Darktrace RESPOND Action or after reviewing all Darktrace RESPOND Actions, you may wish to adjust the parameters of the Darktrace RESPOND Models or tag membership.

Suggested Workflow

1. Identify an action

Open the Darktrace RESPOND Actions menu item and identify an action to investigate.

2. Understand the effect of the Darktrace RESPOND Action on that device or user

Review the description under the Action column.

To see specific connections that have been interfered with, click on the device name under the device column.

3. Understand the conditions of the model that caused the Darktrace RESPOND action

If you are not familiar with the Darktrace RESPOND Model that caused this Action, click on the name of the model to launch the model editor entry.

4. Review the event log for the action

To open the Event Log for the device from within the Darktrace RESPOND Actions page, click 'true' in the blocked column.

5. Identify any Darktrace Model Breaches that prompted the Darktrace RESPOND action

In the Event Viewer (opened in previous step) locate the Darktrace RESPOND Action.

Working down the event log, find Darktrace Alerts that may have prompted the Darktrace RESPOND Action to occur.

6. Understand the conditions of the model that caused the Darktrace Model Breaches

If you are not familiar with the conditions of the Darktrace Model Breach, you can review the model logic by returning to the Model Editor from the main menu.

7. Identify which connections caused the Darktrace Model Breach to occur

Return to the Event Log to review any connections that may have contributed to the Darktrace Model Breach.

8. Consider the risk posed by removing the Darktrace RESPOND Action or by extending the Darktrace RESPOND Action

Are the connections being blocked preventing any business activities?

- If NO then you may wish to leave the Darktrace RESPOND Action in place.
- If YES then consider the anomalous behavior that provoked the Darktrace RESPOND Action.

Are you satisfied that the anomalous behavior reported by Darktrace is benign?

- If YES then you may wish to remove the Darktrace RESPOND Action on the device. You can do this by clicking the Clear button from the Darktrace RESPOND Actions page.

After reviewing the Event Log you may wish to extend the Action for a period of time. (Note that Darktrace RESPOND Actions will automatically extend if further anomalous activity is seen from the device).

9. Optionally consider whether the level of anomaly justified the specific Darktrace RESPOND Action placed on that device in the context of your business logic for that device

DARKTRACE RESPOND/NETWORK DEPLOYMENT SCENARIOS

Tailoring the Darktrace RESPOND/Network Deployment

Tailoring Darktrace RESPOND/Network responses to best fit the network environment can be approached from a number of different angles. Understanding how a Darktrace RESPOND/Network action is triggered on a specified device will ensure that only desired devices will be targeted for autonomous responses. When a device breaches a model contained within one of the Darktrace RESPOND/Network folders, Darktrace RESPOND/Network will look for the corresponding tag on the device before responding.

The 'Antigena All' tag allows Darktrace RESPOND to take action when any default Darktrace RESPOND model breaches, regardless of subfolder. For instance, a model residing within the **Antigena > Network > Compliance** folder will only prompt a Darktrace RESPOND response if it is triggered by a device that either has the tag 'Antigena Compliance' or 'Antigena All'. See the Darktrace RESPOND/Network Models guide for more information.

Default Darktrace RESPOND/Network Models will only fire on devices with one of the Darktrace RESPOND/Network (Antigena) tags applied. For custom models with Darktrace RESPOND responses enabled, it is recommended that the model excludes devices without the appropriate tags. This can be achieved with the following logic, modified as appropriate for the model target:

```
Tagged [Internal Source/Internal Destination] Has Tag [Antigena All/
Antigena Compliance/etc]
```

Deployment Scenarios

Darktrace RESPOND/Network offers granular controls and can be tailored in a number of different ways. The exact methodology for deploying Darktrace RESPOND/Network will depend on the composition of your network - the following provides an outline of simple deployment options.

Key recommendations

- When moving from human confirmation mode to active mode, begin with active mode out of business operating hours using the schedule. Gradually progress to active mode during business operations until 24hr active mode is achieved.
- Frequently review devices eligible for actions and ensure changes in network structure are covered by Darktrace RESPOND response.
- If specific models frequently produce low level alerts, consider tuning the model with one-click defeats or altering the logic to better reflect your organizational requirements.

Each tag has a corresponding tag model which can be used to continually apply tags to new devices which meet the criteria. Editing these models to provide future coverage is explained in Darktrace RESPOND/Network Tags. The scenarios below cover the recommended Darktrace RESPOND tag to apply to the specified device types to produce different levels of configuration. In all cases, substitute 'Any Client Device' for a specific IP range or device type(s) as desired.

Please note, servers tagged with Darktrace RESPOND tags will only be actioned if the relevant setting is enabled on the System Config page. This setting is enabled on deployments initialized from Darktrace Threat Visualizer 5.1 and above. Please see Darktrace RESPOND/Network and Server Devices for more information.

Scenario 1 - Minimum Requirement for Core Protection

DEVICES	ALL	EXTERNAL THREAT	SIGNIFICANT ANOMALY	INSIDER THREAT	COMPLIANCE
Client Devices		✔			
Server Devices		✔			

External Threat includes targeted coverage for ransomware.

Scenario 2 - Expanded Coverage of Client Devices

DEVICES	ALL	EXTERNAL THREAT	SIGNIFICANT ANOMALY	INSIDER THREAT	COMPLIANCE
Client Devices	✔				
Server Devices		✔			

Scenario 3 - Expanded Coverage of Client and Server Devices

DEVICES	ALL	EXTERNAL THREAT	SIGNIFICANT ANOMALY	INSIDER THREAT	COMPLIANCE
Client Devices	✔				
Server Devices		✔	✔	~	

***Insider Threat**: Review the model set and apply to appropriate servers only.*

Scenario 4 - Full coverage

DEVICES	ALL	EXTERNAL THREAT	SIGNIFICANT ANOMALY	INSIDER THREAT	COMPLIANCE
Client Devices	✔				
Server Devices	✔				

DARKTRACE RESPOND/NETWORK AND SERVER DEVICES

Autonomous Darktrace RESPOND/Network response controls unusual activity on devices by restricting connectivity, increasing severity from targeting specific endpoints up to full device quarantine.

The business-critical nature of servers creates ideal opportunities for lateral movement, persistence, or for threats such as ransomware to spread rapidly from a single key asset to many client devices. Darktrace RESPOND autonomous response can provide AI-powered defense of these systems, halting unusual activity whilst ensuring normal operations are maintained.

Many phased deployments of Darktrace RESPOND/Network capability have targeted autonomous response at client devices. Automatic eligibility models available for roll-out prioritized the tagging of device types such as laptops and desktop unless actively edited. This approach should be reconsidered to ensure full network coverage.

How it Works

From Darktrace Threat Visualizer 5.1, Darktrace RESPOND/Network autonomous response has been tailored for devices identified as servers. Darktrace RESPOND will never fully quarantine server devices (unless manually overridden by configuration) and will only block targeted connections. Darktrace RESPOND will not block incoming connections from client devices that are deemed to be a normal part of the server's pattern of life' to ensure business continuity whilst unusual activity is blocked.

Before 5.1, servers may have experienced Darktrace RESPOND/Network control under specific circumstances:

- When server devices had been observed by Darktrace for fewer than 24 hours, they were eligible for RESPOND/Network actions to allow for early device-type changes.
- When the **Block Server Devices** option was manually enabled on the System Config page.
- When server devices were indirectly impacted by actions taken through Active Integrations, such as firewalls.

From 5.1, "Block Server Devices" will be enabled by default on new deployments. Existing deployments will need to enable the setting on the Darktrace System Config page, and a warning will be shown in the Threat Visualizer if server devices are tagged for Darktrace RESPOND actions but the setting is disabled.

Inhibitors

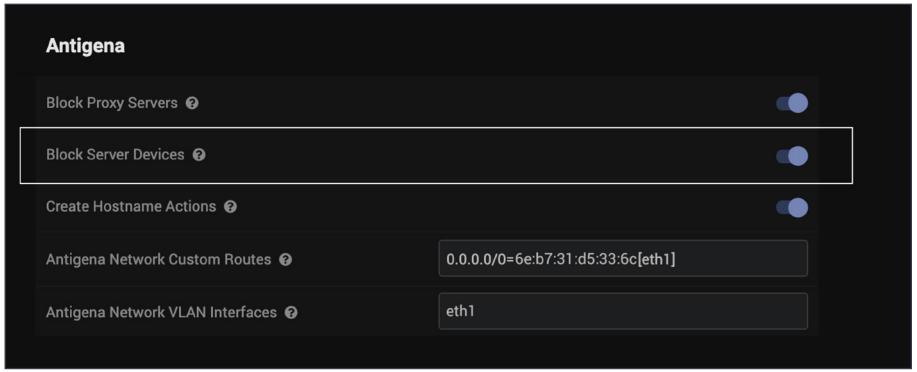
Server devices are eligible for existing Darktrace RESPOND/Network inhibitors with some minor modifications. These modifications will be clear in the final action displayed on the Darktrace RESPOND Actions page.

- In the default configuration, **quarantine** actions applied by a model will be downgraded to **block all outgoing traffic**.

  Configuration options are available to upgrade quarantine actions to full quarantine, or prevent quarantine actions (including downgraded actions) entirely.

- Darktrace RESPOND will only respond to incoming connectivity to a server device in specific circumstances:
  - When a model applying the specific **Block matching connections** inhibitor triggers on incoming traffic to the server device.
  - Where Darktrace RESPOND deems a specific connection from a client to a server is most effectively blocked by targeting the incoming connection on the server end. For example, if a client uses a proxy and the client is unreachable, Darktrace RESPOND will target the incoming connection from the client to the proxy. This case is highly specific and occurs infrequently.

Configuration



1. Access the Darktrace Threat Visualizer on the master instance and navigate to **System Config** page (**Main menu > Admin**).

2. Choose **Modules** from the left-side menu and confirm that the **Darktrace RESPOND/Network** entry is turned on and licensed.

3. Choose **Settings** from the left-side menu and locate the **Darktrace RESPOND/Network** subsection.

   Ensure the **Block Server Devices** setting is turned on.

   Save the changes.

Servers and Darktrace RESPOND/Network (Antigena) Tags

To experience Darktrace RESPOND/Network actions, server devices must be tagged with one of the Darktrace RESPOND/Network (Antigena) tags. Please see Darktrace RESPOND/Network Tags for more information.

Default models which control Darktrace RESPOND/Network (Antigena) tags will continue to target client devices by default; this can be altered by changing the “Client - Any” filter or by following the recommended steps in Darktrace RESPOND/Network Deployment Scenarios.

CHAPTER 11 - THE MOBILE APP

GETTING STARTED

REGISTERING THE APP AS A USER 193
DASHBOARD VIEW 194

CYBER AI ANALYST

AI ANALYST VIEW 195
AI ANALYST INCIDENT OVERVIEW 196

DEVICES AND MODELS

ALERT VIEW 199
MODEL DETAILS VIEW 202
DEVICE DETAILS VIEW 203
BREACH DETAILS VIEW 206

DARKTRACE RESPOND

DARKTRACE RESPOND ACTIONS VIEW 209
ACTION DETAILS VIEW 210

DARKTRACE/EMAIL

DARKTRACE/EMAIL OVERVIEW VIEW 212
DARKTRACE/EMAIL EMAIL LOGS VIEW 215
DARKTRACE/EMAIL EMAIL DETAILS VIEW 218

OTHER VIEWS AND SETTINGS

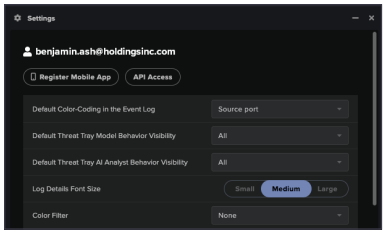
APP SETTINGS 220

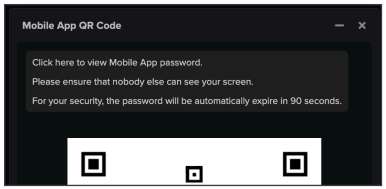
REGISTERING THE APP AS A USER

The Darktrace mobile app, an increasingly popular option that our customers are utilizing in order to get alerts whilst away from the office, can be registered from the Account Settings page. A configured mobile app Service is presumed - if this is not configured, please see Configuring the Mobile App for the setup process.

