




Darktrace Cyber AI

An Immune System for Email

“

More than ever, modern email security requires innovation and a shift in mindset to combat the evolving threat landscape. ”

– Peter Firstbrook, VP Analyst at Gartner



Introduction

Threats at a Glance

Spear Phishing & Payload Delivery	4
WeTransfer Attack	6
Malware Hidden in Fake Invoices	7
Municipality Address Book Compromised	7
Supply Chain Account Takeover	8
Consecutive Supply Chain Attacks	10
Malicious File Hidden in OneDrive Page	13
Social Engineering & Solicitation	14
Impersonation Attack	16
CEO Payroll Request	17
'Finance VP' Spoofing Attack	17
Compromised Employee Credentials	18
Unusual Login at Panamanian Bank	20
Attempted Access from Rural Japan	20
Office 365 Account Compromised	21
Automated Brute Force Attack	21


Email and collaboration platforms represent the connective tissue of any digital business. Information is shared, plans are hatched, and alliances are formed in the digital realm of written correspondence. Yet as a human-driven medium, email will always be fueled by a pervasive assumption of trust that stands as the 'weakest link' in an organization's security strategy.

While this assumption of trust is critical for collaboration and growth, it means that email more than any other area of the business will remain structurally resistant to the modern-day spirit of 'zero trust', and it is therefore hardly surprising that 94% of cyber-threats still originate in email.

To minimize the influence of human fallibility in this area, the industry has generally come around to the idea that technology must be relied upon to identify malicious emails that even the most discerning, well-trained employees fail to spot. However, until recently, traditional defenses have struggled to keep pace with innovations in the cyber-threat landscape.

Spear phishing, impersonation attacks, and account takeovers in particular remain fruitful avenues of attack for cyber-criminals aiming to infiltrate an organization with ease. Targeted email attacks of this kind, together with the limitations of traditional defenses, remain a burning challenge for organizations with even the most layered and mature security strategies.

Peter Firstbrook, VP Analyst at Gartner, summarizes the market dynamic well: "Common controls, such as standard, reputation-based, anti-spam, and signature-based antivirus, are fine for widespread attacks and scam campaigns, but they're not good enough for protection against more targeted, sophisticated, and advanced attacks. More than ever, modern email security requires innovation and a shift in mindset to combat the evolving threat landscape."



Darktrace AI: An Immune System Platform

Thanks to the recent emergence of enterprise-scale AI, however, this 'shift in mindset' has finally crystallized in the form of an 'immune system' approach to email security.

As Firstbrook suggests, traditional email defenses may be adequate for simple and indiscriminate threats, but they are not designed to counter more advanced attacks that have been customized to particular recipients and businesses.

Legacy email gateways and native controls rely on hard-coded rules and a knowledge of historical attacks for detection. Their purview is therefore necessarily confined to threats that have already been seen, or that are at least basic enough to trigger a static and binary rule at the border. But – as many business leaders can tell you by pointing to their scars – this is not the challenge we face.

Fortunately, the paradigm shift that has emerged in email security has fallen out of an important distinction between Firstbrook's 'common approach' and a novel application of enterprise-scale AI. This distinction has been usefully compared to the difference between the 'protective skin' of an organization and its learning 'immune system' for threats that get through.

Whereas your protective skin knows about historical attacks and can stop well-known threats, your 'immune system' knows about the disjointed 'patterns of life' that characterize every employee's digital workflow. Crucially, these 'patterns of life' are made manifest not only in email traffic, but also in network and cloud traffic, and in a way that can be unified into an evolving and comprehensive picture of normality for every user.

This unique, enterprise-wide understanding is enabling organizations to neutralize more targeted attacks than ever before, as it remains the only approach that can furnish enough evidence to accurately determine whether subtle deviations in a targeted email are genuinely malicious.

For the first time, our email defenses can meaningfully ask whether it *would be weird* for a user to receive an email, given what the system knows about the 'patterns of life' of this employee, their peers, and the wider organization not only in email, but also in the cloud and corporate network.

It is also the only approach that can update its decisions and actions in light of new evidence, even after an email has been delivered – whether that evidence is made manifest in email, or in malicious behaviors that emerge in the network.

This White Paper is designed to illustrate why a unified and bespoke understanding of network, cloud, and email traffic represents a paradigm shift in the email security market. Darktrace pioneered this approach with Antigena Email and its Immune System Platform. The case studies that follow will fall into one of four highly sophisticated attack categories that routinely bypass your 'protective skin', but which Darktrace's AI easily neutralizes in seconds:

- Spear phishing & payload delivery
- Supply chain account takeover
- Social engineering & solicitation
- Compromised employee credentials

