DARKTRACE

Analysts' Workflow Guide

Darktrace / EMAIL

2024-10-21 Version 1.0 | darktrace.com

Read me first

Disclaimer

This document is intended to complement the foundational training and product guides available on the Customer Portal. It is strongly recommended that users complete this training before consulting this guide.

Failure to do so may result in changes to your deployment that could expose your organization to risks or vulnerabilities. The information provided in this document is based on the assumption that users possess the requisite knowledge from the training. By using this guide without the appropriate training, you acknowledge that you may not fully understand the implications of your actions, which could potentially lead to security incidents.

The information provided in this document is for general guidance and informational purposes only. It does not constitute professional or legal advice.

Darktrace makes no representations or warranties regarding the accuracy, completeness, or reliability of the information contained in this document. Any reliance you place on such information is at your own risk. Darktrace disclaims any liability for any loss or damage arising from the use of the information provided in this document. It is advisable to consult with legal counsel regarding any applicable regulations or legal matters.

These guides may be updated from time-to-time and the latest version, available from the Customer Portal, should always be used.

Contents

Introduction

Aim

This Playbook is intended to complement the training resources available to all users via the Customer Portal. It sets out an example best practice workflow for Darktrace/EMAIL that should be adapted to requirements.

Note that unlike /NETWORK, it is not expected for security teams to triage actions and emails on a regular basis. This is because by default, the tool provides protection autonomously and is frequently updated with new models and metrics. However, workflows using the tool can still be established to extract more value from the tool or fulfill specific use cases. The purpose of this document is to explain how to use Darktrace/EMAIL for teams that want to interact with the tool routinely.

The document is structured in two parts: exploration and optimization. The exploration section will begin with the quickest, highest level overviews of the tool through to slower, more granular triage. The optimization section will similarly begin with general tool behavior settings and will gradually narrow down to individual features with specific use cases. This is based on the assumption that all users will wish to immediately triage apparent, high impact issues.

The workflows described here are entirely flexible and can be followed by one user as a process, divided up and approached as separate tasks on different days, or allocated as separate responsibilities to different users/teams.

Assumptions

User skills

We assume the user has watched the videos in the Darktrace / EMAIL Part 1 - Familiarization Course and the Darktrace / EMAIL Part 2 - Customization Course.

Adaptation

It is expected that users will adapt the processes to support their own requirements and priorities. As a generic process it cannot be exactly suitable for every use case, and users retain ultimate responsibility for their security.

Using this document is meant to serve as a process guide.

Each section includes a heading, indicative timings and frequency. Steps are structured as an action (in bold) followed by further explanation and detail: This allows more confident users to concentrate on the actions and filter out the background information.

Workflow diagram

For confident users, the process diagram (Figure 1) illustrates the workflow steps only and is intended as a handy "desktop" guide. Note that each "layer" contains tasks from both exploration (white boxes) and optimization (black boxes).

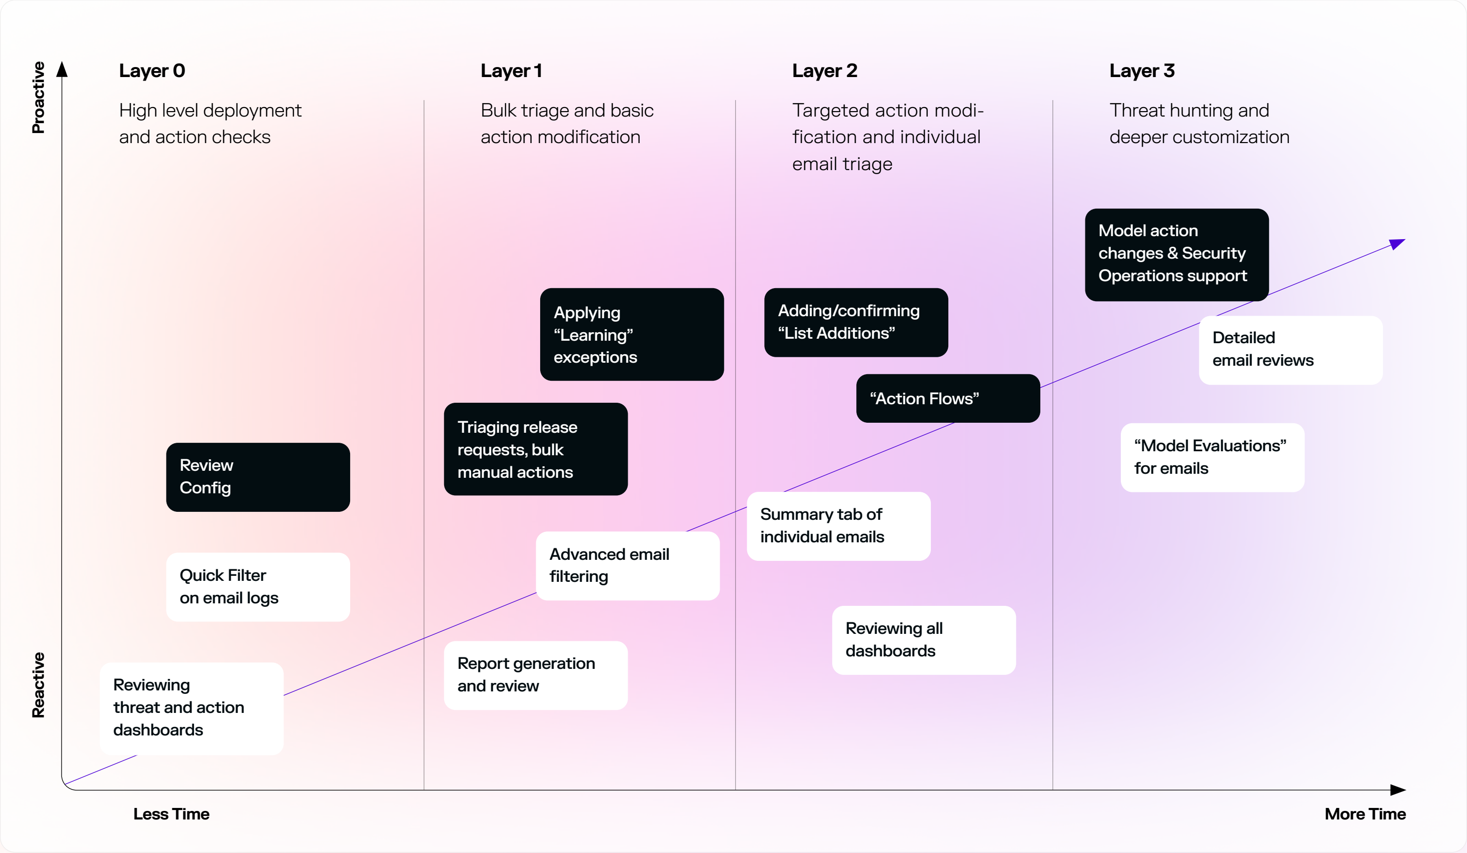
Proactive
Reactive
Layer 0
High level deployment and action checks
Layer 1
Bulk triage and basic action modification
Layer 2
Targeted action modification and individual email triage
Layer 3
Threat hunting and deeper customization
Review Config
Quick Filter on email logs
Reviewing threat and action dashboards
Applying "Learning" exceptions
Triaging release requests, bulk manual actions
Advanced email filtering
Report generation and review
Adding/confirming "List Additions"
"Action Flows"
Summary tab of individual emails
Reviewing all dashboards
Model action changes & Security Operations support
Detailed email reviews
"Model Evaluations" for emails
Less Time
More Time

Figure 01: Workflow Diagram

□ Exploration Workflow ■ Optimization Workflow

Part 1

Deployment Exploration Workflows

Part 1 - Deployment Exploration Workflows focuses on identifying and investigating the emails present on the deployment. Understanding all layers in this section will help in applying the information laid out in **Part 2.**

Workflow Tip

First, review high priority dashboards or configure a custom dashboard with important insights for your team, before diving into an investigation of emails within the logs. Doing so allows security teams to quickly identify key insights and efficiently prioritize tasks.

Afterwards, review the User Interactions dashboard tab in case users have reported malicious emails or have requested an email to be released.

Part 1, Layer 0:

Deployment level triage

Basic Dashboard Review

Time allocation: ~ 5 minutes
Dashboards provide high level overviews of key metrics across different facets of Darktrace / EMAIL detection. Investigation of the dashboards can provide a jumping off point for deeper investigations described later in this document or identify potential optimization opportunities. Dashboards can be configured to look at data in a daily, weekly, or monthly time interval.

Below are the three default dashboards users should start with:

- **Threats:** The Threats tab that appears in the email dashboard is meant to visualize certain characteristics and trends that constitute malicious activity which is detected by Darktrace / EMAIL
- **Actions:** The Actions tab that appears in the email dashboard is meant to serve as an overview of the autonomous actions that Darktrace / EMAIL has taken
- **User Interactions:** The User Interactions tab in Darktrace / EMAIL visualizes interactions with end users and Darktrace / EMAIL or other security tools. When navigating this tab, a user can review user actions such as reporting malicious emails, marking a sender as not junk or requesting email releases. Additionally, users can quickly access pending release requests.