The following setup process requires access to the Threat Visualizer or SaaS Console. For Darktrace/Email customers who do not have access to these interfaces, the relevant settings will be made available in a future update.

How to Register

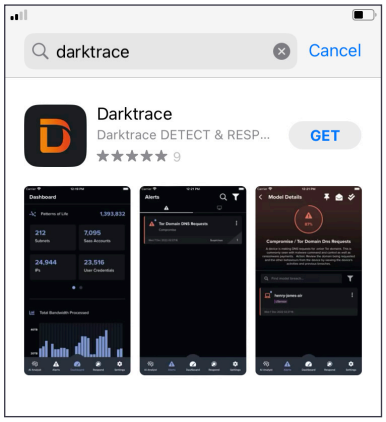




1. Navigate to **Account Settings** from the main menu of the Threat Visualizer or SaaS Console.

   The **'Register Mobile App'** option should now be available. Click **'Register Mobile App'** to reveal a QR code.

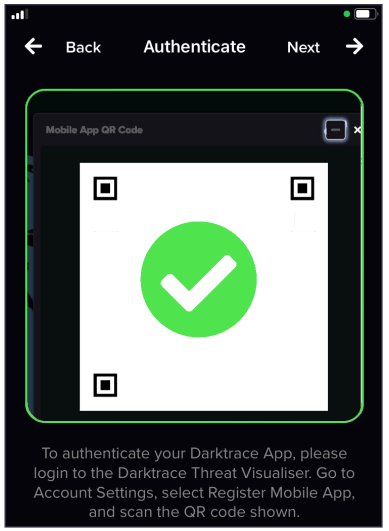
   For SSO users, this QR code dialogue will also display a passphrase section.

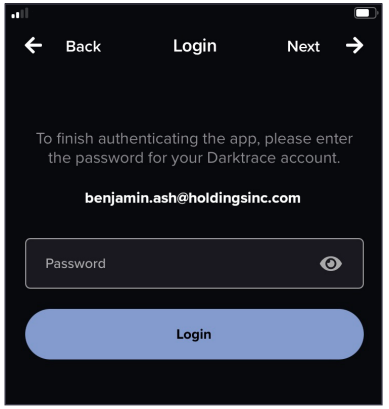


2. On your smartphone, open the appropriate app store and search for **Darktrace**.

   The Darktrace iOS app is available on the App Store and the Android app is available on the Google Play store.

   Download the Darktrace app.







3. Open the Darktrace app.

   You will be prompted to authenticate with the Darktrace instance. Press **'Next'** to proceed.

   The app will request permission to use your smartphone camera. **Use the camera to scan the QR code** in the Threat Visualizer generated above.

4. After the QR has been scanned, you must provide a password to finalize authentication.

   For Threat Visualizer and LDAP users, enter the password used to log into the Threat Visualizer.

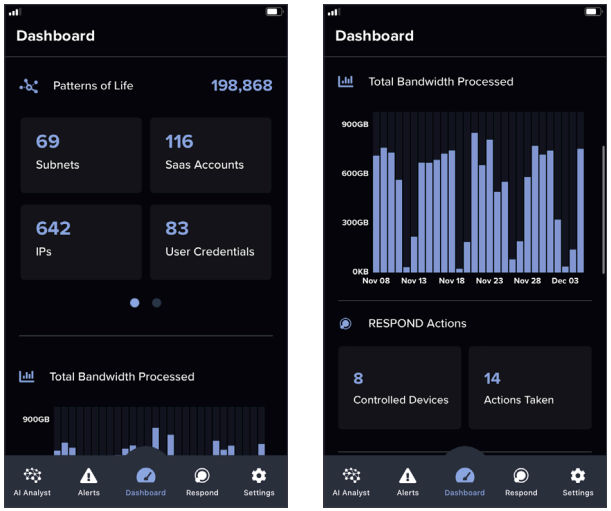
   For SAML SSO users, enter the passphrase displayed in the same pop-up as the QR code in the password field. This passphrase is only valid for 90 seconds. Ensure that all characters including hyphens are entered correctly.

   The app should now authenticate.

   Move between the screens for a guide overlay explaining the functionality. The guide can be re-enabled at any point from the settings page.

DASHBOARD VIEW

The Darktrace Mobile App provides a streamlined Threat Visualizer experience for on-the-go investigation and response. Investigate ongoing anomalous activity highlighted by Darktrace DETECT, approve autonomous RESPOND actions pending human confirmation, share alerts with colleagues and be notified when critical priority AI Analyst incidents are created - on the go, wherever you are.



After authentication, new users will land on the **Dashboard**. This summary mirrors the high-level Darktrace DETECT and Darktrace RESPOND information available on the summary panel on the Threat Visualizer workspace. The page will refresh periodically - pull down on the summary area to trigger a manual refresh.

Device counts display the number of entities - users detected in external platforms (SaaS users), devices, credentials - that Darktrace has observed and is actively modelling a pattern of life for. Statistics are calculated over 7 days (***IPs***[1], ***Clients***, ***Servers***, ***Devices***, and ***SaaS Accounts***), 28 days (***Credentials***) or 12 weeks (***Patterns of Life***[2]).

The total bandwidth processed over the last 28 days is displayed for both network devices and, if present, Darktrace/Endpoint. If Darktrace RESPOND/Network is enabled, statistics on the number of actions taken and devices actioned will also be displayed.

AI Analyst investigates anomalies detected by the Darktrace system at machine speed and surfaces only those that need human attention. The summary displays the equivalent number of hours it would take a human analyst to perform the same investigations. Of those incidents surfaced, the type of activity they represent is also broken down by category.

[1] Peak unique values in any 24hr period over last 7 days.

[2] A detailed description of how the "Patterns of Life" value is calculated is available in Advanced Summary Panel Info from the main Threat Visualizer guide.

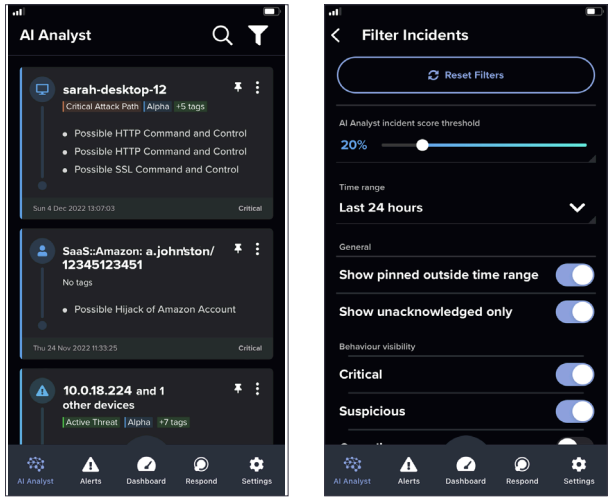
The next step is to review the AI Analyst page - where incidents are created for the most interesting and high priority activity observed by Darktrace.

AI ANALYST VIEW

Incidents are created from AI-powered investigations into anomalies across your digital environment. Every time the conditions for a model are met and a model breach is created, AI Analyst investigates the activity and concludes whether it needs to be surfaced for human analysts to review. The results of this analysis is then grouped in incidents.

Darktrace Mobile App 4.0 onwards uses the intelligent incident grouping method introduced in Darktrace Threat Visualizer 5.2.

Incidents



The **AI Analyst** page displays recent AI Analyst incidents. By default, incidents are sorted chronologically and filtered to show only critical incidents with a score over 20% created in the last 24 hours. Tap the filter icon to change the sorting, filters, or whether pinned or acknowledged incidents should be included in the results.

OPTION	DESCRIPTION
Search	Search the incidents.
Filters	Change the filters, order, time frame and type of incidents that appear.

Incident Tiles

The heading of each entry indicates the initial device - the device that triggered the first activity AI Analyst detected - and if applicable, the number of additional devices involved in the incident. Device details such as device type (indicated by icon), hostname and tags are also displayed for quick reference.

Underneath the heading, the tile lists the title of each event the incident comprises. The time the first event was created is displayed in the bottom left and the behavior category associated with the incident is indicated in the bottom right.

Pinned incidents are always returned regardless of whether they fall within the time range chosen, unless explicitly filtered out. Pinned incidents are indicated by the pin icon.

Filtering options closely mirror those of the Threat Visualizer Threat Tray - please refer to Filtering Alerts with Behavior Categories, Filtering Alerts with Basic Filters and Filtering Alerts with Advanced Filters for more detailed explanations of each of these settings.

Interactions

The following interactions are available:

- Swipe right to pin or unpin the incident.
- Swipe left to acknowledge the incident and all underlying events.
- Tap the three dots to see actions for this incident.
- Tap on an incident to review.

Incident Overview

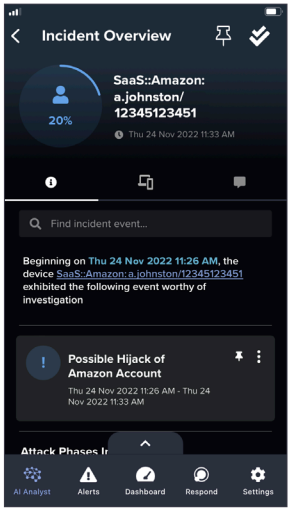
If an AI Analyst incident is tapped from this view, the **Incident Overview** will open. Incident Overview and subsequent views will now be covered in AI Analyst Incident Overview.

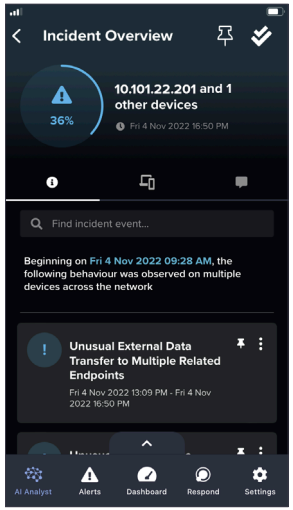
AI ANALYST INCIDENT OVERVIEW

Incident Overview

Now, tap on an individual incident tile to open the overview. Two icons are displayed in the top right.

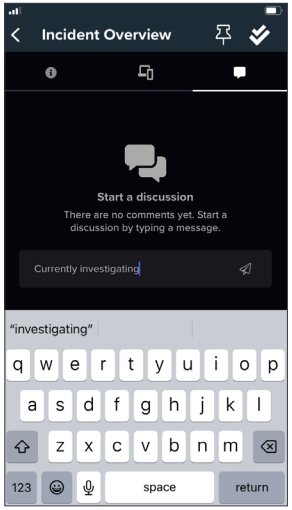
OPTION	DESCRIPTION
Pin / Unpin	Pin or unpin the incident.
Acknowledge All	Acknowledge this incident and all underlying incident events.