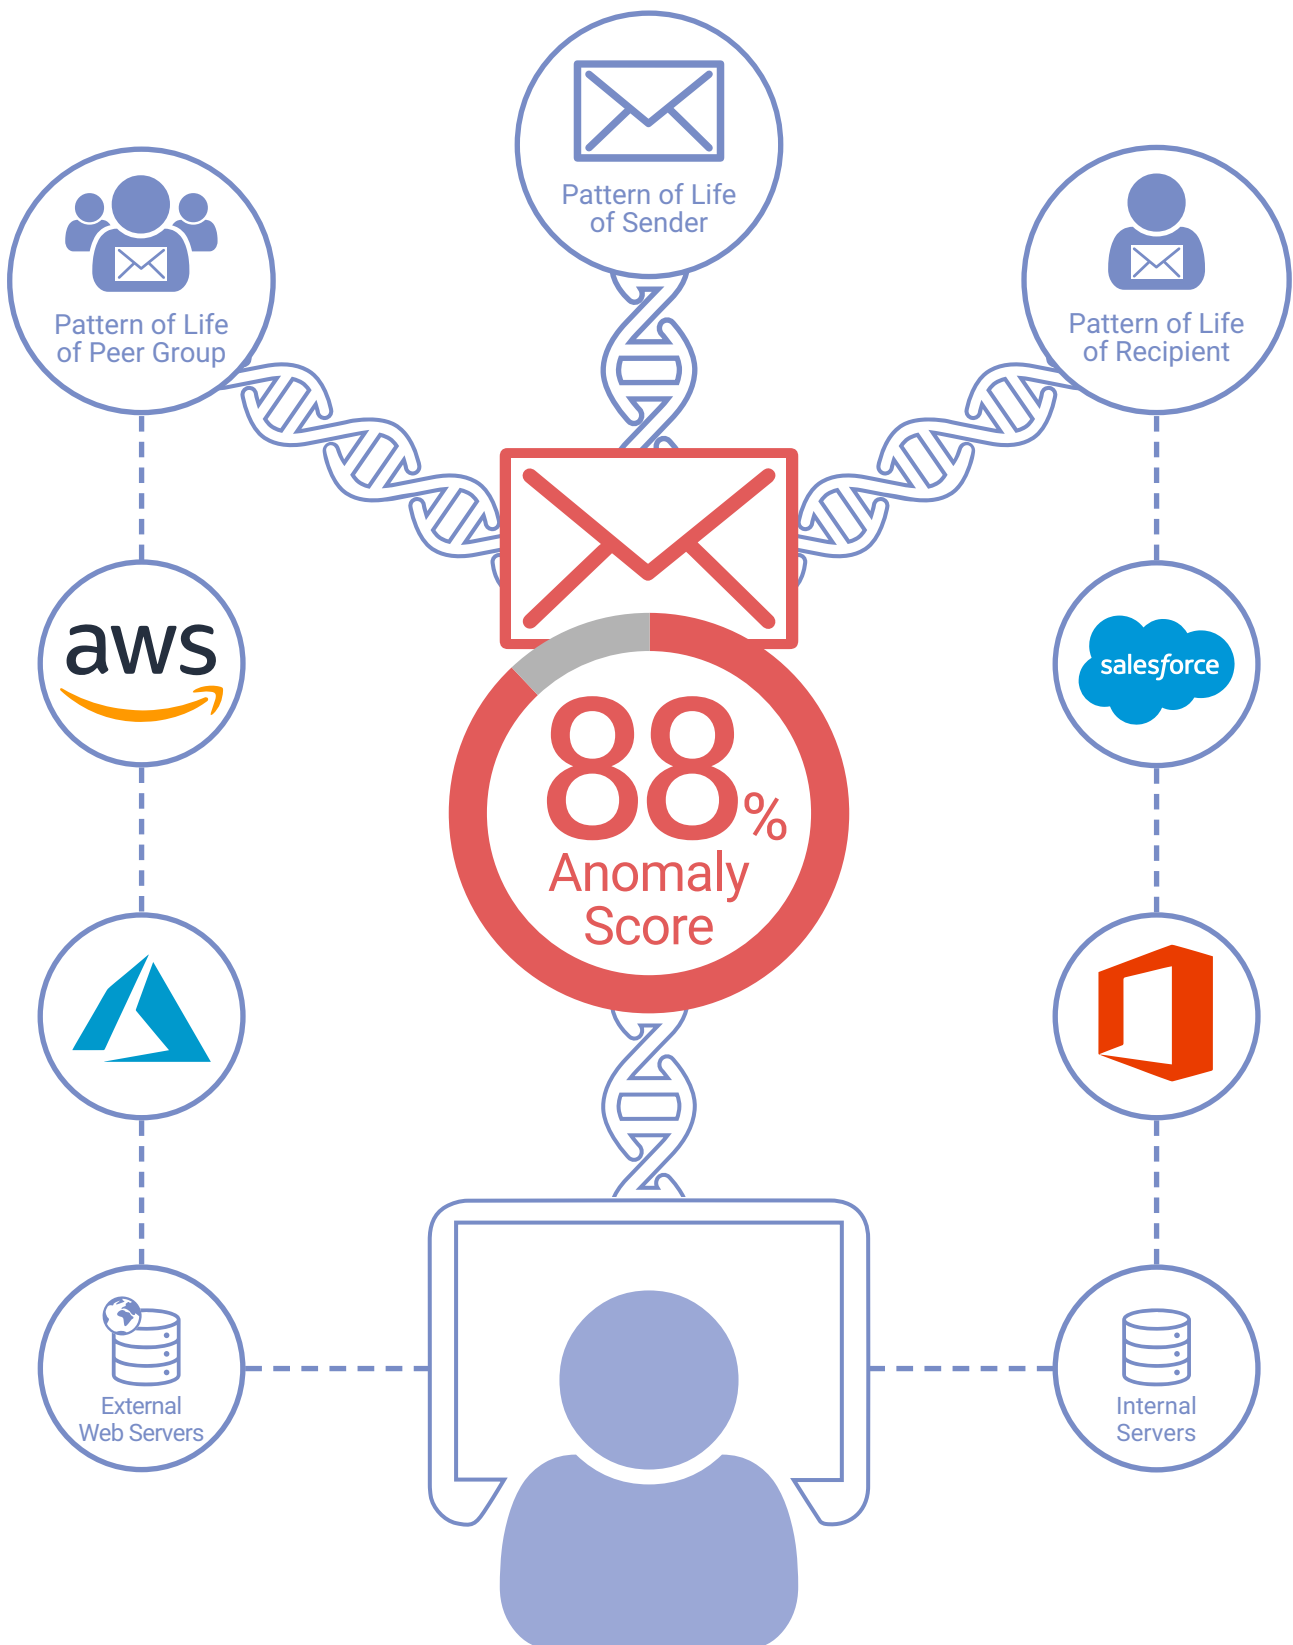


Figure 1: Antigena Email is the only solution that analyzes emails in the context of the wider organization – not just email data. This enterprise-wide understanding allows it to spot malicious emails that evade traditional defenses at the border.


Spear Phishing & Payload Delivery



Antigena Email has been incredibly valuable in catching threats with its understanding of 'normal' for both email and network traffic.

– Head of IT, Entegrus

1 in every 99 emails is a Phishing Attack



Source: Avanan

Most phishing campaigns attempt to deceive users into clicking malicious links or attachments in an email, with the ultimate goal of harvesting credentials or deploying destructive malware in an organization. These attacks can be launched either as indiscriminate 'drive by' campaigns against thousands of organizations, or as crafted, 'spear phishing' attacks that are customized to a particular recipient or business.

To defend against phishing campaigns, traditional defenses typically analyze emails in light of an understanding of historical attacks, blacklists, and signatures. Yet cyber-criminals understand this reactive approach better than anyone, and they have every incentive to leverage novel tactics and techniques that evade legacy defenses by design.

However, while these attacks have never been seen before and will therefore evade traditional defenses at the border, this means that at some level of description they will be highly anomalous for the targeted user or business – at least if 'patterns of life' for the entire digital environment are taken into account. This basic truth is precisely why bridging the traditional security knowledge gap between the external email layer and wider network, as Darktrace's Immune System Platform has, is so critical.