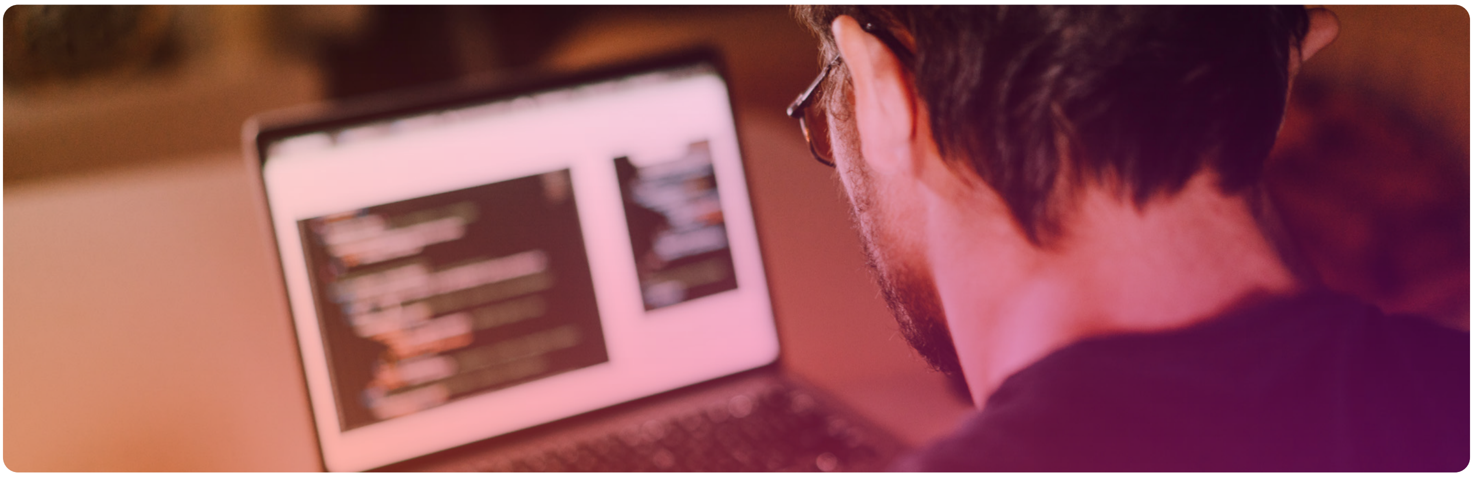

Quick Filters

Time allocation: ~ 5 minutes

Accessing the logs and interacting with email entries forms a core part of the Darktrace / EMAIL workflow. Manipulating email filters is the best way to pivot to specific emails and email categories during investigations.

Utilize Quick Filters to easily sift through email logs. These are available at the top of the Logs page, below the search bar. There are multiple dropdown menus with options that can be selected to narrow down an investigation with basic filtering.

- **Actions:** Filter emails by type of actions taken on them.
- **Tags:** Filter emails by certain tags that were assigned by the Darktrace AI.
- **Flags:** Filter by flags, which are custom identifiers usually manually assigned to emails.
- **End Users:** Filter by emails with specific end user interactions.

Workflow Tip

Filter by Significant Action to identify which emails should be triaged first, since highly suspicious emails tend to receive higher-severity actions.

Please note that, depending on Global Action Severity settings and how Darktrace / EMAIL has been configured for users and groups, this may not always be the case. Regardless, a quick analysis of emails with suggested or enacted significant actions – such as Hold, Junk, Double Lock Link or Strip Attachment – can be beneficial to understand current email security threats seen in the environment.

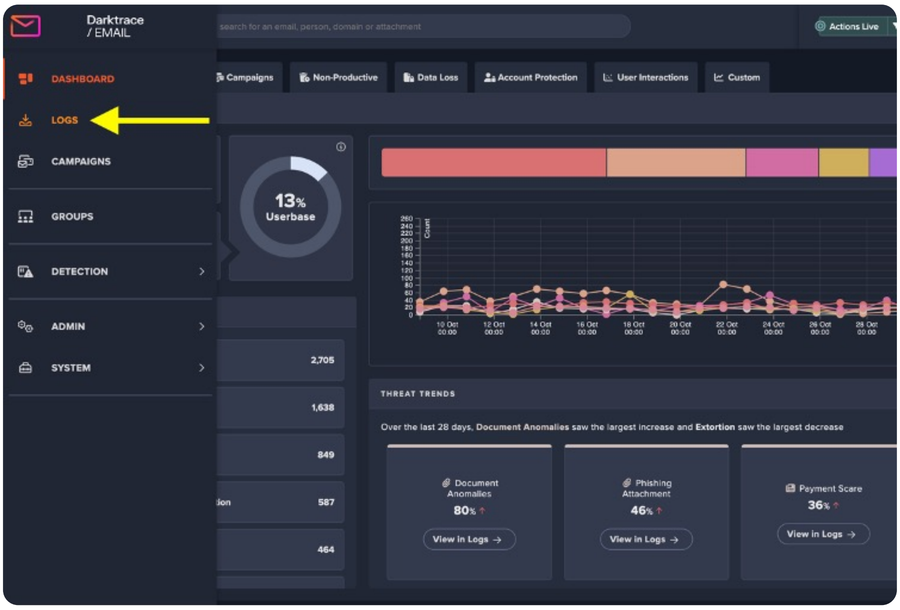

Over the last 28 days, **Document Anomalies** saw the largest increase and **Extortion** saw the largest decrease

Figure 02: Navigating to the Logs page from the main UI.

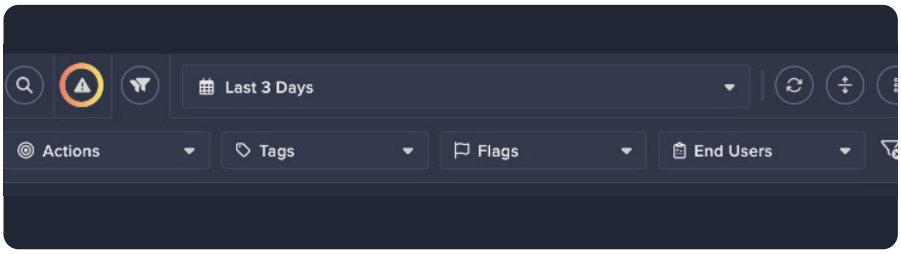

Figure 03: Quick Filters seen in the Logs page.

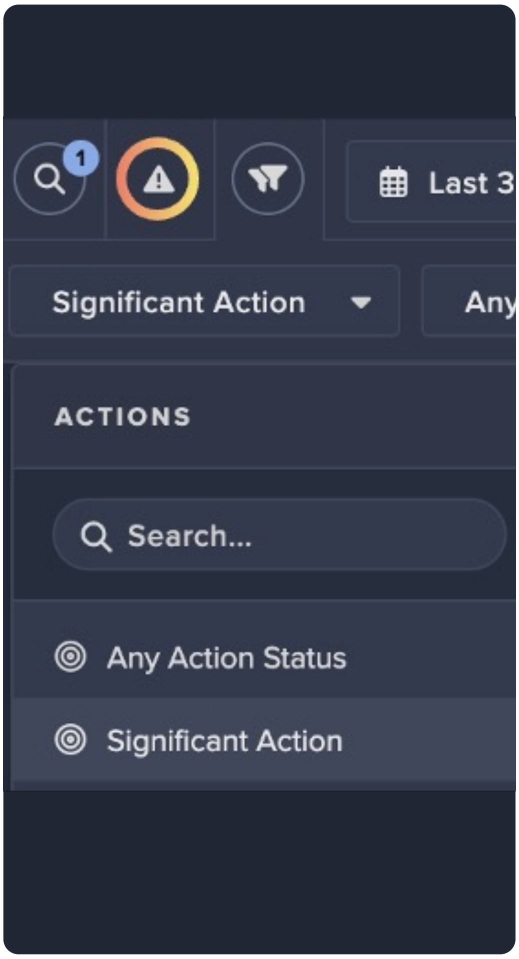

Figure 04: Significant Action Quick Filter option.

Users should now move to Part 2 Layer 0.

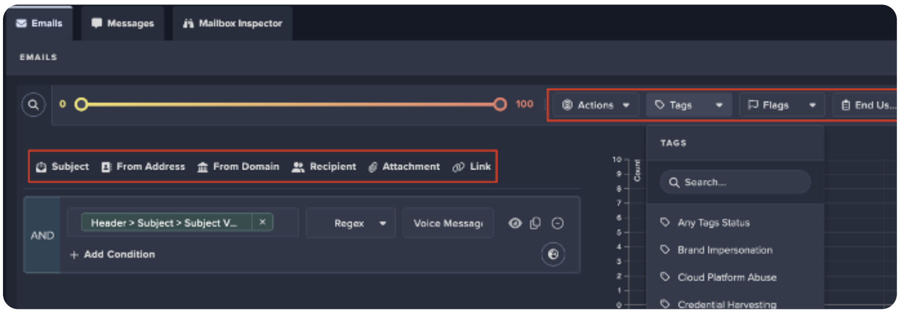

Figure 05: Simple Advanced Filters and Quick Filters in the Logs page.

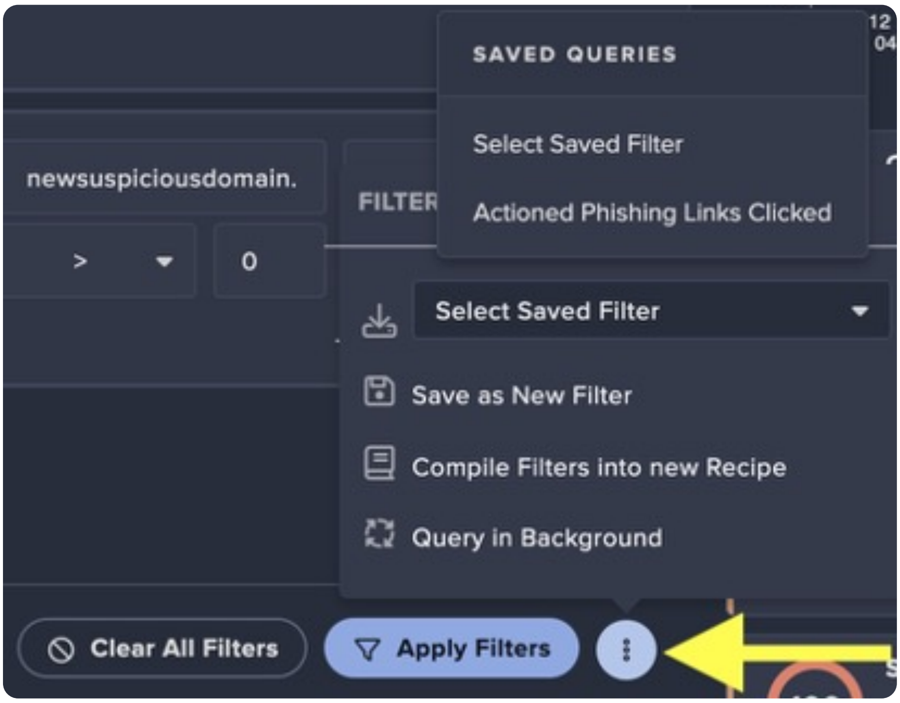

Figure 06: Darktrace / EMAIL save filter options.