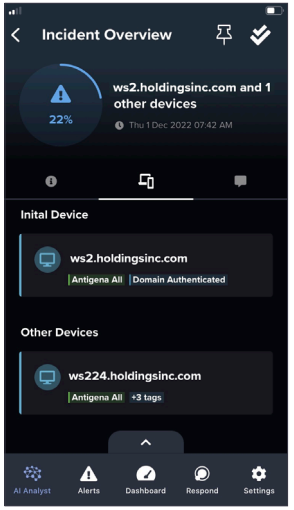




The top of the overview displays the overall incident score from 0-100%, the initial device and number of devices involved, and the date the first event was created. Underneath are three tabs: incident info, devices and discussion.

Incident Overview Tabs





Discussion & Devices

Discussion

The **discussion** tab can be used to post comments on the incident for other Threat Visualizer and Mobile App users to review.

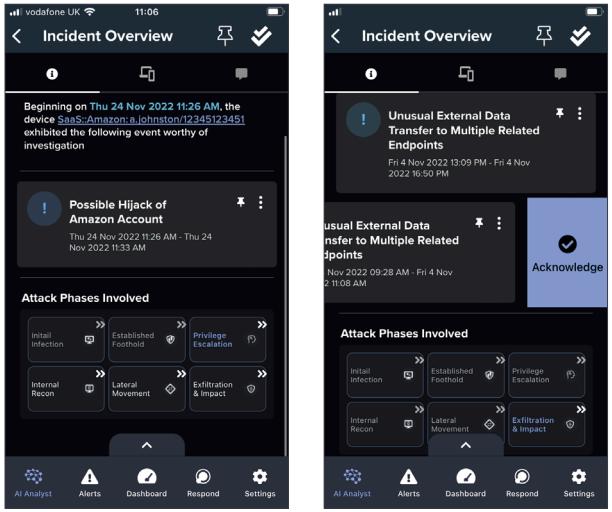
Comments remain within your Darktrace environment - to discuss a model breach with Darktrace analysts, please use Ask the Expert or the Darktrace Customer Portal.

Devices

The **devices** tab details the devices that triggered each incident event in the AI Analyst incident. The device type, hostname and any tags are displayed.

Tap on an individual device to switch to the device details and review model breaches or active RESPOND actions for the device. Device details is covered in Device Details View.

Incident Info



An incident is composed of one or more events; events are a group of anomalies or activity investigated by Cyber AI Analyst that it considers high priority enough to raise to a human operator for attention. The **incident info** tab displays details of these incident events; each event describes the type of activity found and the time frame over which the activity was identified. Individual events can be pinned - pinned events are indicated by the pin icon.

Below the list of events, the **Attack Phases Involved** visualization associates those events with the stages of a cyber attack they likely represent.

Interactions

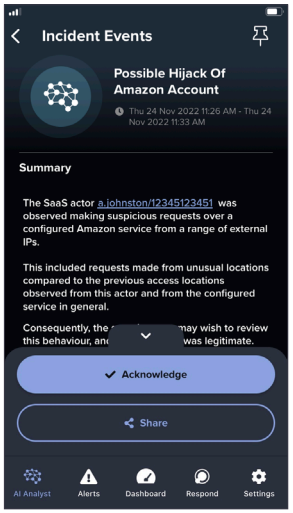
The following interactions are available:

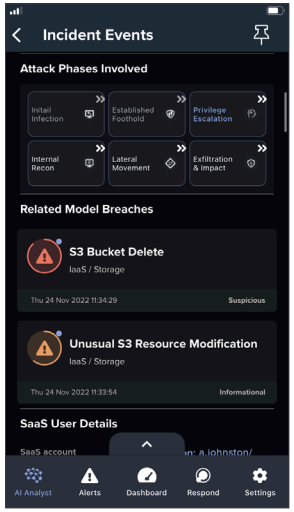
- Swipe left to acknowledge the specific incident event.
- Tap the three dots to see actions for this incident.
- Tap the up arrow to review more actions:
  - Tap the share icon to create a short message with a link to the incident in the Mobile App and in the Darktrace Threat Visualizer.

    The share option can be used to make a note for later follow-up or to email a colleague with the request that they investigate further.
- Tap on an event to review more details on the incident events page.

Incident Events

Tap on an individual event to review. Each incident is comprised of events which detail a specific anomalous activity - or chain of activities - investigated by AI Analyst. The “Incident Events” page gives a detailed overview of each of these events. Where there are multiple events, swipe left or right to move between the event screens.

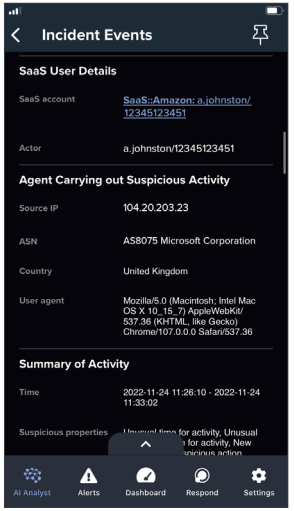


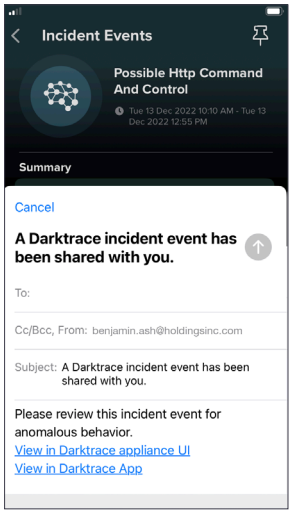


OPTION	DESCRIPTION
/ Pin / Unpin	Pin or unpin the incident.

The event title and the time period over which the activity was detected will appear at the top. Underneath, the summary describes the type of activity AI Analyst detected, the device involved, any endpoints connected to and an explanation of why the activity is anomalous. The summary is the best way to quickly understand the activity AI Analyst has flagged for human review, and can be localized into any of the current twelve language options in the global settings.

Attack Phases Involved associates the event activity with the stages of a cyber attack it may likely represent. Below, the model breaches that triggered the initial investigation are listed. AI Analyst is not reliant on model breaches, they primarily provide a starting point for its deeper analysis of the device activity around the anomaly time. Each model breach can be tapped to see further details - these details are covered in Breach Details View.





Depending on the event type, multiple sections containing contextual information will appear. These may include connection information, devices triggering connections, transferred file information.

ALERT VIEW

Interactions

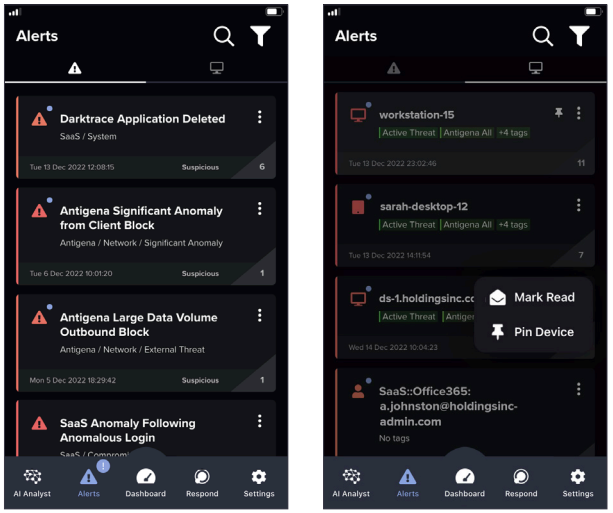
The following interactions are available:

- Swipe left or right to move between the event screens.
- Tap a related model breach to see more details.
- Tap the name of a device that appears in the summary or event details so see more detailed device info.
- Tap the ^ up arrow to review more actions:
  - Tap the share icon to create a short message with a link to the incident in the Mobile App and in the Darktrace Threat Visualizer.
  - Tap the ✓ check icon to acknowledge this specific event. Events can be acknowledged individually, rather than as part of the whole incident - this is a common workflow as a threat analyst investigates each underlying event in turn.
- Tap the pin icon to pin the incident to the top of the Cyber AI Analyst incident list; pinned events will display a filled icon. Pinned incidents will always appear at the top of the list and will persist until unpinned.

The alerts page contains Darktrace model breaches, grouped by model or device. AI Analyst already investigates every model breach that occurs on the system and reports only the most interesting activity as incidents. Users who wish to dive down deeper and investigate specific behaviors after investigating AI Analyst incidents can move to this view of the Mobile App.

"Alerts" combines the "Models" and "Devices" screens from the previous Mobile App version.

Alerts



The alerts page displays recent Darktrace Model Breaches. By default, alerts are sorted by threat score and filtered to show only critical or suspicious model breaches with a score over 60% created in the last 24 hours. Tap the filter icon to change the sorting, filters, or whether acknowledged model breaches should be included in the results.

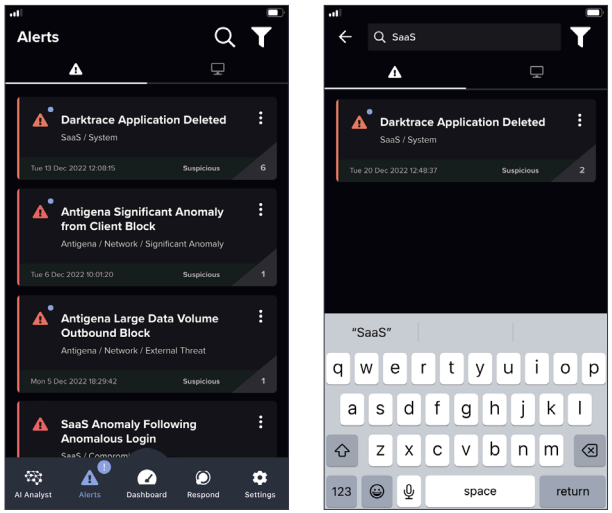
OPTION	DESCRIPTION
Search	Search the model breaches.
Filters	Change the filters, order, time frame and type of model breaches that appear.

The alerts page has two tabs: ⚠ models and 🖵 devices. For devices with current model breaches **and** AI Analyst incidents - the devices tab will display both from the same view. This is not applicable to the models view.

Any defined filters are applied to both tabs. Filtering options closely mirror those of the Threat Visualizer Threat Tray - please refer to Filtering Alerts with Behavior Categories, Filtering Alerts with Basic Filters and Filtering Alerts with Advanced Filters for more detailed explanations of each of these settings.

Models Tab

In model view, the heading of each entry displays the name of the Darktrace model that was triggered. The color of the ⚠ model breach icon and the left bar indicate severity - a yellow line and icon indicate a low score and red a high score. Models with unread model breaches are indicated by a blue dot.



The time of the most recent model breach for that model is displayed in the bottom left. On the bottom left of the tile are the behavior category associated with the incident and the number of alerts for the model.

Pinned model breaches are always returned regardless of whether they fall within the time range chosen, unless explicitly filtered out. Pinned model breaches are indicated by the 📌 pin icon.

Interactions

The following interactions are available:

- Swipe right to mark the model breach as read or unread.

  The main alerts view can be filtered to prioritize (display first) or only display unread model breaches.

- Tap the ⋮ three dots to see actions for this model.

- Tap on a model tile to review the model breach(es) in more detail.