With enterprise-scale AI, Antigena Email can analyze links, attachments, domains, content, and other elements of an email alongside 'patterns of life' in the cloud and network, correlating a rich constellation of data points that reveal seemingly benign emails to be unmistakably malicious.



Unlike any other solution, Antigena Email and the Immune System can correlate network, cloud, and email data to identify whether domains associated with a payload and sender are abnormal, the location of a link in an email is strange, the topics of discussion and content are unusual, and even whether patterns in the URL pathway are suspicious.

This fundamentally unique approach means that Darktrace's decision-making is drastically more accurate than that of other tools, such that it can take highly proportionate and targeted actions to neutralize phishing attacks at scale.

The Immune System is also in the unique position of being able to detect an infection in any environment, and automatically perform a root cause analysis to see if it originated via email. If so, it will instantly protect all other employees targeted by the same attack. We call this strategic autonomous response – where learning from Patient Zero enables the strategic protection of the rest of the business without human intervention. From a security team's perspective someone still needs to clean up the laptop of the first victim, but that is much better than cleaning up 200 or worse.

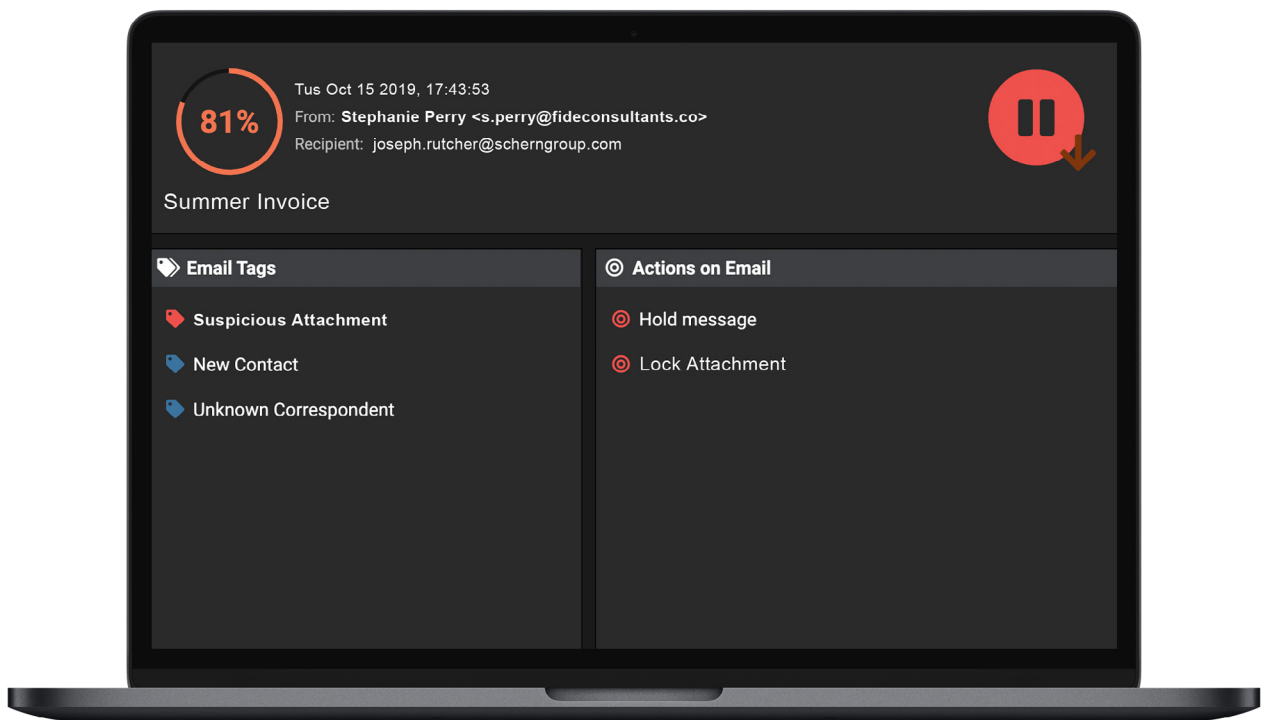
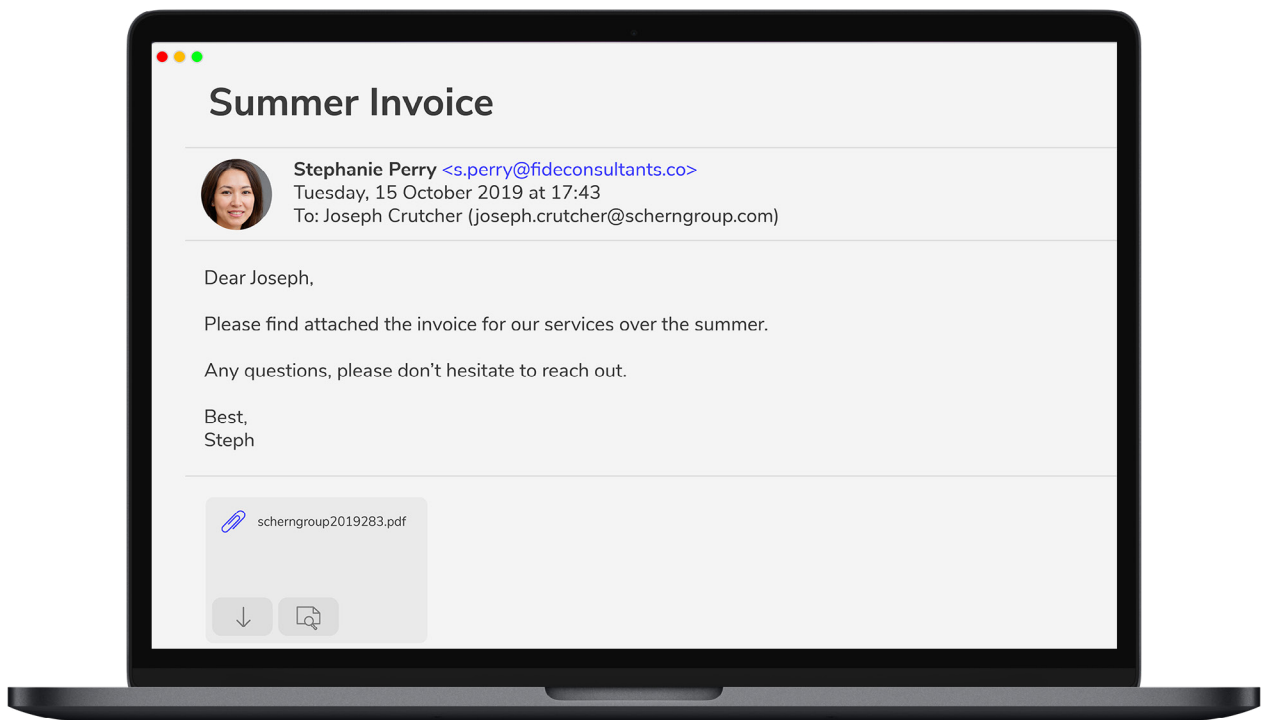


Figure 2: An email coaxing an employee to click on an attachment containing a malicious payload, and the corresponding view within Darktrace's user interface, showing the anomaly tags and actions taken.