Part 1, Layer 1:

Bulk email triage

Advanced Filters

Time allocation: ~ 10 minutes

Conduct a deeper analysis with Advanced Filters. If needed, more advanced filtering can be accomplished by clicking on the eyeglass icon seen in Figure 3. Another panel will slide open with more options.

The red rectangle on the left-hand side of Figure 4 highlights the Simple Advanced Filters. These are shortcuts to add common filter queries – such as Subject, From Address and From Domain – to your current search condition.

Below them, more complex conditional statements can be created to filter emails. On the right-hand side of Figure 5, also encapsulated in a red rectangle, are the Quick Filters mentioned beforehand.

If a filter condition from the Quick Filters or the Simple Advanced Filters has been added, its logic will also be visualized in the advanced filters panel as a conditional “AND” or “OR” statement.

More information on constructing advanced filter queries can be found in the Customer Portal.

Workflow Tip

Save commonly used filter conditions to quickly re-use them.

Doing this can save analysts time from having to re-construct complex queries. This can be accomplished by clicking on the icon with the vertical ellipsis at the bottom of the Advanced Filters panel. From there, click on "Save as New Filter".

As you can see, clicking on that same icon also allows analysts to select a previously saved filter to re-load it. Please note that filters are saved per user account and are not visible for all users.

Reading Threat Reports

Time allocation: ~ 15 minutes
Use the reporting feature to generate a document containing high level overview of Darktrace / EMAIL detection as well as highlights of specific emails and users. These can be generated on an ad-hoc basis, or scheduled for specific times.

- **Period:** The designated time range that the user wishes to generate the report for
- **Timezone:** Timezone used for the report. Users can input a region and select the best match
- **Include Threat Examples:** Enable to showcase specific emails in the selected time period within the report
- **Use Recipient Count:** The numbers displayed in the report will be based on the number of recipients of processed emails rather than total number of emails processed
- **Restriction Type:** Limit the type of emails the user wishes to see in the threat report by Domain or Groups
- **Custom Filters:** Include Dashboard enabled filters in the report. Some examples include Brand Impersonation, Credential Harvesting, MultiStage Payload and Email Account Takeover

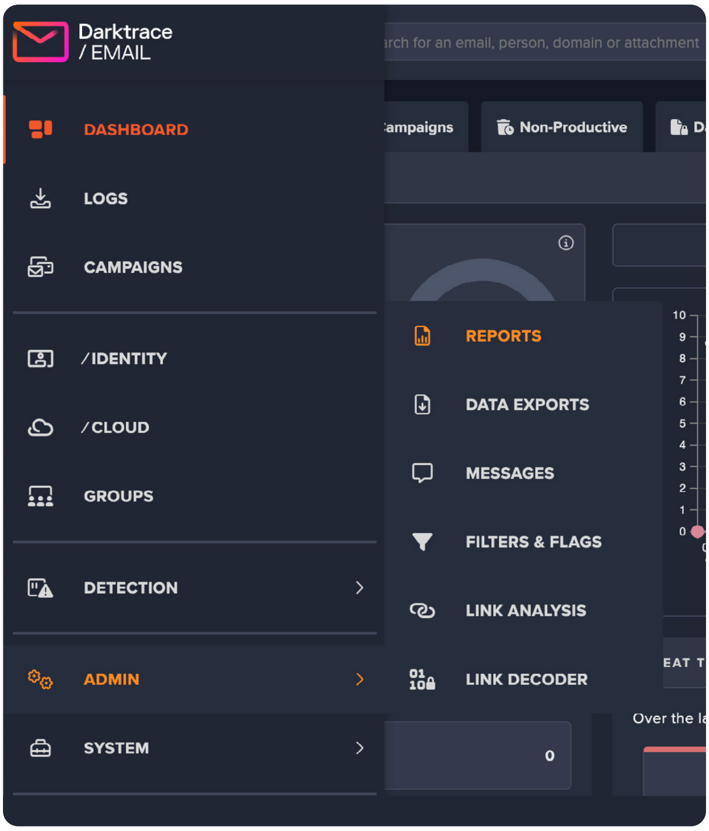

Figure 07: Location of Reports tab and "Generate Threat Report" button

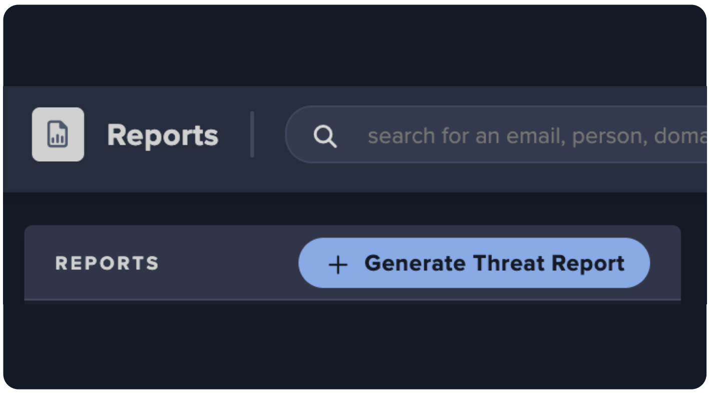

Figure 08: Location of Reports tab and "Generate Threat Report" button

Users should now move to Part 2 Layer 1.

Part 1, Layer 2:

Targeted email triage

Reviewing All Dashboards

Time allocation: ~ 15 minutes

Following review of the three basic dashboards in Part 1, Layer 0, users may wish to review the remaining ones for insight into more specific threats.

- **Campaigns:** The campaigns tab will provide a high level overview of campaigns that are identified and the primary type of malicious communications that being used within each campaign
- **Non-productive:** The non-productive tab in Darktrace / EMAIL is meant to highlight the different types of unsolicited emails that are identified by Darktrace / EMAIL and showcase the increase in productivity by actioning said emails. By default, emails categorized as "Graymail" or "Cold Calls" are not actioned. However, there are default action flows that can be turned on to enable actions towards these unwanted or unproductive emails. More information can be found in the "Action Flows" subsection of this guide.
- **Data Loss:** The data loss tab in Darktrace / EMAIL is meant to showcase potential data loss incidents. This tab will highlight outbound email communications which may contain sensitive business data that was sent out to external recipients whether intentional or accidental.
- **Darktrace/IDENTITY:** Darktrace / IDENTITY Account Protection tab centralizes Darktrace / EMAIL account protection alerts and enables users to review suspicious email communications against what is considered anomalous user behavior and historic login locations.

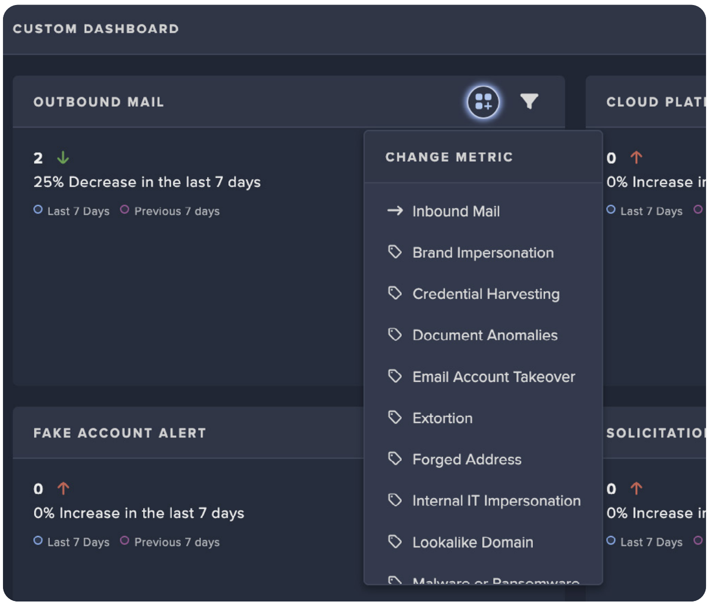

Figure 09: Custom dashboard options

- **Custom:** The Custom Tab in Darktrace / EMAIL allows the user to customize the activity that they wish to observe in the Darktrace / EMAIL user interface. Users can customize the dashboard by time of activity and filter on the metric they wish to see on the dashboard by clicking the "Change Metric" button as showcased in the screenshot. Examples of metrics include Forged Address, Extortion, and Tag, among others.

Email overview and AI Analyst tab

Time allocation: ~ 5 minutes/email

After filtering for emails using the methods in either Layer 1 or Layer 2, users can hone in on specific emails and obtain key information about them. The Overview Tab will open when clicking on an email log.

The Overview Tab will display the following information:

- A summary of the anomaly indicators provided by Darktrace / EMAIL and Cyber AI Analyst
- Actions taken on the email
- Tags categorizing the email based on models the email triggered
- Highlighted content and/or payloads that contributed to the anomaly indicators

The Cyber AI Analyst Tab provides additional high-level details on the email. This includes Anomaly Indicators, the history of the emails being received by the sending organization, the history of communication between the recipient and the sender, and validation metrics.

Users should now move to Part 2 Layer 2.