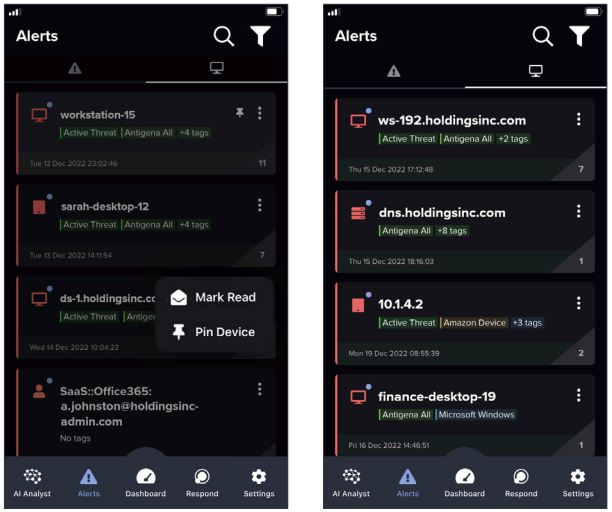
Model Details

If a model breach is tapped from the model view, the **Model Details** page will open.

Model Details is covered in Model Details View.

Devices Tab

In devices view, the heading of each entry displays the identifier of the device that has triggered the model breaches; the device may be identified by hostname, IP or label. The color of the device icon and the left bar indicate severity - a yellow line and icon indicate a low score and red a high score. The icon type also corresponds to the device type - for example, indicates a desktop device. Devices with unread model breaches are indicated by a blue dot.



The time of the most recent model breach for that device is displayed in the bottom left and the number of model breaches on the right. If the device has triggered AI Analyst incidents during the timeframe, these are also included in the count.

Pinned devices are always returned regardless of whether they fall within the time range chosen, unless explicitly filtered out. Pinned model breaches are indicated by the pin icon.

Interactions

The following interactions are available:

- Swipe right to mark the device and underlying model breaches as as read or unread.

  The main alerts view can be filtered to prioritize (display first) or only display unread model breaches.

- Tap the three dots to see actions for the device.

- Tap on a device tile to review the model breach(es) for that device in more detail.

Device Details

If a model breach is tapped from the device view, the **Device Details** page will open. Unlike model details, device details includes both model breaches and AI Analyst incidents.

Device Details is covered in Device Details View.

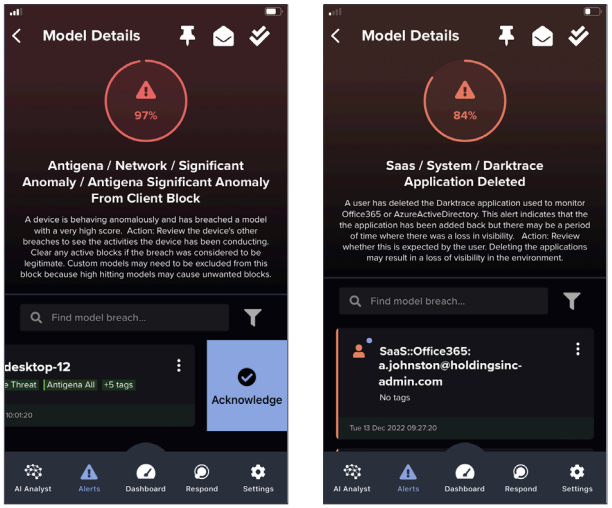
MODEL DETAILS VIEW

The **Model Details** view can be opened from a number of locations in the mobile app including:

- Tapping a model name from the Alerts page in model view
- Tapping the "Go to model" button from a RESPOND action
- Tapping the "Go to model" button from the Breach Details view

The most common workflow to access this view is from the alerts page in model view. The alerts page was covered in more detail in Alert View.

Model Details



Three icons are displayed in the top right:

OPTION	DESCRIPTION
Pin / Unpin	Pin or unpin all model breaches for this model.
Acknowledge All	Acknowledge all model breaches for this model.
Mark as Read / Unread	Mark all model breaches for this model as read or unread.

Model details shows all the recent breaches for the selected model (depending on filter options) and a short description of the model itself.

The central model breach icon displays the highest threat score of all underlying model breaches for the model. Below this icon is the full model name and a short description of the type of activity detected by it.

Each model breach for the model appears as a separate entry. Unread model breaches are indicated by a blue dot. Entries can be filtered or searched. Tap the filter icon to filter the results using the same filters as the main alerts page - this filter will not persist if the view is closed.

The heading of each entry indicates the device that triggered the model breach. Device details such as device type (indicated by icon), hostname and tags are also displayed for quick reference. The color of the device icon and the left bar indicate severity - a yellow line and icon indicate a low score and red a high score. The time the model breach was triggered also appears in the bottom left of the tile.

Interactions

The following interactions are available:

- Tap any of the three icons in the top right as described above - / **Pin / Unpin**, **Acknowledge All** or **Mark as Unread**.
- Tap the filter icon to filter the results.
- Filter the results with the search.
- Swipe left on a model breach entry to acknowledge or unacknowledge that entry.
- Swipe right on a model breach entry to mark that entry as read or unread.
- Tap the three dots on a model breach entry to see actions for that entry.
- Tap on a model breach entry to review more details on the breach details page for that entry.

Breach Details

The breach details view can be opened from multiple locations in the mobile app. Tapping on a model breach entry from model details will open this view.

Breach Details is covered in Breach Details View.

DEVICE DETAILS VIEW

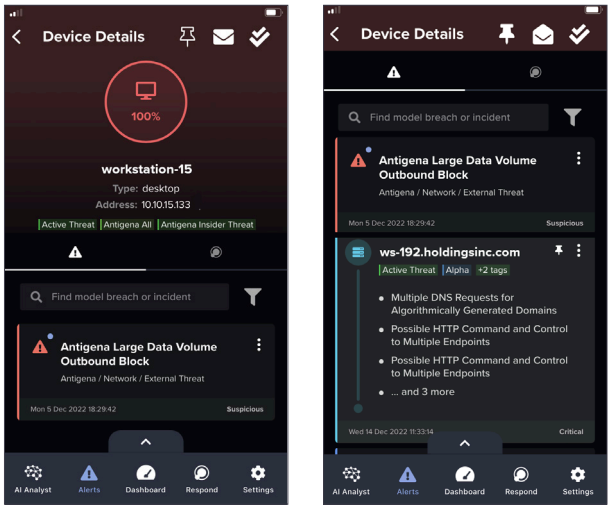
The **Device Details** view can be opened from a number of locations in the mobile app including:

- Tapping a device name from the **Alerts** page in devices view
- Tapping a device name from an AI Analyst event
- Tapping the “Go to device” button from the Breach Details view

The most common workflow to access this view is from the alerts page in devices view. The alerts page was covered in more detail in Alert View.

Device Details

Device view shows all the recent model breaches - and, if applicable - AI Analyst incidents for the selected device (depending on filter options). Darktrace RESPOND actions are also displayed in a separate tab.



The central device icon displays the highest threat score of all underlying model breaches for the model - the icon itself corresponds to the device ***type***. Below are details of the device such as the type, IP address, hostname and any tags applied.

There are two tabs: Alerts and Darktrace RESPOND.

Three icons are displayed in the top right:

OPTION	DESCRIPTION
Pin / Unpin	Pin or unpin this device device.
Acknowledge All	Acknowledge all model breaches for this device. This will not impact any AI Analyst incidents.
Mark as Read / Unread	Mark all model breaches for this device as read or unread. This will not impact any AI Analyst incidents.

Interactions

The following interactions are available:

- Tap any of the three icons in the top right as described above - Pin / Unpin, **Acknowledge All** or **Mark as Unread**.

  Please note, these options will not impact any AI Analyst incidents associated with the device.

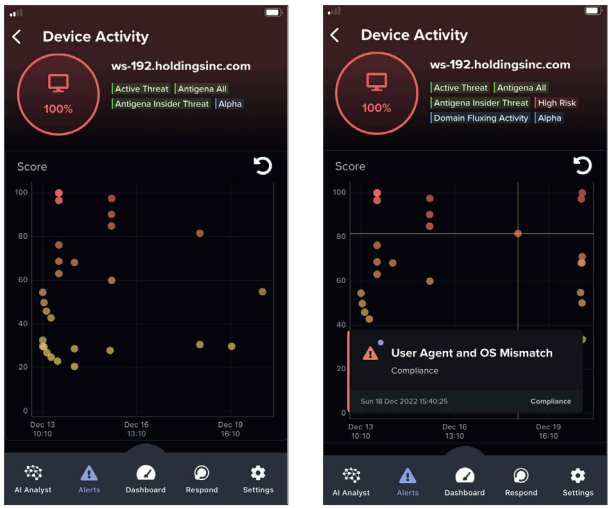
- Tap any of the tabs to review different information about the device - Alerts or Darktrace RESPOND

- Tap the up arrow to review more actions:

  - Tap the share icon to create a short message with a link to the device in the Mobile App and in the Darktrace Threat Visualizer.

  - Tap the chart icon to open the Device Activity view (described below).

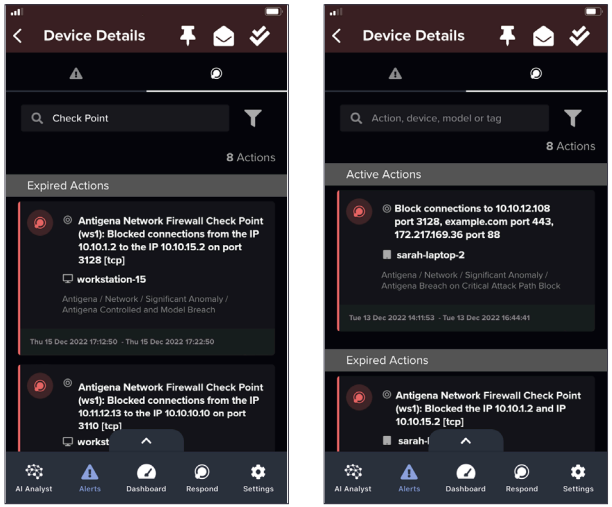
Device Activity



The device activity graph displays all model breaches for the given device. The y axis and dot color represents the severity score of each model breach. Tap on a dot will produce a summary pop-up - tap the summary to open the Breach Details View for the model breach.

Drag with two fingers to zoom. Return to the original view with the arrow.

Darktrace RESPOND tab

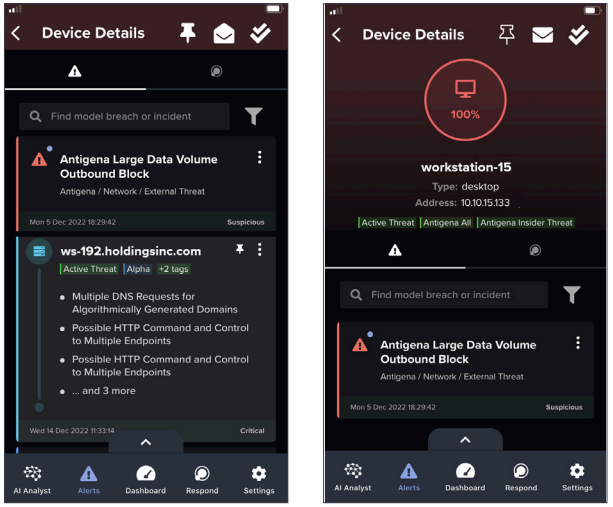


For each device, Darktrace RESPOND actions can be quickly reviewed from this tab. RESPOND actions can be filtered to show a longer time period or show only those associated with a model breach above a certain score. The sorting order - by default, oldest expiring action first - can also be altered.