WeTransfer Attack

Darktrace detected a phishing attack targeting five high-profile users of an academic organization in Singapore, carefully constructed to trick them into clicking a malicious link.

Antigena Email assigned these emails a 100% anomaly score and took action to 'Hold' them back, preventing delivery. It also identified the subtle indicators of service spoofing, despite the organization having a known relationship with the sender.

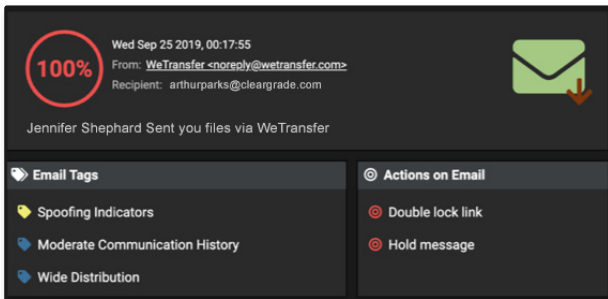


Figure 3: The user interface showing the model breaches and actions

1. From the header data, there were no clear signs that this email had any source other than WeTransfer, and would have appeared perfectly normal to the recipient. The 'Width' and 'Depth' indicate that this email address has communicated with many people in the organization, across multiple days.



Figure 4: The connection data of the relevant emails

2. However, Antigena Email was able to pick up on a range of subtle anomalies given its understanding of 'normal' for the user and organization, together with additional context gained from the network layer.

a. First, the 'Address IP Anomaly Score' was high (63%). This metric indicates how unusual it is for this email address to send from this IP given historical sending patterns, and it is typically an indication of a spoof or hijacked account.

b. In addition, as Darktrace is constantly modeling 'normal' behavior for every external sender, it was able to pick up on a key anomaly in the body of the email – a link that was highly inconsistent with what Darktrace had seen from WeTransfer previously, allowing Antigena Email to identify it as the malicious payload in the email.

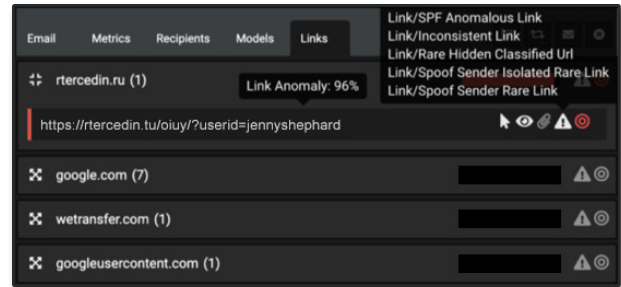


Figure 5: A breakdown of the links shown in the emails

c. The link in question was given a 96% anomaly score, and it was hidden behind 'click here'-style buttons in several parts of the email, including a fake 'https://wetransfer.com/...' link (pictured below) and the text 'Inquiry Sheet.xls' and 'Get Your Files'.

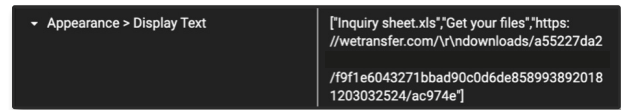


Figure 6: Antigena was able to determine where the link appeared within the email

This attack was completely novel and bypassed every other signature-based tool that the university had in place. Equally, because the link utilized a completely benign domain and did not lead to an obviously malicious payload, even heuristic detection and sandboxing likely would have failed.



Malware Hidden in Fake Invoices

A major law firm became one of the key targets in an advanced phishing campaign, which sought to disguise credential-stealing malware within ISO files attached to fake invoices. Traditional email defenses typically whitelist ISO files, while operating systems automatically mount their images upon a single click, affording them an obvious appeal for threat actors.

Yet when a score of the offending emails got past the firm's traditional email defenses, Darktrace caught the campaign by recognizing a wide range of anomalous indicators. For instance, one of the AI models that the emails triggered was "Attachment/Unsolicited Anomalous MIME," which means that the MIME type of the attachment was highly unusual for the user and their peer group, and that the recipient had never communicated with the sender to request the file.

By pinpointing the precise provenance of the threat, Darktrace took surgical action to disarm it, rather than merely marking all potentially suspicious emails with generic warnings that were likely to be ignored. To counter the harmful ISO files, Darktrace converted the attachments into harmless PDFs and moved the emails to the junk folder. And crucially, upon detecting the first email in the campaign, the technology automatically neutralized 20 others before they had a chance to impact the business.

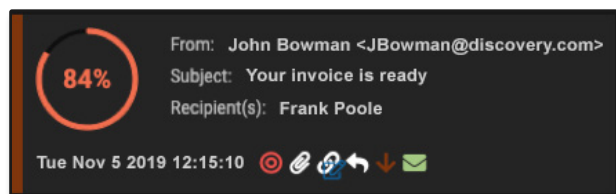


Figure 7: Header of the malicious emails, showing suggested action

Municipality Address Book Compromised

A threat actor managed to get hold of the address book of a US municipality, delivering an attack to recipients alphabetically, from A to Z. Each email was well-crafted and customized to the recipient, and the messages all contained a malicious payload hiding behind a button that was variously disguised as a link to Netflix, Amazon, and other trusted services.

When the first email came through, Darktrace immediately recognized that neither the recipient nor anyone in his peer group or the rest of the city's staff had visited that domain before. The system also recognized that the way the links were hidden behind each button was highly suspicious. It raised a high-confidence alert, and suggested autonomously locking each link as it entered the network.