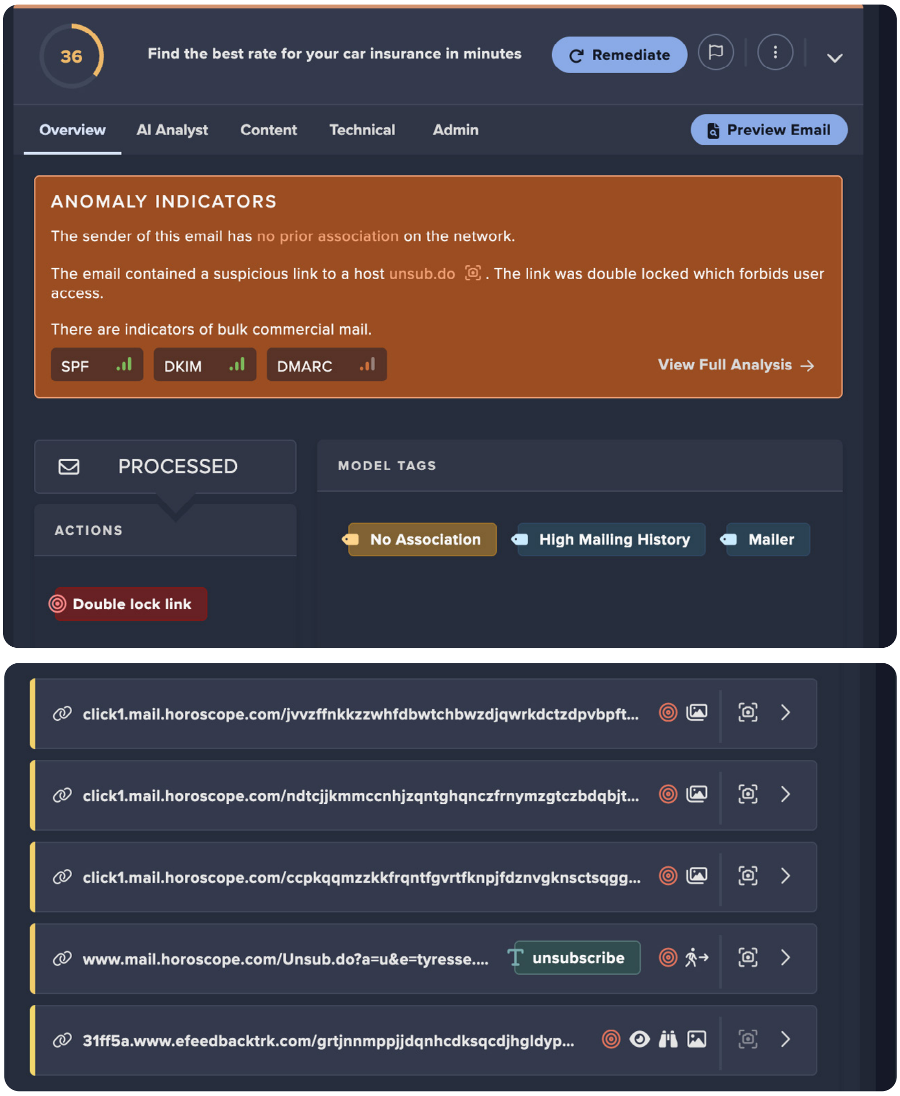

The sender of this email has no prior association on the network.

The email contained a suspicious link to a host unsub.do. The link was double locked which forbids user access.

There are indicators of bulk commercial mail.

click1.mail.horoscope.com/jvvzffnkkzzwhfdbwtchbwzdjqwrkdctzdpvbpft...

click1.mail.horoscope.com/ndtcjjkmmccnhjzqntghqnczfrnymzgtczbdqbjt...

click1.mail.horoscope.com/ccpkqqmzzkkfrqntfgvrtfknpjfdznvgknsctsqgg...

www.mail.horoscope.com/Unsub.do?a=u&e=tyresse....

31ff5a.www.efeedbacktrk.com/grtjnnmppjjdqnhcdksqcdjhgldyp...

Figure 10: Overview tab of an email with highlighted content towards the bottom of the page.

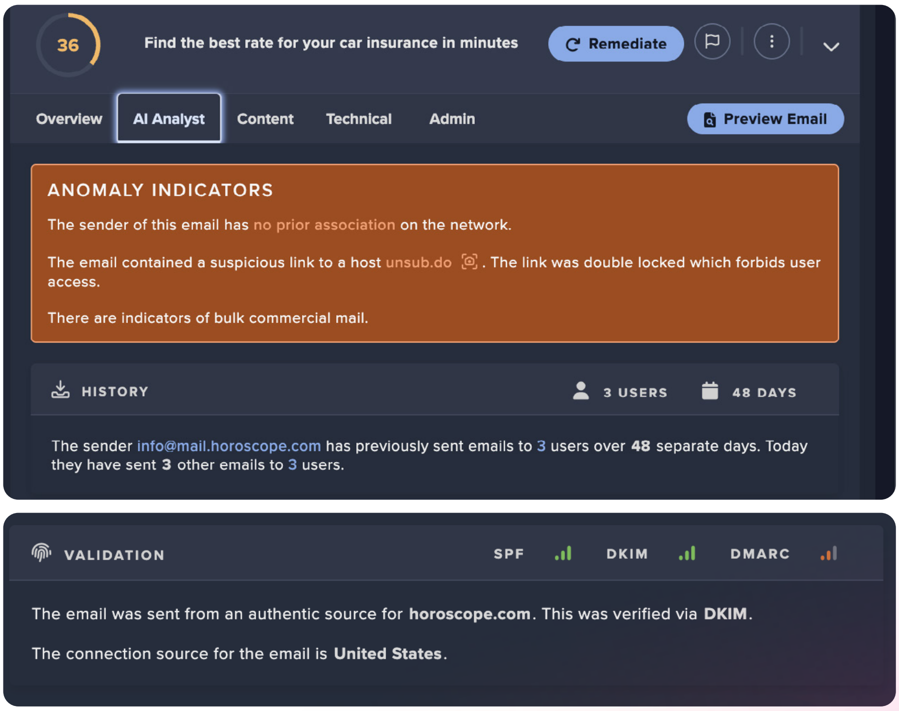

The sender of this email has no prior association on the network.

The email contained a suspicious link to a host unsub.do. The link was double locked which forbids user access.

There are indicators of bulk commercial mail.

The sender info@mail.horoscope.com has previously sent emails to 3 users over **48** separate days. Today they have sent **3** other emails to 3 users.

The email was sent from an authentic source for **horoscope.com**. This was verified via **DKIM**.

The connection source for the email is **United States**.

Figure 11: AI Analyst tab of an email showing the anomaly indicators summary and the three metric summaries.

Part 1, Layer 3:

Advanced email triage

Reviewing email metrics

Time allocation: ~ 20 minutes/email

- **Content Tab:** This tab in the EMAIL log entry will always contain a “Correspondents” sub-tab which contains information on all participants in the email. Additional tabs that can appear here are “Links”, “Attachments”, “Phones Numbers” and “Financials Details” if the corresponding content is contained within the email. Each of these list the individual piece of content in each sub-tab as well as metrics specific to each.
- **Technical Tab:** This Tab in the Darktrace / EMAIL User Interface showcases Metrics, Model Alerts, Email Headers and Authentication Headers
  - **Metrics:** This tab showcases metrics derived from email meta-data. Key metrics are used to gather more insight about an email and its evaluation. Within the Metrics subtab, to the right of the search bar you can click the toggle next to “Advanced Metrics” to reveal all available metrics.
    - **Key Headers:** Contains the header from domain, address, personal and subject values
    - **Action:** Contains information on actions applied as well as campaign metrics if email was added to campaign
    - **Attachment:** Contains two metrics indicating if there was an attachment and if so, how many there were
    - **Connection:** Primarily contains information on the envelope address and the sender infrastructure
    - **Email:** Primarily contains metrics on the body content of the email.
    - **Header:** Primarily contains more detailed information about the header from address and reply-to address
    - **Identity:** Primarily displays metrics that indicate shifts in sender behavior
    - **Link:** Contains link counts, link specific metrics
    - **Model:** Contains list of models triggered and tags applied
    - **Recipient:** Contains similar metrics for the Header field but applied to recipient address instead
    - **Validation:** Contains metrics specific to email validation such as whether the email passed SPF.

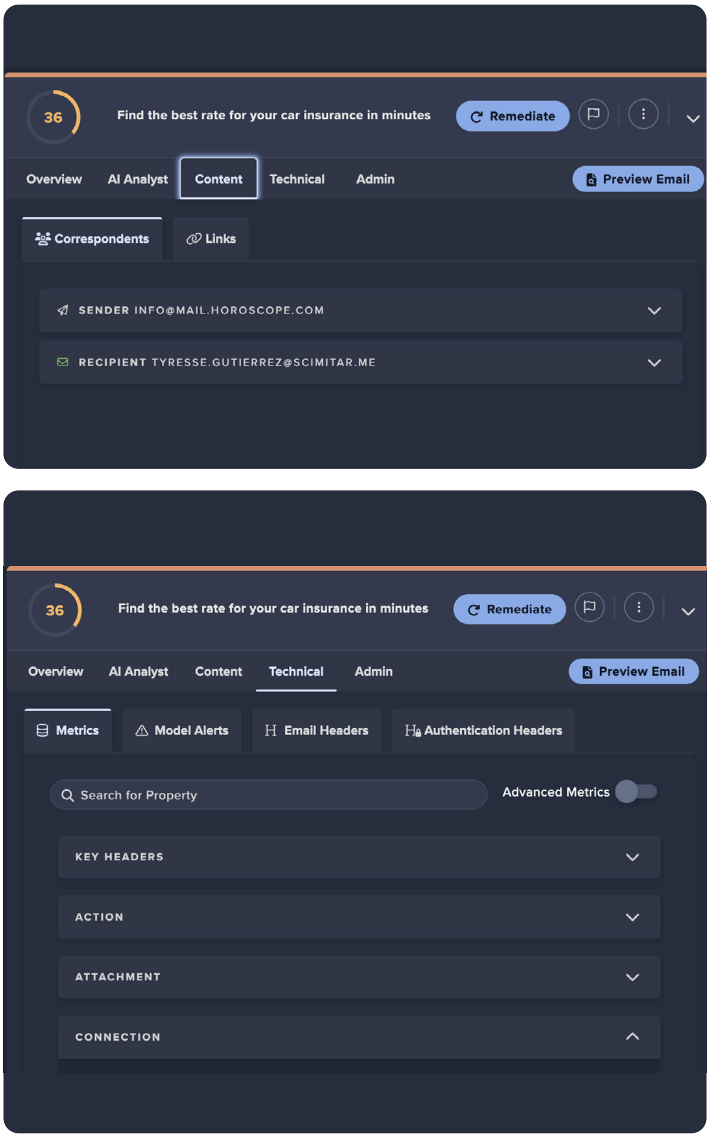

Find the best rate for your car insurance in minutes

SENDER INFO@MAIL.HOROSCOPE.COM

RECIPIENT TYRESSE.GUTIERREZ@SCIMITAR.ME

Figure 12: The ”Technical” tab selected for an email followed by selection of the ”Metrics” subtab.