Each entry represents an individual action - the number of actions associated with the device over the timeframe appears in the top right. The color of the RESPOND icon and left hand bar indicate the severity of the model breach that triggered the action. In this case, this severity will match that of the model breach under investigation. The entry details describe the device, action triggered, and the start and end time of the action itself.

To see further details on the action or to interact with the action, tap on the entry to open Action Details. This page is covered in more detail in Action Details View.

Alerts Tab



The alerts tab combines both AI Analyst and Model Breach alerts for the device.

Each model breach or AI Analyst incident for the device appears as a separate entry. Unread model breaches are indicated by a blue dot. Entries can be filtered or searched. Tap the filter icon to filter the results using the same filters as the main alerts page - this filter will not persist if the view is closed.

Model Breaches

The heading of each model breach entry indicates the model that was triggered. The color of the model icon and the left bar indicate severity - a yellow line and icon indicate a low score and red a high score. The time the model breach was triggered appears in the bottom left of the tile and the model threat category appears on the right.

Tapping on a model breach will open the Breach Details page. This page contains detailed information about the model breach and is covered in Breach Details View.

AI Analyst Incidents

The heading of each entry indicates the initial device - the device that triggered the first activity AI Analyst detected - and if applicable, the number of additional devices involved in the incident. Device details such as device type (indicated by icon), hostname and tags are also displayed for quick reference.

Underneath the heading, the tile lists the title of each event the incident comprises. The time the first event was created is displayed in the bottom left and the behavior category associated with the incident is indicated in the bottom right.

Tapping on an incident from this view will open the **Incident Overview**. Incident Overview and subsequent views will now be covered in AI Analyst Incident Overview.

Interactions

- Tap the filter icon to filter the results.
- Filter the results with the search.
- Swipe left on a model breach entry or AI Analyst incident to acknowledge or unacknowledge that entry.
- Swipe right on a model breach entry to mark that entry as read or unread.
- Tap the three dots on a model breach or AI Analyst incident entry to see actions for that entry.
- Tap on a model breach entry or AI Analyst incident to review more details for that entry.

Breach Details

The breach details view can be opened from multiple locations in the mobile app. Tapping on a model breach entry from device details will open this view.

Breach Details is now covered in Breach Details View.

BREACH DETAILS VIEW

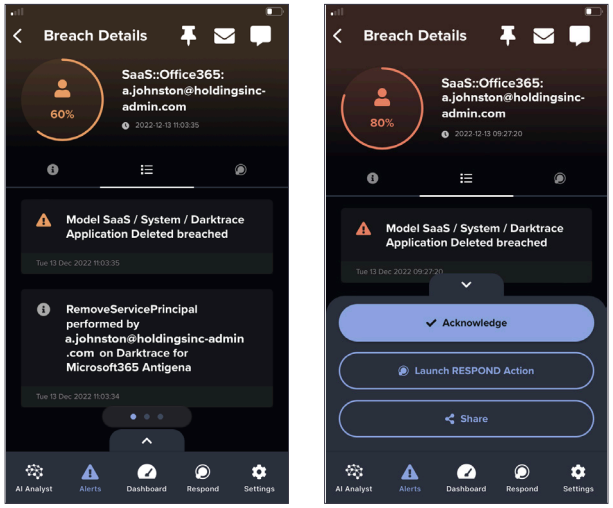
The Breach Details view can be opened from a number of locations in the mobile app including:

- Tapping a related model breach from an AI Analyst event
- Tapping a model breach entry from the device details or model details view.
- Tapping a related model breach from a Darktrace RESPOND action

This is a key view for performing preliminary investigations into model breach activity and reviewing any associated Darktrace RESPOND actions.

Breach Details

Tap on an individual model breach entry to open the breach details.



The top of the overview displays the overall model breach score from 0-100%, the triggering device and device details, and the date the model breach was created. Underneath are three tabs: Model Breach info, Model Breach Event Log and Darktrace RESPOND.

Three icons are displayed in the top right.

OPTION	DESCRIPTION
Pin / Unpin	Pin or unpin this model breach.
Mark as Read / Unread	Mark this model breach as read or unread.
Comment	Review the discussion for this model breach and add comments.

Interactions

The following interactions are available:

- If there are multiple model breaches, swipe left or right to move between the model breaches.

  This is only valid if the breach details view was accessed from device details or model details, where there may be multiple model breaches associated with the device or model.

- Tap any of the three icons in the top right as described above - Pin / Unpin, Mark as Read / Unread or Comment.

- Tap the up arrow to review more actions:

  - Tap the check icon to acknowledge this model breach.

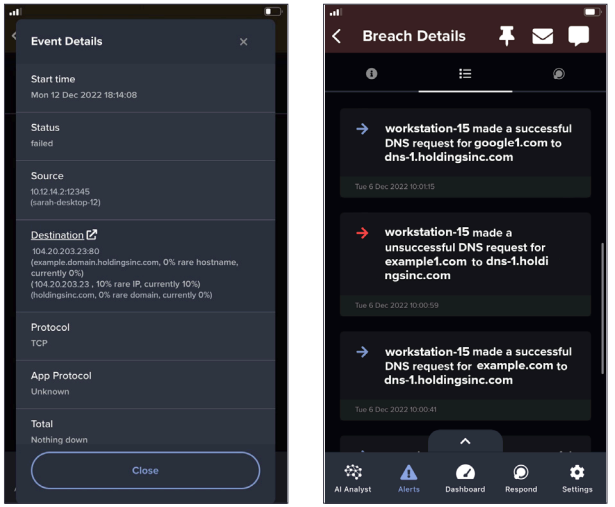
  - Tap Launch Darktrace RESPOND to create a manual Darktrace RESPOND action.

    This action will only appear for users created by a valid Darktrace/Apps, Darktrace/Cloud or Darktrace/Zero Trust module.

  - Tap the share icon to create a short message with a link to the model breach in the Mobile App and in the Darktrace Threat Visualizer.

- Tap any of the tabs to review different information about the model breach - Model Breach info, Model Breach Event Log or Darktrace RESPOND

Model Breach Event Log

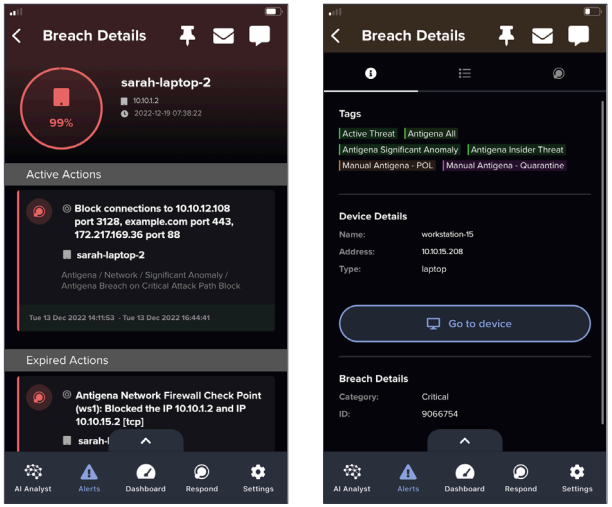


The Model Breach info tab details the events and activity that triggered the model breach. Darktrace RESPOND actions associated with the model breach may also appear here.

For example, for a model breach relating to unusual SSH connectivity, the model breach event log might display a single anomalous SSH connection (or possibly more than one). In contrast, for a model breach relating to excessive HTTP errors or port scanning, the event log will display far more events, because for these types of behavior it is a number of events in combination that all contribute to the detection of an anomaly. For most models, the event log will contain a small number of entries that are relevant.

Tap on any of the breach details to produce a pop-up with further information about the connection or event. For external locations such as IP addresses or hostnames, Darktrace will attempt to derive a geolocation; the icon indicates a geolocation is available. Tapping this icon will open a Maps application at the derived coordinates.

Model Breach Info & RESPOND



Darktrace RESPOND

Certain models can also trigger Darktrace RESPOND actions; if a model breach has triggered an action, this will appear in the Darktrace RESPOND tab.

Each entry represents an individual action - one model may trigger multiple actions if, for example, both Darktrace RESPOND/Network and a Darktrace RESPOND firewall integration are in place. The number of actions triggered by the model appears in the top right.

The color of the RESPOND icon and left hand bar indicate the severity of the model breach that triggered the action. In this case, this severity will match that of the model breach under investigation. The entry details describe the device, action triggered, and the start and end time of the action itself.

To see further details on the action, review the action state, or interact with the action, tap on the entry to open Action Details. This page is covered in more detail in Action Details View.

Model Breach Info

This tab displays more detailed information about the user or device that triggered the model breach and the model breach itself. Any tags applied to the triggering device are included alongside details of the device itself such as hostname and type.

Interactions

The following interactions are available:

- Tap the **Go to Device** button to open **Device Details**.
- Tap the **Go to Model** button to open **Model Details**.

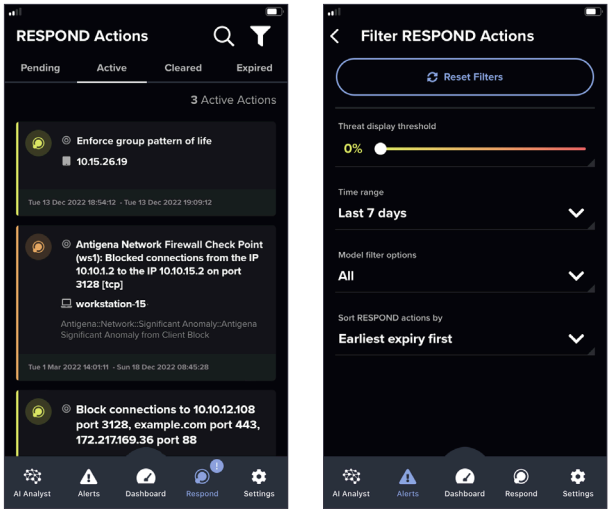
Depending on where the breach details view is opened from, one or both of these buttons may appear.

DARKTRACE RESPOND ACTIONS VIEW

Darktrace RESPOND can take a range of proactive, measured, automated actions in the face of confirmed cyber-threats detected in real time. Darktrace understands the 'pattern of life' of users, devices, and networks and so Darktrace RESPOND can take action in a highly targeted manner, mitigating threats while avoiding over-reactions.

RESPOND Actions

The Darktrace RESPOND actions page lists all current and expired Darktrace RESPOND Actions. Actions from Darktrace RESPOND components such as Darktrace RESPOND/Network, Darktrace RESPOND/Apps and Darktrace RESPOND/Endpoint will appear in this window.

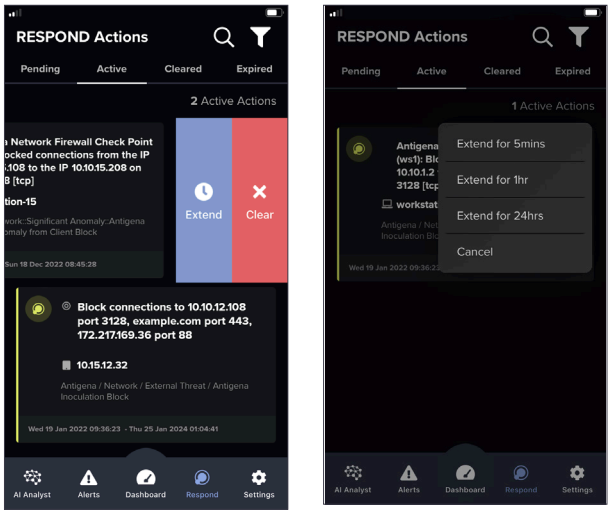


Two icons are displayed in the top right:

OPTION	DESCRIPTION
Search	Search the actions.
Filters	Change the filters, order, or time frame of actions that appear.