Interestingly enough, the fact that Antigena was deployed in 'Passive Mode' provided plain and concrete evidence of the system's ability to thwart subtle attacks that other tools miss: while Antigena spotted and sought to neutralize the campaign at the letter 'A', the security team's legacy tools woke up to the threat at 'R'. In 'Active Mode', Antigena would have neutralized the attack before it could reach a single user.

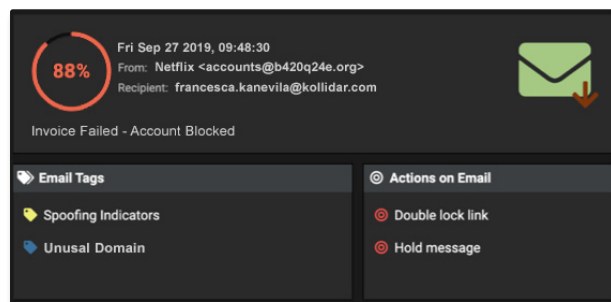




Figure 8: Antigena Email showing an 88% anomaly score



Supply Chain Account Takeover

Account takeover losses have more than tripled in the last year to \$5.1 billion



Source: Javelin

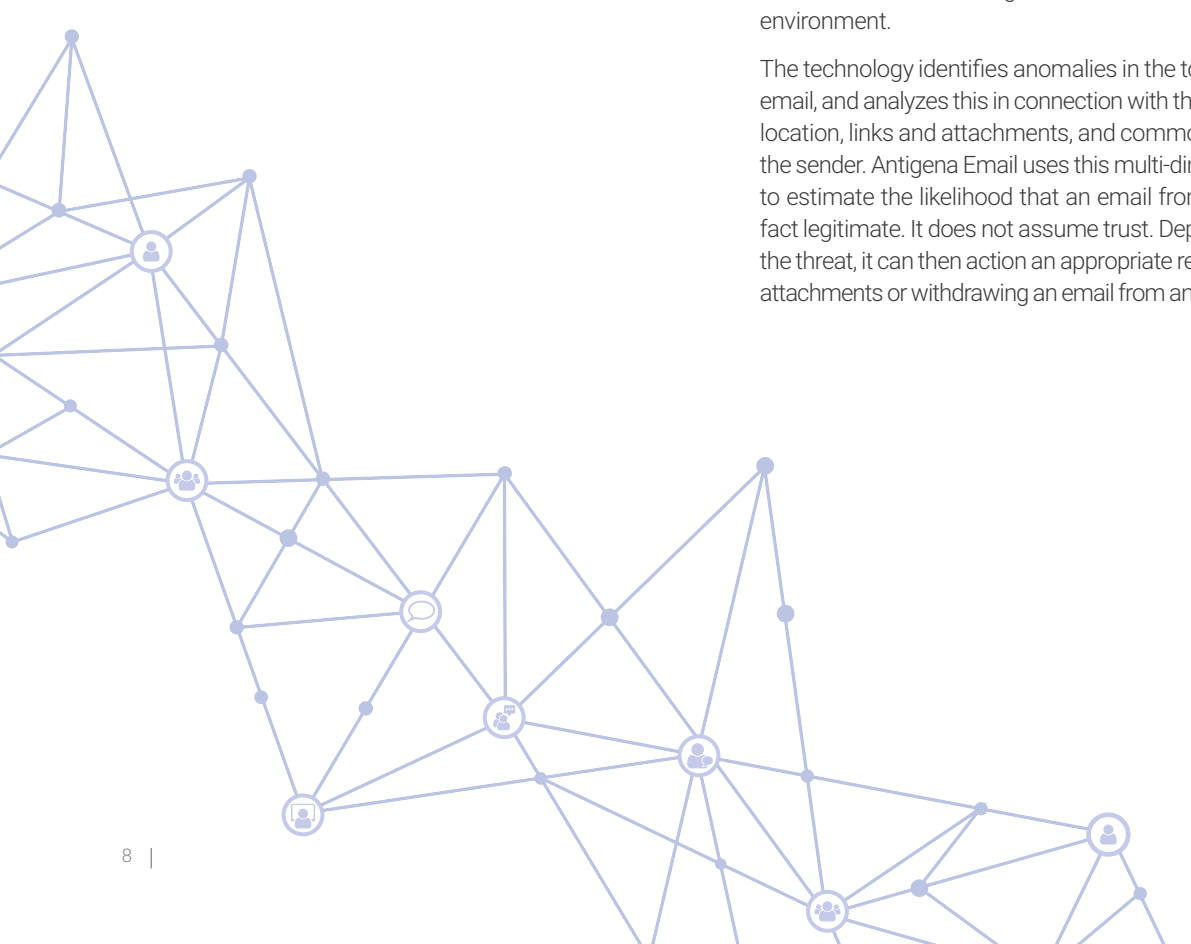
By hijacking the account details of a trusted contact in your supply chain, threat actors can easily gain the trust of a recipient in the network and coax them into clicking a malicious link or transferring millions out of the business. Legacy email defenses assume trust, which means that sophisticated account takeovers often go completely unnoticed.

Compromised accounts have been responsible for several high-profile attacks on large organizations in recent years. Cyber-criminals are increasingly leveraging supply chains – comprised of vendors, partners and contractors – in their attacks to infiltrate an organization or establish offline communication. Earlier this year, a report on so-called ‘island hopping’ – where attackers try to expand on a breach through supply chains – found that this method accounts for half of today’s attacks.

Attackers who have total access to a supplier’s email account are able to study previous email interactions and produce a targeted response to the latest message. The language they use will often appear benign, so legacy email security tools searching for key words or phrases indicative of phishing will fail to pick up on these attacks.

Antigena Email is able to formulate a comprehensive notion of word normality for every internal user, so regardless of how plausible the phrasing might be to most observers, human or machine, it can identify irregular distributions of words and phrases. Analyzing patterns of communication with the full context of all email and network traffic, Antigena Email uses a range of metrics to confidently identify cases of account takeover, something that is impossible to detect without a detailed understanding of ‘normal’ behavior for the entire digital environment.

The technology identifies anomalies in the topic and content of every email, and analyzes this in connection with the consistency of the login location, links and attachments, and common previous recipients for the sender. Antigena Email uses this multi-dimensional understanding to estimate the likelihood that an email from a trusted supplier is in fact legitimate. It does not assume trust. Depending on the severity of the threat, it can then action an appropriate response, locking links and attachments or withdrawing an email from an employee’s inbox entirely.