- **Email Headers:** Contains the raw email header-key values. These can be downloaded using the “Download Headers” button
- **Authentication Headers:** Contains meta-data on the raw authentication header-key values

Reviewing email models and model evaluation

Time allocation: ~ 30 minutes/email
The “Model Alerts” sub-tab under technical will list all of the models that any given email triggered. Models are what increase/decrease the anomaly score of an email and advise the actions. The anomaly score is a key metric used in the models “Antigena Anomaly” and “High Antigena Anomaly” which in practice, produce most of any given deployment’s delivery actions.

By default, the former will junk emails above 60 anomaly and the latter will hold emails above 95 anomaly. Any applicable action flows will also be listed at the top (see Part 2, Layer 2 for more information on action flows.)

Higher anomaly models will be placed towards the top of the list. Conversely, models that remove anomaly will be placed towards the bottom. There are also many informational models that provide additional context, but don’t change the anomaly of an email.

Any model seen in this tab can be evaluated. This provides users with a way to see why a particular model triggered based on existing metrics and recipes. To do so, click on the small triangle next to the model name then click on the button next to “View Evaluation Log”. This will then show a visual representation of the model logic, where the trigger requirements for the model can be inspected at each stage.

Users should now move to Part 2 Layer 3.

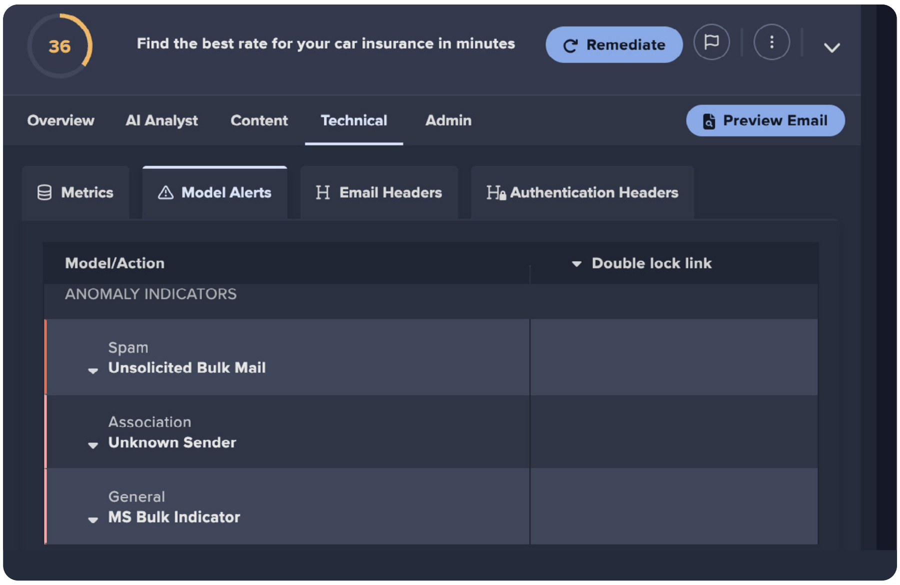

Figure 14: The ”Technical” tab selected for an email followed by selection of the ”Model Alerts” subtab.

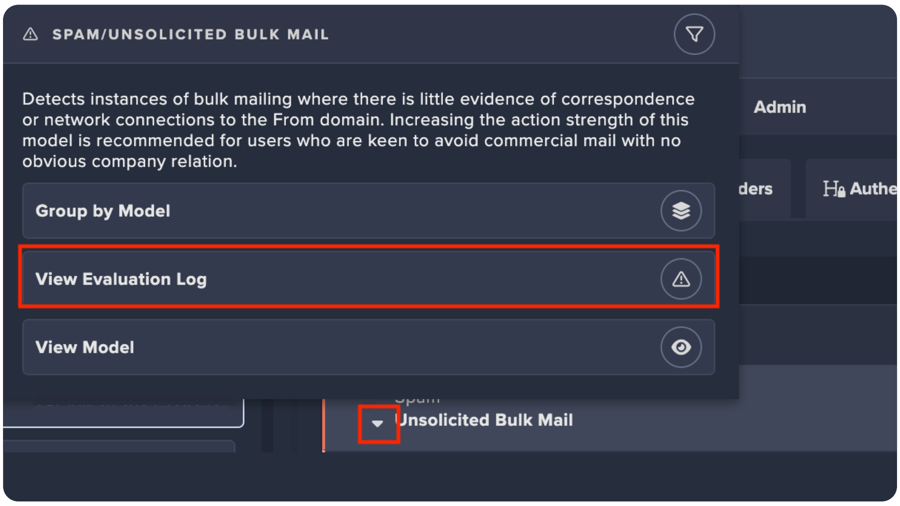
SPAM/UNSOLICITED BULK MAIL

Detects instances of bulk mailing where there is little evidence of correspondence or network connections to the From domain. Increasing the action strength of this model is recommended for users who are keen to avoid commercial mail with no obvious company relation.

Figure 13: Location of the ”View Evaluation Log” button for evaluating specific models.

Part 2

Deployment Optimization Workflows

Part 2 – Deployment Optimization Workflows focuses on settings and processes that tailor the behavior of the tool to user requirements. Most changes can be configured holistically for all users and groups, or individually on a per-user or per-group basis.

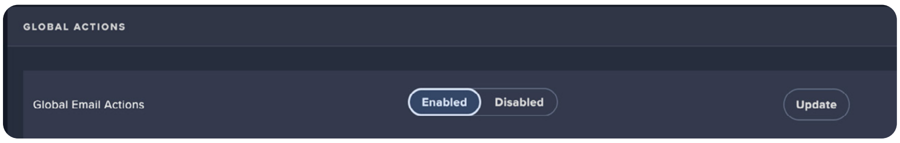

GLOBAL ACTIONS

Global Email Actions [✓ Enabled] Disabled

Figure 15: Global email actions setting set to "Enabled"

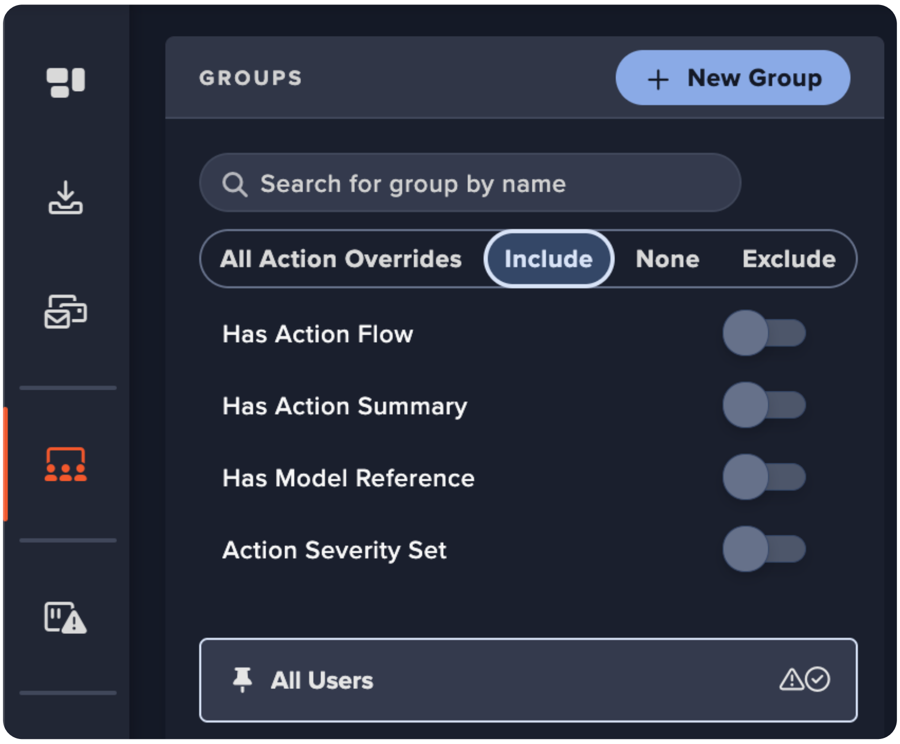

GROUPS

All Action Overrides [✓ Include] None Exclude

Has Action Flow ○

Has Action Summary ○

Has Model Reference ○

Action Severity Set ○

All Users

Figure 16: Example of the All Users group in an "Include" action override

Part 2, Layer 0:

Reviewing Configurations

Time allocation: ~ 6 minutes

Check Global Email Actions

Time allocation: ~ 30 seconds

The "Global Email Actions" controls whether Darktrace / EMAIL is able to take any actions whatsoever. When disabled, all actions are advisory. When enabled, actions can be taken only on users that are eligible due to group membership [16]

Check All User Inclusion

Time allocation: ~ 1 minute

For all users to receive actions if Global Email Actions is enabled, the "All Users" group needs to be set to the "Include" action override.