By default, actions are sorted to display earliest expiry first and filtered to show the last seven days; tap the filter icon to change the sorting and filters. Filters can be applied to only display actions triggered by model breaches above a certain score or triggered by models within specific folders. No score filter applied by default as manual Darktrace RESPOND actions do not have an associated model breach score and would therefore be filtered out.

Action Tabs



Actions are split into four tabs: pending, active, cleared and expired.

- **Pending Actions** are Darktrace RESPOND Actions that have not yet been approved by a human operator. Actions can be activated from this view.
- **Active Actions** are current Darktrace RESPOND Actions which are in place and have not yet expired.
- **Cleared Actions** are Darktrace RESPOND actions that were manually cleared by an operator. Clear will inform Darktrace RESPOND to stop controlling the device or user and suppress the combination of Darktrace RESPOND Action and model breach conditions for 24 hours.

- **Expired Actions** are previous Darktrace RESPOND actions for devices and users in your environment. Actions can be reactivated and manually extended if desired.

The number of actions under each tab is listed in the top right.

Each tab contains action entries; each entry details the action taken (*inhibitor*), the device or user affected, the model that triggered the action (if applicable) and the start and expiry time of the action. The color of the RESPOND icon and left hand bar indicate the severity of the model breach that triggered the action.

Interactions

- Tap the filter icon to filter the results.
- Filter the results with the search.
- Swipe left on a **pending** action to activate it.
- Swipe left on a **cleared** action to reactivate it for the chosen time period.
- Swipe left on an **expired** action to reactivate it for the chosen time period.
- Swipe left on an **active** action to show further options:
  - Tap on **Extend** to extend the action for the chosen time.
  - Tap Clear or swipe fully to clear it.
- Tap on an action entry from any tab to review more details on the **action details** page.

Action Details

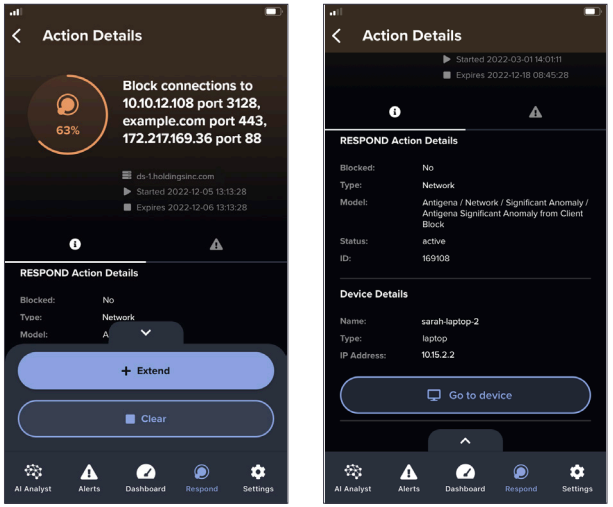
The action details view can be opened from multiple locations in the mobile app. Tapping on an action from the RESPOND actions page will open this view.

Action Details is covered in Action Details View.

ACTION DETAILS VIEW

Action Details

The action details page can be opened from a number of locations in the mobile app. There are two tabs: action info and model breach info.



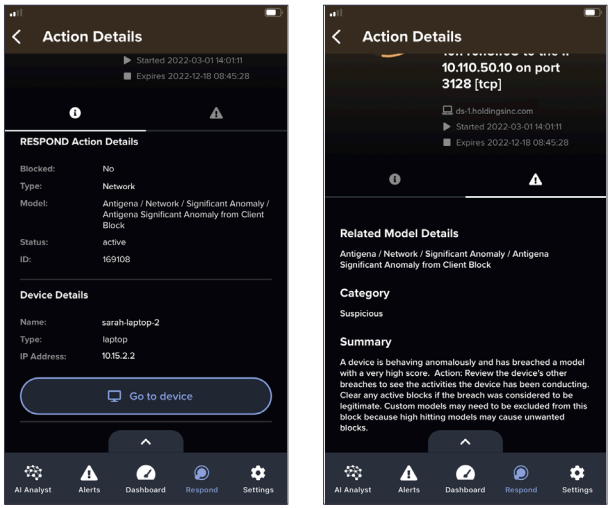
The top of the overview describes the *inhibitor* - the type of action RESPOND has taken - and if triggered by a model breach, the score of the associated model breach. Manual actions created by users do not have a score associated as they are not triggered by a model breach. Underneath the action are details of the device impacted (or potentially impacted in the case of pending actions), the start time of the action and the expiry of the action

Interactions

The following interactions are available:

- Tap the ^ up arrow to review more actions:
  - Tap ▶ Activate to reactivate an expired or cleared action.
  - Tap ✚ Extend to extend an active action.
  - Tap ■ Clear to clear an active action.
  - Tap ▶ Activate to activate a pending action.
- Tap any of the tabs to review different information about the action - ⓘ action info or ⚠ model breach info.

Model Breach Info & Action Info



Model Breach Info

The model breach info tab includes details of the model that triggered the Darktrace RESPOND action. The model name, model behavior category and model description are listed.

Tap the **Go to Model** button to open the Model Details view for the model. Model details is covered in more detail in Model Details View.

Action Info

This tab displays more detailed information about the action and the user or device affected by the action. Tap the **Go to Device** button to open Device Details for the affected device.

If the action was triggered by a model breach, the related model breach will appear on the tab - tap this entry to open Breach Details for the model.

DARKTRACE/EMAIL OVERVIEW VIEW

The new Darktrace/Email expansion extends the quick investigation and response capabilities of the Darktrace mobile app to the inbox.

The **Emails** page is the landing page for interacting with Darktrace/Email functionality - investigate, hold or release emails on the go. The email interface includes an overview dashboard, interactive actions, fully searchable email logs, and the key parts of the Darktrace/Email narrative for each email entry.

There are two main views - Emails and Overview.

- The Overview page displays key metrics from the Intent Dashboard of the main Email Console, such as mailflow volumes, most anomalous users, and the number of Darktrace/Email actions taken.

- The Emails page provides full filtered logs - hold emails, release emails, and review the email narrative to quickly understand and action individual emails.

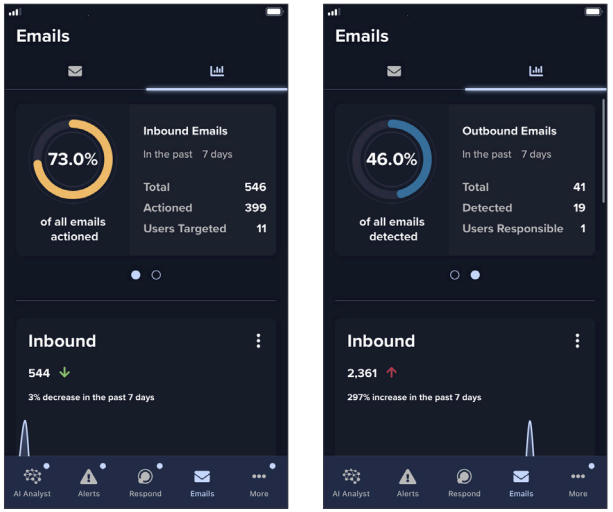
Overview

The Darktrace/Email overview page displays an at a glance overview of the most at-risk users, current threats identified in organizational email traffic and any trends in malicious mailflow over the last seven days. The elements on this page mirror the *Dashboard* and *Intent Dashboard* of the Email Console.

The first element is a summary of actions applied to inbound emails and activity detected in outbound email over the last 7 days. Swipe to switch between views.

Inbound Emails

The percentage of actioned emails contextualizes the number altered in some way by Darktrace/Email against all inbound emails. "Actioned emails" are those that received some reaction by Darktrace/Email that has altered their appearance, content, final location in the inbox, or have triggered an alert to operators of some form.



- **Total emails** are all inbound emails observed by Darktrace/Email.

- **Actioned emails** are inbound emails with some form of Darktrace/Email response.

- **Users targeted** is the total count of internal email addresses targeted by the malicious communications Darktrace/Email identified.

Outbound Emails

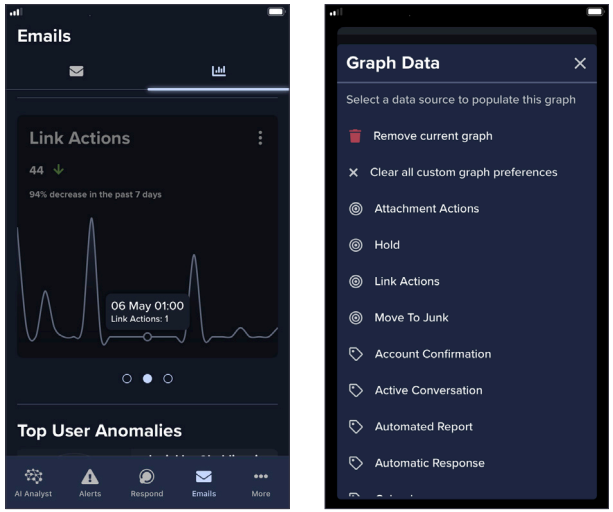
"Detected emails" are those that Darktrace/Email has detected to have some anomalous or otherwise concerning properties. The percentage of detected emails contextualizes these against all observed outbound emails.

- **Total emails** are all outbound emails observed by Darktrace/Email.

- **Detected emails** are outbound emails with unusual properties.

- **Users responsible** is the number of users to whom the anomalous communications Darktrace/Email detected can be attributed.

Overview Graphs



By default, the overview graphs display trends in inbound and outbound emails over time - this can useful to identify spikes that may represent a targeted campaign against the organization.

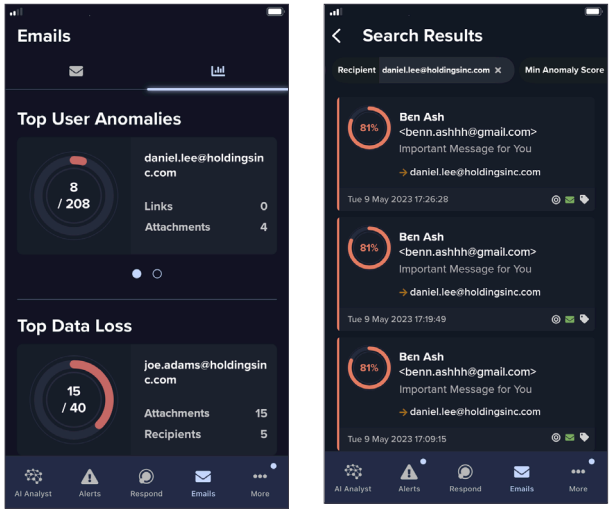
The overview graphs can also be customized to graph significant Darktrace/Email actions or the quantity of emails tagged with Darktrace/Email tags over the last seven days. Tags attempt to characterize the behavior or sender intent observed by Darktrace/Email in the message.

Interactions

The following interactions are available:

- Swipe to move between graphs; the number of tiles is indicated by dots below the element.
- Tap on a graph at any point to display the number of matching emails at that time.
- Tap the ⋮ three dots to remove the graph, change the data displayed on the current graph tile, or to reset all tiles to default.
- Alternatively, swipe to the final blank tile with a ✚ plus icon and tap the icon to add a new graph tile.