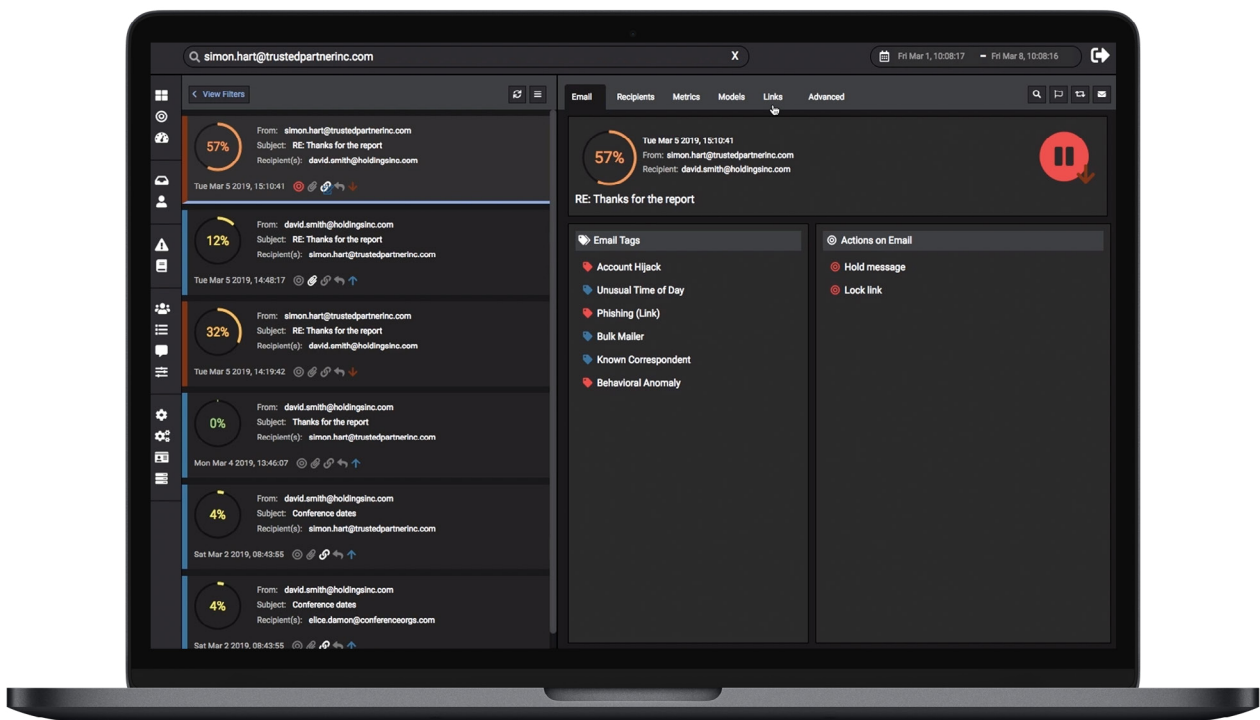
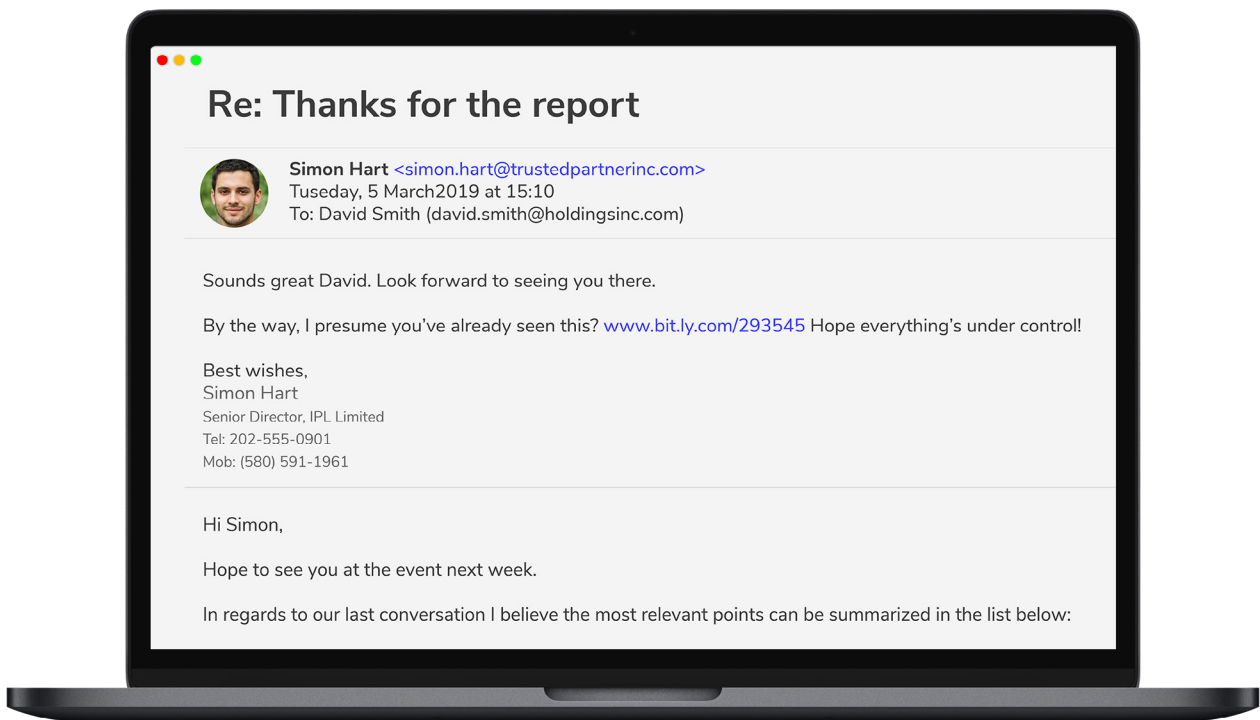


Figure 9: A plausible reply sent from a trusted supplier’s compromised account following a thread of email correspondence. The link contained a malicious payload.

Consecutive Supply Chain Attacks

A customer trialing Antigena Email experienced two serious incidents on successive days, when the email accounts of trusted suppliers became the source of a malicious campaign – very likely after these accounts were compromised.

Antigena Email had not yet been configured to take autonomous actions and so the users were thus fully exposed to the content of the emails. However, in every case Antigena Email advised that it would have held the emails back and double locked the link payloads, while Microsoft’s inbuilt security tools detected nothing suspicious and let everything through without action.