This can be checked by navigating to "Groups", then the "Include" tab. Note that by default the "All Users" group does not list members, but it does contain all users. Any groups in "None" that are not in "Include", and groups that are in "Exclude" will not receive actioning. "Exclude" group membership takes priority over the other two.

Check Action Severity Settings

Time allocation: ~ 5 minutes

Maximum Action Severity Settings can be configured to turn off significant action types while still otherwise allowing Darktrace / EMAIL to function as normal. These settings can accessed via Menu > System > Config page and under the Global Actions sub-window. When changing action severity settings, be sure to click the "Update" button next to each option to commit the change.

When increasing the severity, all actions in lower tiers are also still permitted. Darktrace / EMAIL can take four different types of actions: Delivery, Link, Attachment and Header actions.

For additional details on what each level allows for each action, see the following documentation. Lastly, note that action severity settings can be configured per group and per user, but this falls outside the scope of this section.

Top Tip

Consider allowing Darktrace / EMAIL to take the **highest severity actions** possible. Security teams should aim to globally permit all types of actions in their environment to ensure proper risk mitigation of possible email compromises and properly allow Darktrace / EMAIL to act on anomalous correspondence. Special evaluation for a smaller set of specified users or groups can be configured accordingly, if needed.

For assistance configuring Darktrace / EMAIL actions, a ticket can be opened with Darktrace's Security Operations Support service.

Users should now move to Part 1 Layer 0.

Part 2, Layer 1:

Manual actioning

Time allocation: ~30 min per unique email

If the "Release Requests" option is enabled in the global config, users of Darktrace/ EMAIL should expect to receive requests to "release" emails. Additionally, users can go to security teams directly with emails they believe should have been actioned more harshly. The simplest way to deal with both is to apply manual actioning.

Individual Email Action

Time allocation: ~10 min per unique email

To manually hold or release an individual email, select an email from the tray and then click on the "Remediate" button as seen on Figure 17.

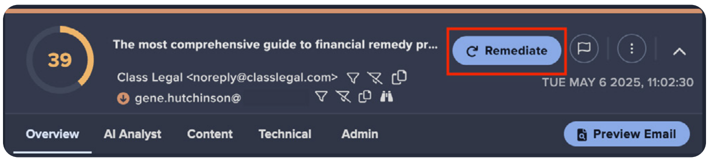

Figure 17: Location of "Remediate" button.

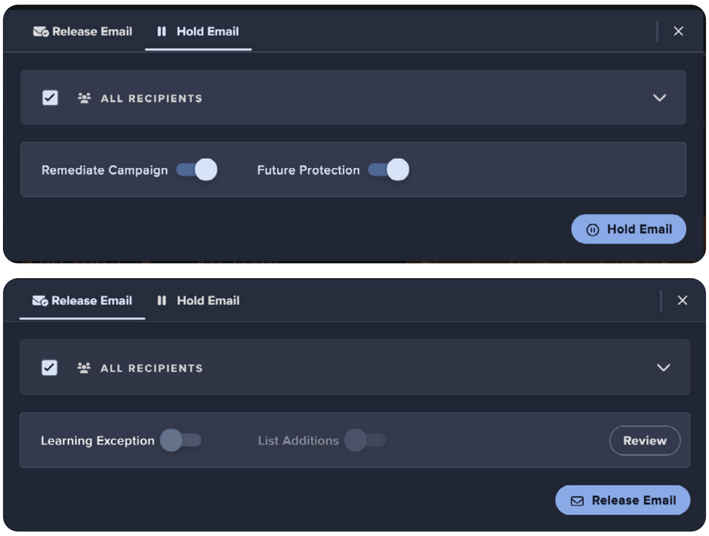

Figure 18: Hold and release dialog options

Once clicked, you will be presented with a dialog box with two tabs as seen in Figures X: hold and release. The recipients of either action can be configured via the dropdown arrow on the "All recipients" option for both.

By hovering over additional options, you will see a dialog box with additional information on the action. Learning exceptions are covered in more detail in the subsequent section and list additions will be covered in more detail in layer 2 of this section.

Learning Exceptions

Time allocation: ~10 min per unique email
Prevent specific actions against certain senders, depending on how their correspondence has been evaluated, with learning exceptions. Open the remediate window for any email and navigate to the "Release Email" pane. Optionally choose to release the emails (selected by default). Then select "Learning Exception". Click "Review" to get a written summary of the changes caused by the learning exception.

Top Tip

Use learning exceptions when you would still like to monitor the correspondence of a certain sender. Learning exceptions have been designed to produce strategic exceptions, whilst not disabling all monitoring of the addresses involved.

The aim is to still provide protection in the event of address spoofing or supply chain compromises, for example.

Bulk Actions

Time allocation: ~10 min per unique email
Perform bulk actions to efficiently action multiple emails. These actions allow you to flag, unflag, release or hold emails, and export logs [14].

Bulk actions can be applied by clicking on the button on the right-most icon of the Quick Filters panel, as seen in Figure X. All emails that are currently being filtered for will receive the action.

Top Tip

If you find yourself having to consistently apply manual actions, use the knowledge from the explore section of this knowledge to build a profile of the types of emails that receive the most manual actions and apply the next layers workflow to that profile. That way, changes in actioning will be more generalized and result in less end user reporting and less administrative work for the tool user(s).

Users should now move to Part 1 Layer 1.

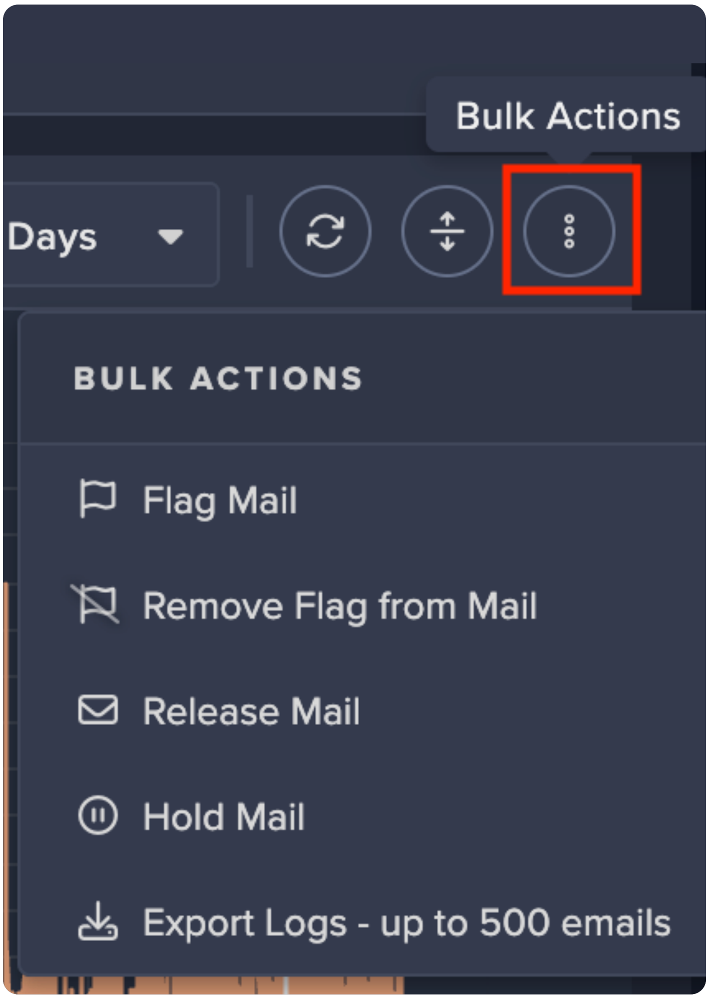

Figure 19: Bulk actions dropdown menu.

Part 2, Layer 2:

Proactive action optimization

In this section, we begin examining proactive ways to change Darktrace/ EMAIL behavior.

List Additions

Time allocation: ~20 min per entry
Leverage list entries to prevent or augment model evaluation.

They can be subdivided into the following categories:

- **Permit:** These are used to reduce actions against specific senders, links, or attachment types.
- **Restrict:** These are used to increase actions or block senders entirely.
- **Dynamic:** These are used for automatic tracking of certain pieces of information about the local email environment but can still be manually edited if needed.
- **Advanced:** This type covers various lists - it includes those used for increasing or decreasing actions which were not included under "Restrict" or "Permit" because they may not be considered relevant for most customers.
- **Custom:** Lists created locally that are personalized to the environment.

Each list contains a description of what the valid inputs are and the impact of each entry would have. Additionally, you can see the models and recipes that use the list additions input under the "Models Using", "Models Inserting", and "Recipes Using" tabs within the list description.

There are three main ways to add list additions:

- **Add list additions manually.** Menu -> Detection -> List. Then select a list from any sub-tab and click on it. Review the list description to verify effects and proper inputs. Then click “Add to current entries”, optionally selecting an expiry time and commit the change.
- **Add list addition via learning exception.** Create a learning exception for any email as described in Part 2, Layer 1. Click on the “Review” button to see the specific list additions advised. Then, hover over the “i” symbol to get a quick summary of the list description.
- **Confirm automatic list additions.** Menu -> Detection -> List -> Dynamic Lists -> Click on checkmark symbol on any of the proposed list additions.

Top Tip

Include known correspondents in an Association Boost list to reduce actioning, and unknown ones in an Association Drain or Block list to increase actioning. If certain users are expected to receive emails from unknown correspondents (such as a sales team), add recipient addresses to the Public Facing Inbox list. This is more effective than continuously adding senders to Association Boost.