User Intent and Data Loss



These elements display key metrics from the *Intent Dashboard* of the Email Console.

User Anomaly Tile (Inbound)

An anomaly score is produced for all emails observed and indicates how unusual Darktrace/Email considers this email to be in comparison to the normal pattern of life. Users who appear on the user anomalies tile are those who received the most anomalous mail over the selected timeframe. The threshold for an email message to be considered for the intent dashboard graph is a Darktrace/Email anomaly score over 60%.

For each user, the number of anomalous emails are displayed against the total number of emails received. The total number of links and attachments associated with the anomalous mailflow is also displayed.

Data Loss Tile (Outbound)

To qualify as a potential data loss incident, an outbound email must breach one of Darktrace's carefully curated Data Loss models. These models detect when a user communicates with or sends files to their personal home email, to a new freemail, address or to entirely new correspondents which may have been created for the purpose of exfiltration.

For each user, the number of emails which may involve some measure of data loss are displayed against the total number of emails sent. The total number of recipients and attachments associated with these emails is also displayed.

Interactions

The following interactions are available:

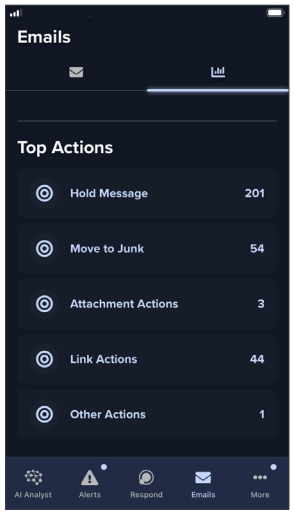
- Swipe to move between senders or recipients on the user anomaly or data loss tiles.
- Tap on an entry to open the email logs filtered for the anomalous behavior or data loss activity for the given user.

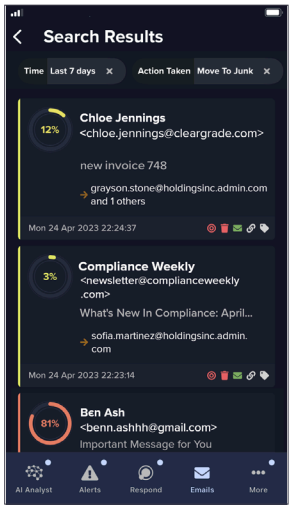
  The email logs are covered in more detail in Email Logs View.

Top Actions

Here, the total actions taken across all email messages is separated into the type of action taken; each category of action can be tapped to filter the email logs by messages actioned in the specified way.

- **Hold Message** and **Move to Junk** are delivery actions which change the final location of the email, typically taken when an email presents a high level of anomaly.
- **Link Actions** are those which alter visible and invisible links that appear in an email messages or store information about links that are seen. Links may be redirected, blocked or removed.





- **Attachment Actions** flatten or remove files attached to emails which exhibit a high level of anomaly.
- **Other Actions** comprise other, *non-significant* actions such as those which alter the visual appearance of the email, like as adding banners or removing spoofed sender information, and those which may send a notification to the user regarding held content.

Email Logs View

Interactive elements described above can open the email logs view with a set of predefined filters applied, such as tags, anomaly level or sender/recipient.

Filtering the email logs and interacting directly with emails is covered in Email Logs View.

DARKTRACE/EMAIL EMAIL LOGS VIEW

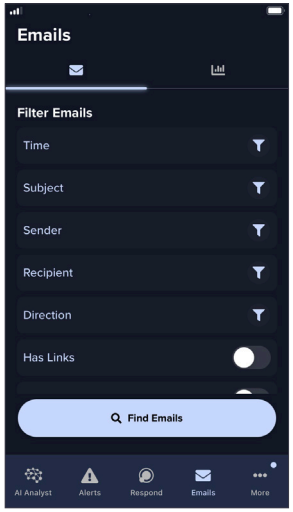
The Emails page is the landing page for interacting with Darktrace/Email functionality in the Darktrace mobile app. There are two main views - Emails and Overview.

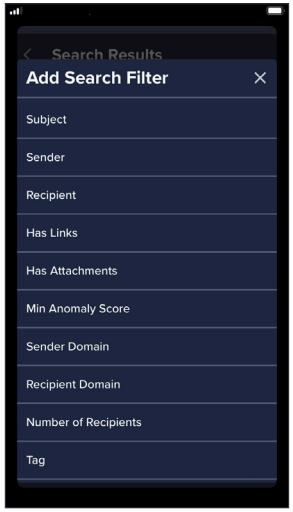
The Overview page was covered previously in Email Overview View. The Emails tab of the page provides full filtered logs - hold emails, release emails, and review the email narrative to quickly understand and action individual emails.

The landing view of the email logs page will differ depending on how it was accessed:

- If accessed from the Emails icon, a filter view will open.
- If accessed from an overview tile, it will open on already filtered logs (section "Email Logs" below).

Email Log Filters





By default, the Darktrace/Email logs will open on a filter page. The functionality is similar to that of the main Email Console filters but the available options are streamlined for on-the-go workflows.

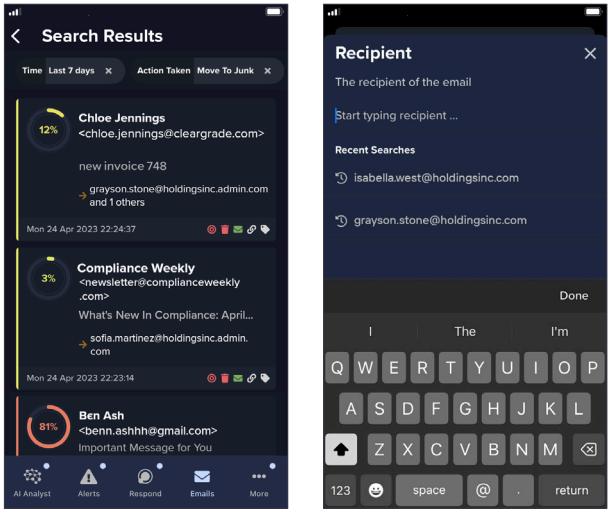
Each filter displays the current value; applied filters are highlighted with an outline. Tap filter elements to add or modify the filter applied. Some filter elements are free text and will maintain a limited history of previous search terms.

Interactions

The following interactions are available:

- Tap a filter entry to add or modify the current filter.
- Tap the x icon to remove the filter.
- Tap "Show More" to display further filters.
- Tap the **Find Emails** button to begin searching the email logs.

Email Logs



After a filter is applied, the email logs will open to display matching email messages. This view will also open if the email logs are accessed from an overview element such as the ***Data Loss*** tile or ***Top Actions*** section.

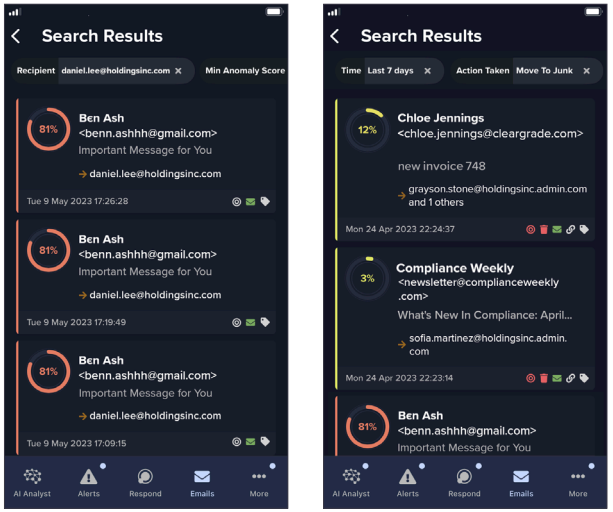
The currently applied filters are displayed across the top - swipe to view all filters. Matching emails are displayed as tiles below.

Interactions

The following interactions are available:

- Swipe right or left on the filters element to review all filters
- Tap ⊕ beside the filters to add an additional filter.
- Tap the ✖ x icon to remove a filter.

Email Log Entries



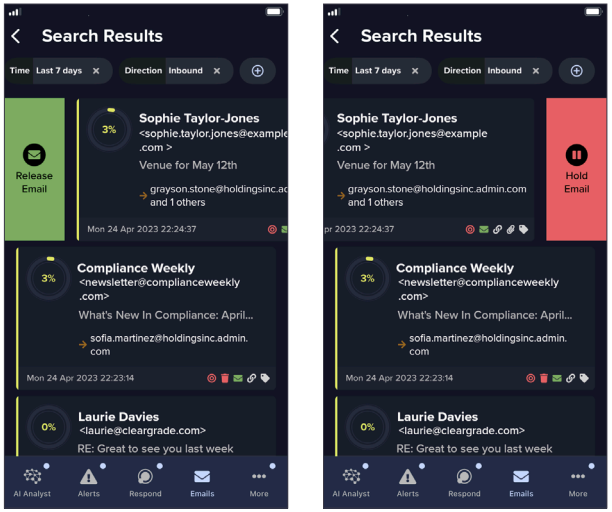
At the top of each email log entry tile is the email sender. The subject is displayed below. The Darktrace/Email anomaly score is displayed on the left of the tile; the color of the vertical bar on the tile is tied to this anomaly score - a yellow line and icon indicating a low score and red a high score.

The email direction is indicated by the colored arrow beside the recipient - a blue arrow indicates outbound, a tan brown arrow indicates inbound. Where there are multiple recipients, one recipient is listed alongside the number of additional recipients.

The time the email was sent (outbound) or received (inbound) is displayed in the bottom left. On the bottom right of the tile are "Quick View Icons" - icons that describe the composition of the message and its final delivery status. These icons are described in depth in the Darktrace/Email guide - Quick View Icons.

Tap on an email entry to open the Email Details.

Interactions



The following interactions are available:

- Tap on an email entry to review more details on the **email details** page.
- Swipe left on an inbound email to **hold** it. If an email has multiple recipients, a dialog will open to select which recipient(s) to hold the email for.

  Held emails are removed from the user inbox until released by an operator.
- Swipe right on an inbound email to **release** it.

  Allows the operator to release an original, unactioned copy of the email to the chosen recipients. If an email has multiple recipients, a dialog will open to select which recipient(s) to release the email to.

Email Details

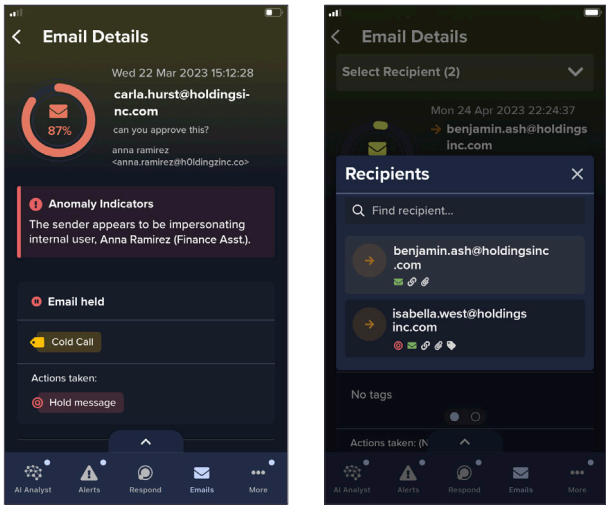
If an email entry is tapped from the email logs view, the **Email Details** page will open. This contains detailed information about the email and further interactions.

Email Details is covered in Email Details View.

DARKTRACE/EMAIL EMAIL DETAILS VIEW

The Email Details view is opened from an email entry in the Email Logs View. It displays key information about the email and allows actions such as hold or release to be taken and for learning exceptions to be taken.