Incident 1 – Consultancy Firm

In the first case, Antigena Email recognized that the sender was well known to the company, with a number of internal users having corresponded directly with them previously. In fact, earlier that day one of these users was engaged in normal correspondence with the soon-to-be-hijacked account. The supplier in question was a UK-based environmental consultancy firm.

Less than two hours after this routine exchange, emails were then rapidly sent to 39 users, each containing a phishing link. There was variation in the subject lines and links contained in the emails, suggesting highly targeted emails from a well-prepared attacker. The purpose of the links could have been to solicit payments, harvest passwords, or deploy malware.

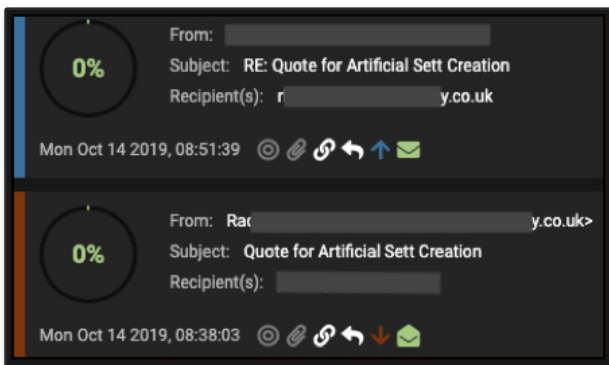


Figure 10: Earlier ‘normal’ correspondence with the sender – with a 0% anomaly score

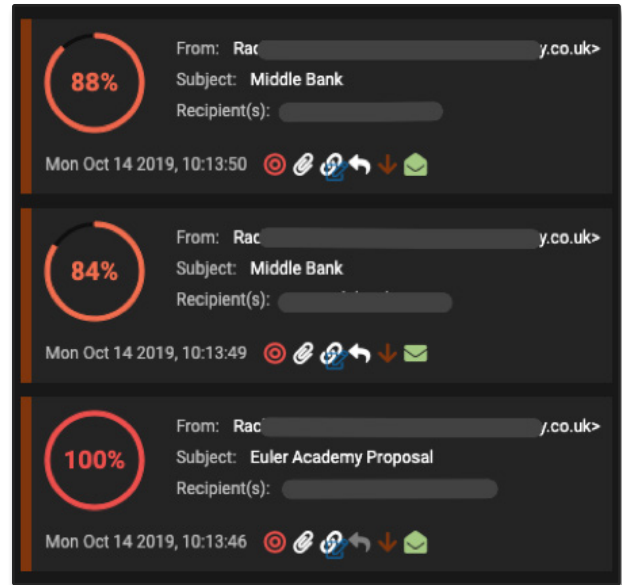
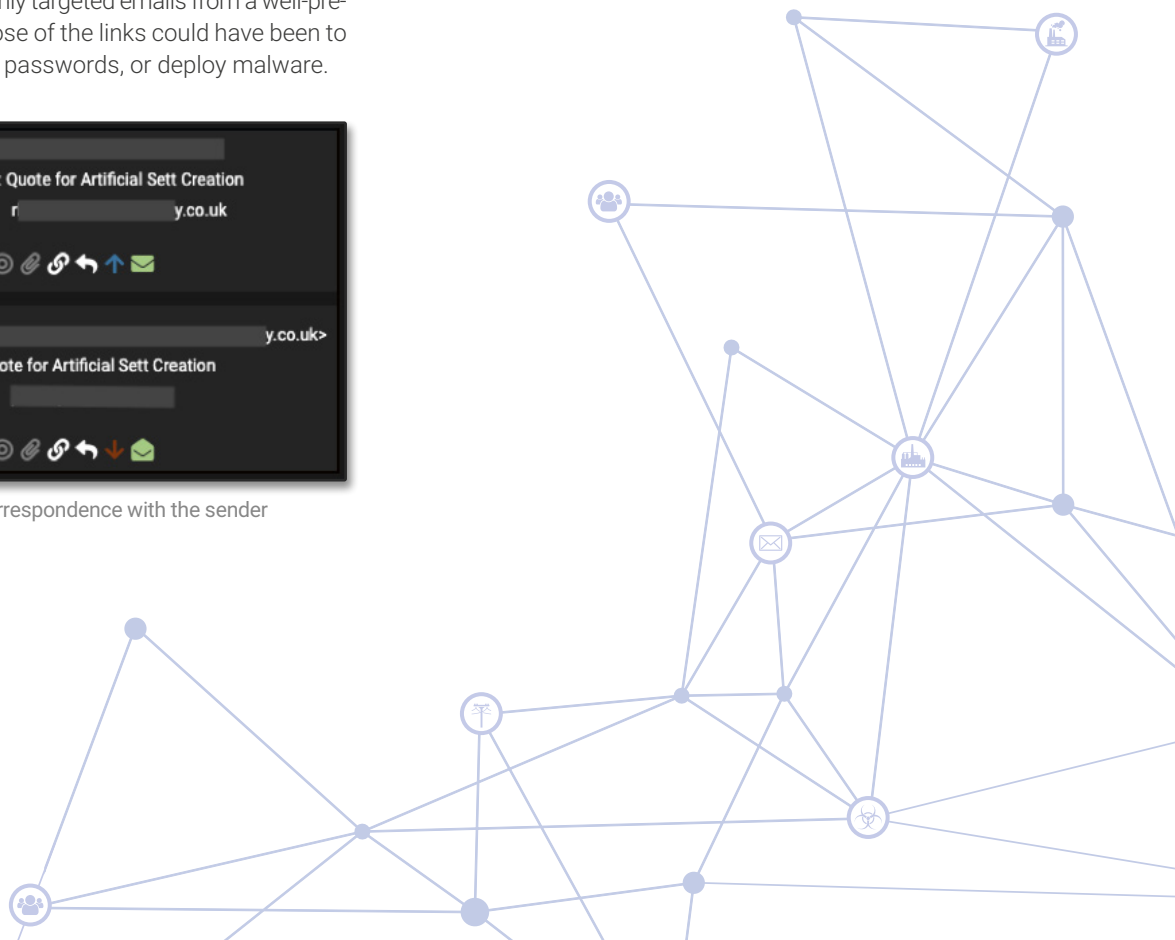


Figure 11: Emails sent later the same day containing malicious attachments



Antigena Email identified the full range of red flags that are typically associated with supply chain account takeovers:

1. Unusual Login Location: Antigena Email determined that the emails had been sent from an authentic Outlook web server. This itself was not unusual for the supplier, but within this connection data it was also possible to extract the geo-locatable IP address, revealing that the attacker initiated their login from an IP in the US, as opposed to their usual login location in the UK.