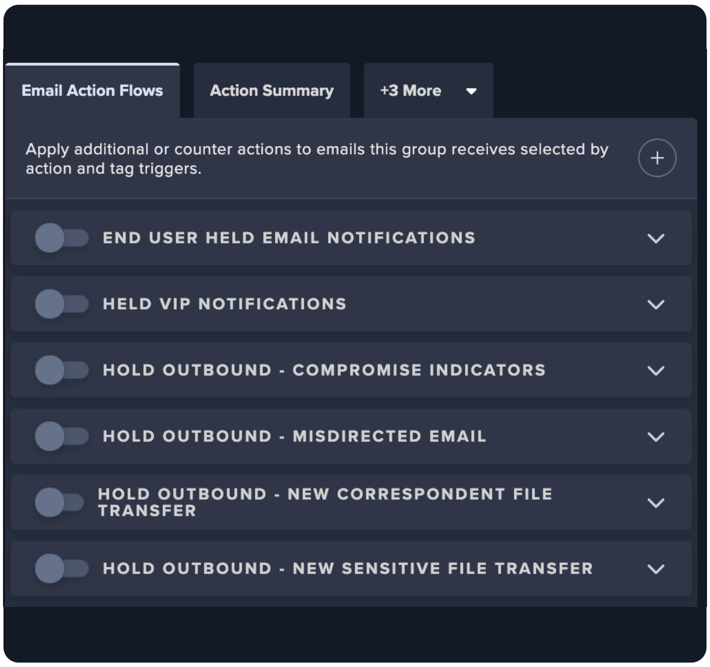

Apply additional or counter actions to emails this group receives selected by action and tag triggers.

END USER HELD EMAIL NOTIFICATIONS

HELD VIP NOTIFICATIONS

HOLD OUTBOUND - COMPROMISE INDICATORS

HOLD OUTBOUND - MISDIRECTED EMAIL

HOLD OUTBOUND - NEW CORRESPONDENT FILE TRANSFER

HOLD OUTBOUND - NEW SENSITIVE FILE TRANSFER

Figure 20: Examples of disabled action flows

Action Flows

Time allocation: ~60 min for All Users group per flow
Unlike list entries and learning exceptions, Action Flows trigger in response to Tags or Model Actions, and they are configured at the Group level. In other words, an Action Flow essentially states: “if an email receives a certain action or tag, then apply a consequent specified second action or counter-action”. Note that Action Flow triggers are “AND” gated, meaning that all triggers must be true for the Action Flow to execute.

- **Open** the groups page. Menu -> Groups
- **Filter** by “Include” groups by selecting the “Include” sub-tab.
- **Select** a group, then select the “Email Action Flows” tab on the right
- **Click** on the drop down arrow on the right of existing action flow names to check and edit details.
  - Alternatively, click on the plus sign below the “Email Action Flows” tab to create a new action flow
- **Click** on “Enabled” to activate action flow.

Workflow Tip

There are six pre-made stock Action Flows that satisfy common use cases. By default, these Action Flows are disabled, however they can easily be activated within the Groups page.

01 **Junk** “Cold Call” emails

02 **Move** “Graymail” to a “Non-Productive Mail” folder

03 **Notify** security team of emails held for a VIP

04 **Notify** security team of outbound emails labeled as Potential Data Loss

05 **Add** a warning to the subject of emails from new senders

06 **Notify** recipients of Held emails

Users should now move to Part 1 Layer 2.

Part 2, Layer 3:

Further customization via models

Model editing ensures email actions align with the user priorities. The end goal is to reduce unwanted actions as much as possible and increase actioning on identified threats. Note it is highly recommended that users NOT edit stock model logic or recipes directly due to their high interdependence. As such, this section only tackles changing model actions.

Identify models suitable for editing

Time allocation: ~240 minutes monthly

Identify models that are triggering on expected correspondence or models that almost exclusively trigger on malicious correspondence. Before attempting model edits, verify that any of the solutions in Part 2, Layer 2 would not achieve the same goal. Users should then use their own judgment on what constitutes a suitable model for optimization.

- **Review current learning exception list.** Open the Learning Exceptions list via Menu -> Detection -> Learning Exception, filter for emails of the senders, and review the models that caused actions on the emails. Filter for those models in the Email Logs tab and review other emails that model has triggered on.
- **Review emails that users have requested release on (if release requests enabled), marked safe/unsafe.** Open the email logs page by navigating to Menu -> Logs. Select the "End Users" quick filter option and select either "Release Requested", "User Safe", or "User Reported". Review the models triggered in the filtered emails, note models of interest, then review the emails where those model(s) were seen triggering.
- **Map threats to model detection.** Identify a particular email threat of interest such as spoofing. Filter for models in the email logs corresponding to that type of detection and review the actioning applied by those models.

Model action changes

Time allocation: ~60 minutes monthly

Once a list of models has been identified, users can increase/decrease the existing actioning applied on said models to fit their use case. This section assumes that severity settings from Part 2, Layer 0 are properly configured. Model actions can be changed by navigating to any one model via Menu -> Detection -> Models. Below are some of the most common action changes users may wish to apply.

- **Change anomaly contribution.** For models that have a "Log" action category besides "Info", adjusting "Bad" or "Good" levels will change the anomaly contribution of the model proportional to the level applied. I.e. shifting from Bad (low) to Bad (high) will significantly increase the anomaly contribution of a model when it fires.
- **Change delivery action.** All models at a minimum will have a "Deliver" action. This can be changed to "Hold", "Move to junk", or "Move to folder" if users want the model to apply a delivery action. Anomaly thresholds can be adjusted for delivery actions. Alternatively, the delivery action can be changed to "Do not hold or alter", "Do not hold", or "Do not junk" if delivery actions want to be prevented for the specific model.
- **Change payload action.** Only models of the "Link" or "Attachment" type have the option of applying individual payload actions. Anomaly thresholds can be adjusted for delivery actions. Alternatively, payload actions can be set to "Take no action on link" for link type models and "Take no action on attachment" for attachment type models.

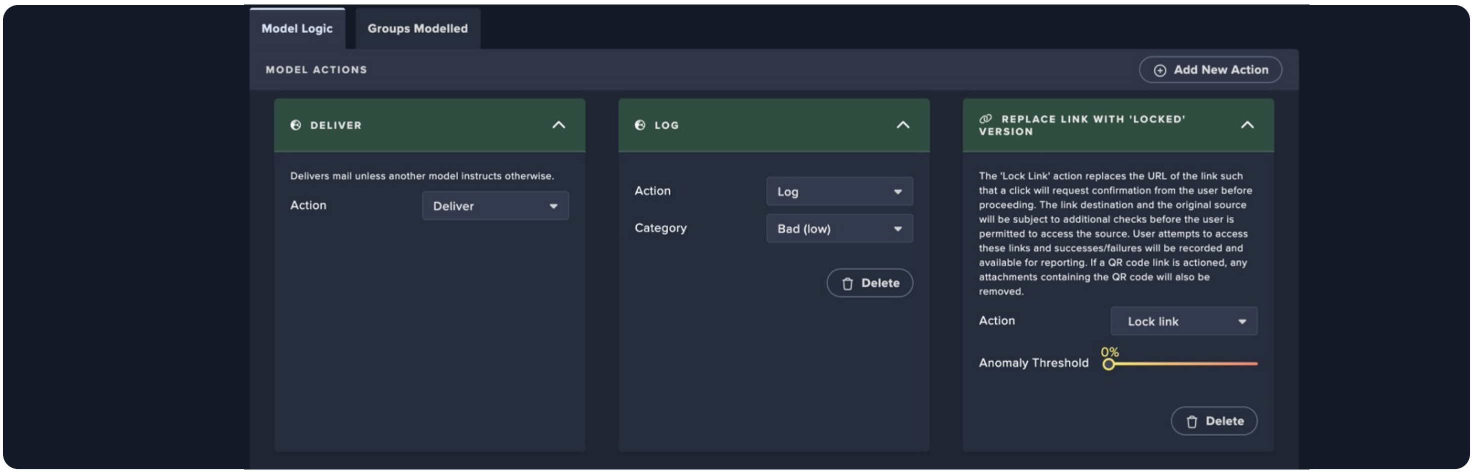

Delivers mail unless another model instructs otherwise.

Action: Deliver

Action: Log

Category: Bad (low)

The 'Lock Link' action replaces the URL of the link such that a click will request confirmation from the user before proceeding. The link destination and the original source will be subject to additional checks before the user is permitted to access the source. User attempts to access these links and successes/failures will be recorded and available for reporting. If a QR code link is actioned, any attachments containing the QR code will also be removed.

Action: Lock link

Anomaly Threshold: 0%

Figure 21: Action configuration for the model "Hidden Link To File Storage"

Contacting Security Operations Support

Time allocation: ~30 minutes per ticket

As it is not advised to edit model logic directly, when a more bespoke solution is required that can't be addressed efficiently by any of the previous sections, a Security Operations Support ticket should be created so a member of the analyst or EMAIL model development team can assist with changes.

When creating a ticket, the following information should be included so that Darktrace can assist as quickly as possible:

- **Clear problem statement.** State whether Darktrace/ EMAIL needs to be actioning an email or category of email more/less.
- **Provide UUIDs or other identifying information of emails.** The UUID of any email can be acquired by navigating to the three dots next to a selected email and selecting "Copy UUID" (as seen in figure below). If the UUID is not provided, any email identifiers (such as subject) should also contain a timeframe.
- **Provide the raw EML file.** As Darktrace staff only have access to the email metadata via the UI, providing the email file provides additional visibility into the issue. This is especially useful for cases in which EMAIL did not provide the expected action. Emails can be downloaded via the Message Actions dropdown menu by users with the appropriate permissions.

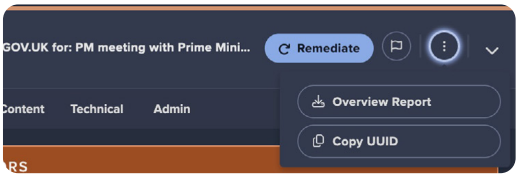

Figure 22: Location of "Copy UUID" button

Workflow Tip

If a member of the Darktrace team creates a custom model as part of the response to a ticket, it will always be contained in a folder called "Site Specific", and the model description and/or commit message(s) will contain a ticket reference. This is so changes to client deployments can be easily found and traceable.