Email Details



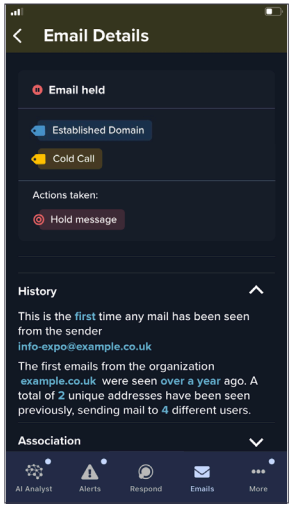
The email details opens when a specific email is selected from the logs.

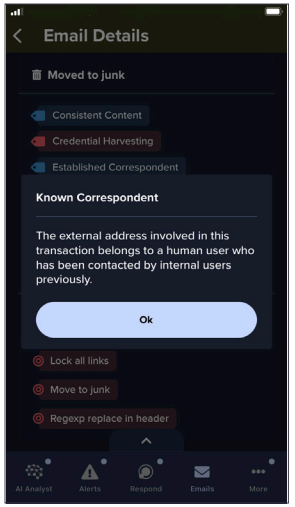
The top of the overview displays the time the email was received (inbound) or sent (outbound). In ths view, the recipient is always displayed first. The arrow beside the recipient indicates direction - a blue arrow indicates outbound direction, a tan brown arrow indicates inbound. Where there are multiple recipients, an additional dialog will appear to swap between recipients.

The Darktrace/Email anomaly score is displayed on the left and the email subject is displayed below. Underneath this element is the recipient (outbound) or sender (inbound).

The ❗ anomaly indicators section expresses the most anomalous, suspicious or potentially threatening aspects of the email. The color of the section indicates the level of anomaly, with red the highest. It will appear for both inbound and outbound emails if there are highly anomalous properties detected by Darktrace/Email. Where an email contains a rare link, or the content indicates the sender is attempting some kind of manipulation, these indicators are displayed here.

Status and Tags



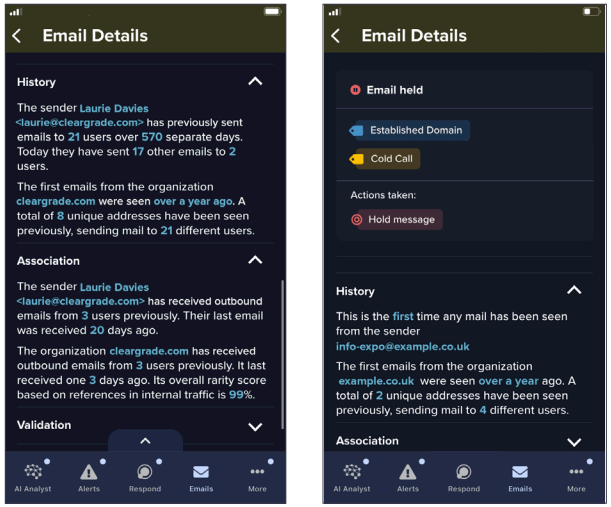


Below the anomaly indicators section is the current delivery status of the message. This status will change if the email is modified, or the final location altered. See the **Status Icon** under Quick View Icons above.

Underneath are a list of tags. Tags are applied to emails by Darktrace/Email to characterize the possible content or intent of the message. Severity is indicated by the icon color: malicious intents are represented in red and those requiring caution are indicated in yellow. Informational tags which give important contextual information (such as indicating a bulk mail or a known correspondent) will appear blue Tap on a tag to view a short description.

The actions taken section describes any significant Darktrace/Email actions taken on the email. A red icon indicates an action was performed and grey icon indicates the action was not taken, whether due to global configuration settings or the user's eligibility for automatic actions. Tap an action for a short description of the action and why it was prevented, if applicable.

Other Email Details Sections



The email details tab breaks down Darktrace/Email's analysis of the message into short, textual summaries that allow a user to understand the context and possible threat posed at a glance. Sections will only appear when relevant and are collapsed automatically. Tap to expand.

History

The history section summarizes the sending history of the external party in the correspondence. For inbound emails the external party is the sender of the email, for outbound emails it is the recipient. This section is useful to understand how the external user and the domain they are associated with have interacted with your organization.

Association

The association section summarizes the relationships that have been created with the external party by internal users. For inbound emails the external party is the sender of the mail, for outbound emails it is the recipient. Association is a reversal of the history section - it covers how your organization has interacted with the external user and their domain.

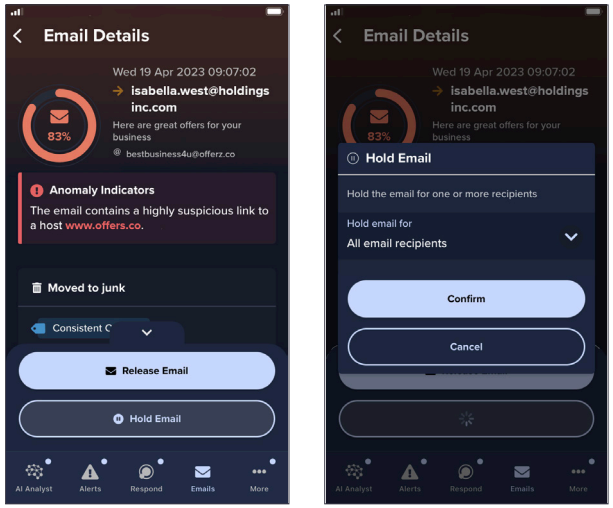
Validation

Email messages must pass validation in some way to be delivered. If they do not pass validation, or are unable to be verified successfully, this may indicate a concern. The validation section describes the results of these conventional authentication checks on the email. For inbound emails only.

Content

The final section, content, describes any further characteristics of the email which may be interesting. This may include information about the email type (e.g., bulk mailing), any low-level anomaly in the message content, or the composition of links and attachments.

Interactions



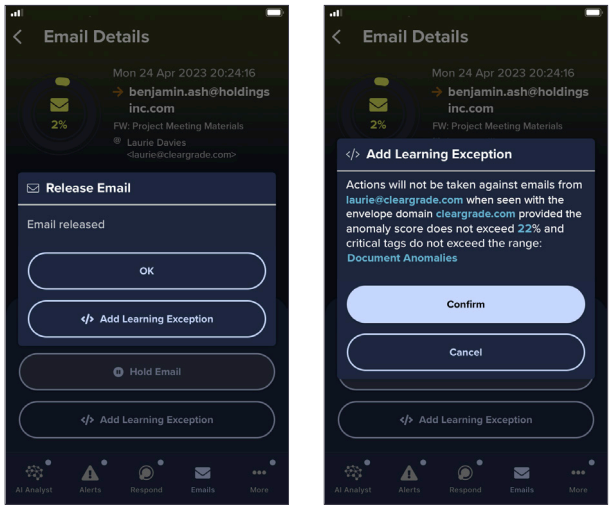
The following interactions are available:

- If there are multiple recipients, swipe left or right to move between the recipients or use the recipient selector at the top of the page.

- Tap the ^ up arrow to review more actions:
  - Tap the **Release Email** option to release the email to one or more recipients.

    Release an original, unactioned copy of the email to the chosen recipients. If an email has multiple recipients, a dialog will open to select which recipient(s) to release the email to.
  - Tap the **Hold Email** option to hold the email for one or more recipients.

    Held emails are removed from the user inbox until released by an operator.
  - If the email is released, the additional **Add Learning Exception** option will appear. Tap this option to create a learning exception to influence Darktrace/Email response to emails from this sender in the future.

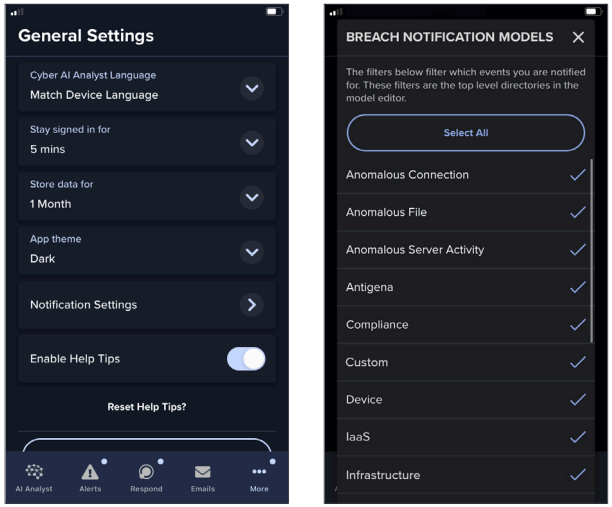


This option will open a confirmation dialog describing the learning exception and any changes to *lists* made as a result. Please see the full Email Console guide "Learning Exceptions" for more information on learning exceptions.

APP SETTINGS

The **Settings** screen offers multiple filtering options that allow you to customize how data is displayed and stored within the app. In 4.0, the number of settings has been reduced in favour of native filtering within the AI Analyst, Alerts and RESPOND pages.

Settings



Cyber AI Analyst Language

Changes the default language for the AI Analyst view. Tap the language to move through all available languages or long press to select from a list.

Stay Signed in For

Set the time required between logins to 5 minutes (default), 15 minutes or 1 hour.

Store Data

Controls how long data is stored within the application on the mobile device.

App Theme

Control whether the app should follow system theme or always operate in light or dark mode.

Enable Help Tips

Enables the workflow descriptions that appear when the app is first opened.

Reset Tips

Resets all workflow descriptions to display as if the app is newly installed.

Re-authenticate App

Re-scan the QR code in the main Threat Visualizer Account Admin page.

Change Pin Code

Alter the (minimum four digit) code created during authentication for an alternative log-in method.

Log Out

Ends the current session in the app and returns to the locked screen, forcing the user to login before continuing their investigation.

Please note, this does not end authentication with the Darktrace instance. The user can continue operating the app once a successful login is performed.

Notification Settings

Incident Notification Categories

Controls the threat behavior categories that can create push app notifications - visible in the Notification Centre (iOS) and on the lock screen. AI Analyst incidents which are not part of these categories will not create notifications.

Incident Score Notification Threshold

Controls the minimum score that an AI Analyst incident must reach before it creates an app notification. Applied in conjunction with the category setting above.

Breach Notification Categories

Controls the threat behavior categories that can create push app notifications for model breaches. Model breaches which are not part of these categories will not create notifications.

Breach Score Notification Threshold

Controls the minimum threat score that a model breach must reach before it creates an app notification. Applied in conjunction with the model breach category setting above.

Breach Notification Models

Controls the model categories that can create push app notifications for model breaches; model categories are derived from the folders the models are saved in. Model breaches which are not part of these folders will not create notifications.

ABOUT DARKTRACE

Darktrace (DARK.L), a global leader in cyber security AI, delivers complete AI-powered solutions in its mission to free the world of cyber disruption. We protect more than 7,400 customers from the world's most complex threats, including ransomware, cloud, and SaaS attacks. Darktrace is delivering the first-ever Cyber AI Loop, fueling a continuous security capability that can autonomously spot and respond to novel in-progress threats within seconds. Darktrace has 115+ patent applications filed. Darktrace was named one of TIME magazine's "Most Influential Companies" in 2021.

LAST UPDATED

May 15 2023

US: +1 415 229 9100 UK: +44 (0) 1223 394 100 LATAM: +55 11 4949 7696 APAC: +65 6804 5010 info@darktrace.com darktrace.com