This is the end of the Desktop Workflow.

The Darktrace Mobile App

This section lists a few functionalities available within the Mobile App that can be applied to the desktop workflow. For more information on the EMAIL mobile app, see the following documentation. Note that not all desktop features are available on the Mobile App such as: List Additions and Learning Exceptions

An example workflow using the app is outlined below.

Layer 0:

Dashboard Review

Check the "Overview" tab to gain broad situational awareness: "Menu" > "Overview" icon. Here you can get general information on mail flow and actioning.

Layer 1:

Email filtering and actioning

- Check the "Emails" tab to view full filtered logs. "Menu" > "Emails" icon.
- Tap filter entries to add/modify filters then tap the "Find Emails" button to begin filtering.
- Filtered emails can be manually held or released by swiping left or right (respectively) on them.

Layer 2:

Email details and learning exception

- Tap on a filtered email to open the detail overview.
- If an email has multiple recipients, swiping left and right moves between recipients.
- Scrolling down, the history, association, validation, and content summaries of the email can be found.
- Tapping the ^ arrow on the details page allows users to hold or release the email as well as create a learning exception for the email.

Workflow Process Diagram

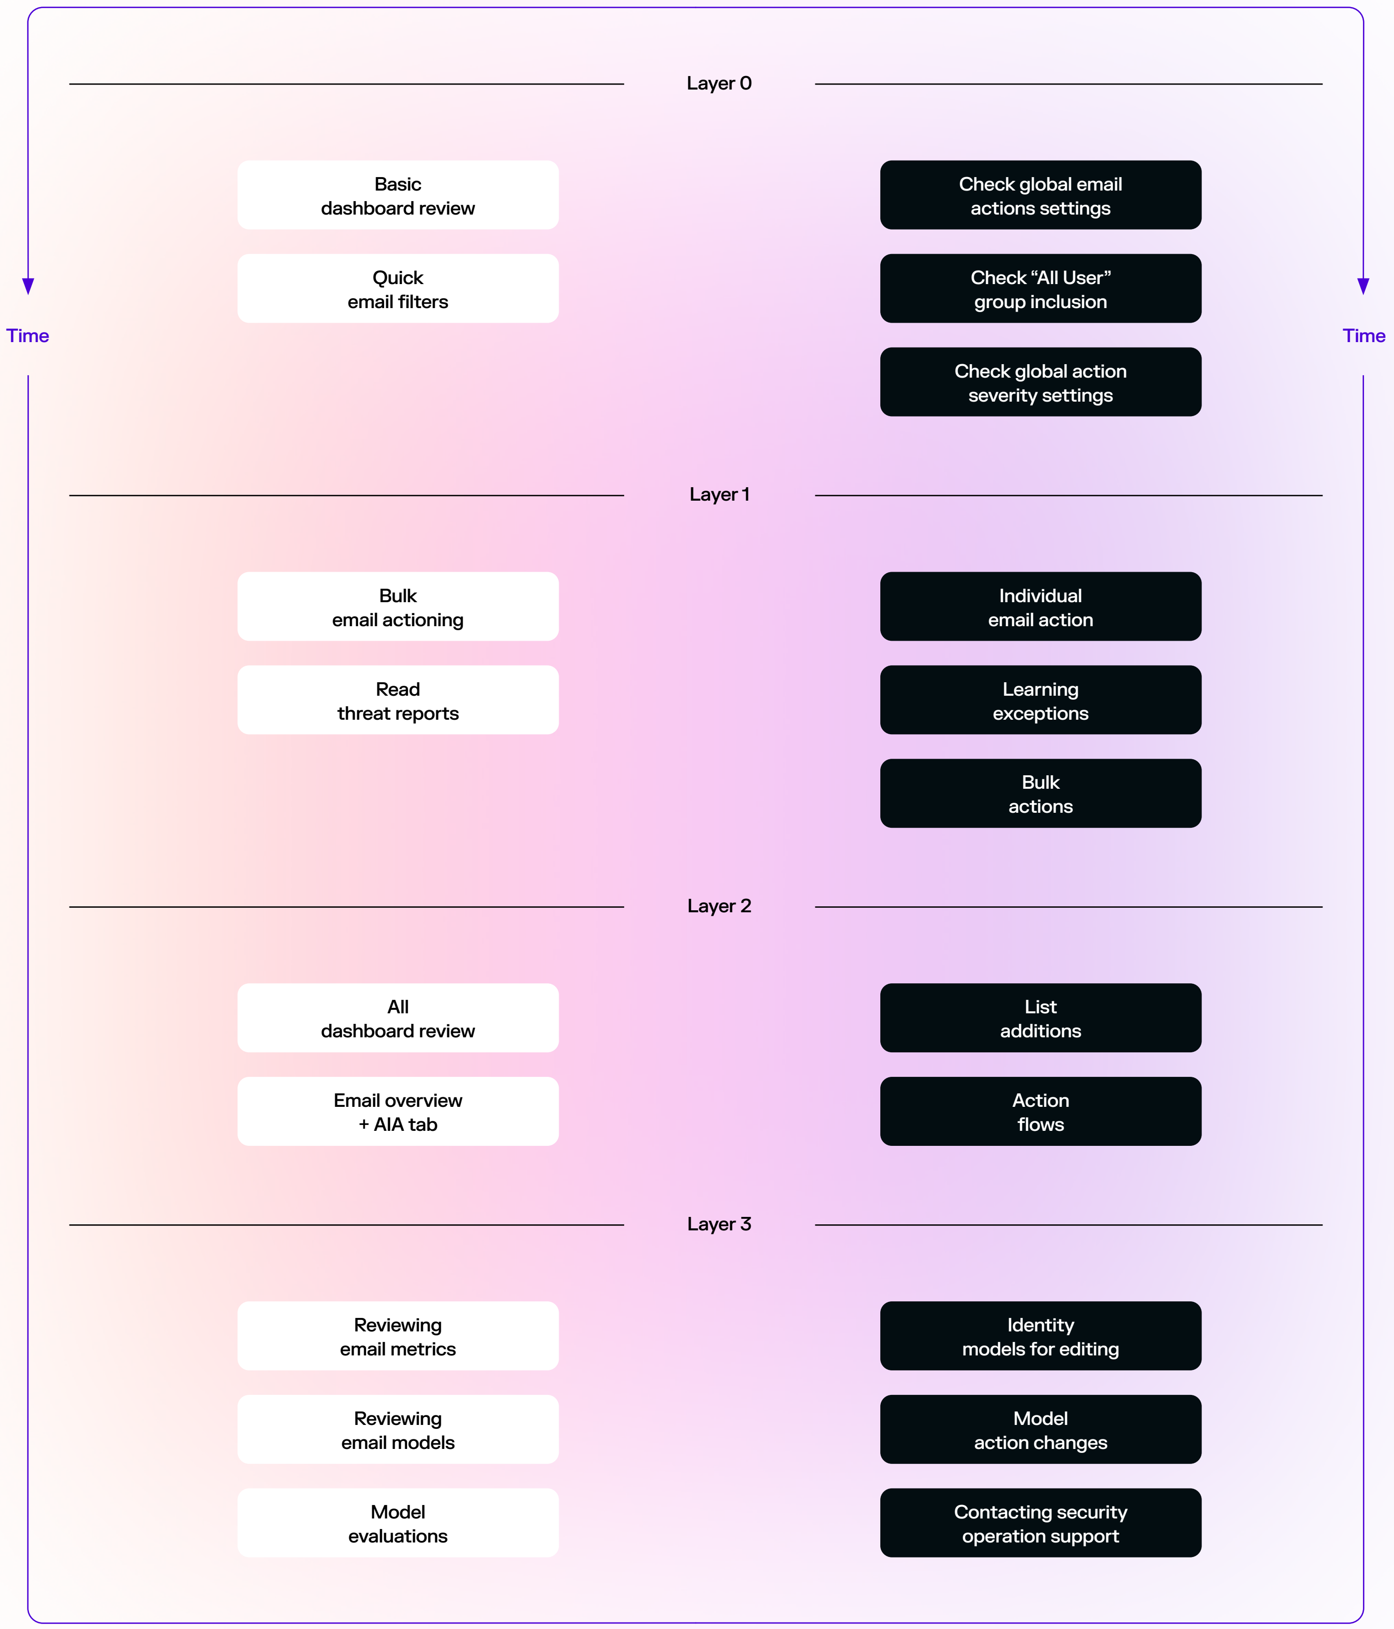

Time

Layer 0

Basic dashboard review

Quick email filters

Check global email actions settings

Check "All User" group inclusion

Check global action severity settings

Layer 1

Bulk email actioning

Read threat reports

Individual email action

Learning exceptions

Bulk actions

Layer 2

All dashboard review

Email overview + AIA tab

List additions

Action flows

Layer 3

Reviewing email metrics

Reviewing email models

Model evaluations

Identity models for editing

Model action changes

Contacting security operation support

Time

□ Exploration Workflow ■ Optimization Workflow

■ **About Darktrace**

Darktrace is a global leader in AI cybersecurity that keeps organizations ahead of the changing threat landscape every day. Founded in 2013 in Cambridge, UK, Darktrace provides the essential cybersecurity platform to protect organizations from unknown threats using AI that learns from each business in real-time. Darktrace's platform and services are supported by 2,700+ employees who protect nearly 10,000 customers globally. To learn more, visit www.darktrace.com.

North America: +1 (415) 229 9100 Europe: +44 (0) 1223 394 100 Asia-Pacific: +65 6804 5010 Latin America: +55 11 4949 7696

darktrace.com | info@darktrace.com