2. Link Inconsistency: The phishing links contained in the emails were all hosted on the Microsoft Azure developer platform – likely to skirt reputation checks on the host domain. Despite the widely assumed legitimacy of azurewebsites.net across the web, Antigena Email was able to detect that this domain was highly inconsistent for the sender based on previous correspondence history. The unusual subdomain also meant that the hostname had a maximum rarity score in the context of the organization’s network traffic. Because other email security products do not benefit from this contextual intelligence, it would have been impossible for them to come to this conclusion.

3. Unusual Recipients: A recipient ‘association anomaly’ score is assigned to estimate the likelihood that this particular group of recipients would be receiving an email from the same source. Adding context to its investigation over time, Antigena Email deduced that this recipient group was 100% anomalous by just the third email.

Usage > Darktrace Host Rarity	100
Usage > Domain External User Hostnames	0
Usage > Domain Inconsistency Score	88

Figure 12: Metrics triggered by the rarity and inconsistency of the link

4. Topic Anomaly: The subject lines for these emails suggest an attempt to appear low-key and professional, and consequently any signature-based attempts to look for keywords associated with phishing would have failed. However, Antigena Email recognized that these recipients do not typically receive emails about business proposals using this style of phrasing.

Property	Value
Recipient > Metrics > Association Anomaly	100

Figure 13: Antigena Email rapidly detected that this group of recipients was not closely related





Figure 14: The summary view for the Topic Anomaly metric



Incident 2 – Compromised SaaS Provider

A second attack the following day involved emails being sent to 55 internal users from a SaaS provider which was known to the company. In the absence of any actions by Microsoft, over 50% of these emails were read by the recipients. Antigena Email advised that these emails should be held back, preventing them from reaching the inbox.

1. As before, the emails sent from the compromised account each contained a malicious phishing link. In this case, however, the link remained active for a longer period of time, allowing a precise reconstruction of what the end users would have encountered.

2. Fortunately, those who had interacted with the emails were easily found and the accounts recovered, thanks to the shared intelligence of Antigena Email and Darktrace's Immune System Platform in the network. The Immune System could also see that devices on the physical network were connecting to the phishing host. Working in sync with Antigena Email, the Immune System flagged these interactions with suspected phishing domains in the network.

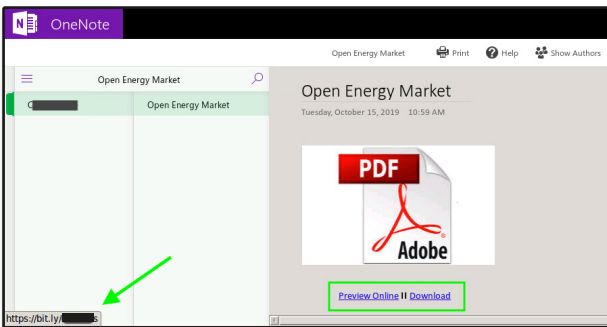


Figure 15: Screenshot exposing a hidden link

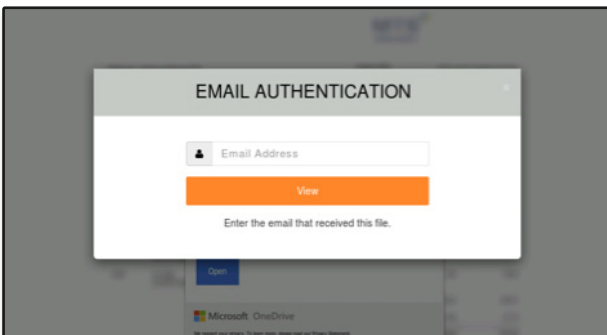


Figure 16: This led to a form which would harvest the user's credentials

3. Although the links were embedded in Microsoft ATP 'safe-links' (meaning that Microsoft would have run a real-time check on the links when clicked by the user) the connections to the actual endpoints in network traffic confirmed that the intelligence available to Microsoft at the time led it to conclude that the links were safe, exposing the users to the malicious endpoint.

4. The link itself was hosted on the well-known file sharing platform SharePoint. Upon visiting the link the user was taken to a document which presented itself as a report on the energy market. However, a button soliciting the user to download the file redirected to another convincing webpage which was set up to solicit the user's email and password – and send them straight to the attacker.

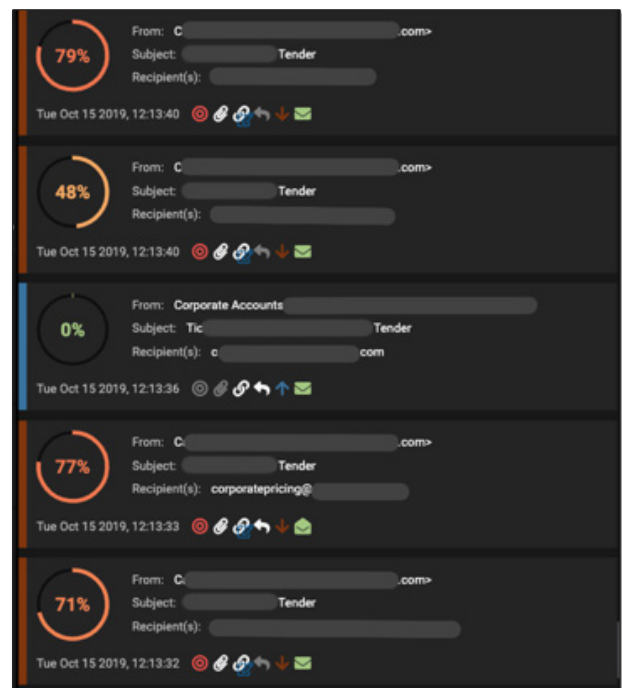


Figure 17: Emails from Incident 2 as they appear in the Antigena Email console, including those that were sent outbound in response. It reveals the 'corporate accounts' user acknowledged the email by opening a ticket.

Malicious File Hidden in OneDrive Page

An advanced threat actor hijacked the email account of a supplier for a large hotel group, using the trusted account to send a malicious payload into the organization. While the attack managed to evade the company's legacy defenses, Antigena Email neutralized the threat in seconds.

1. Analysis of a previous email reveals Antigena Email's understanding that there was a relationship between the two senders.

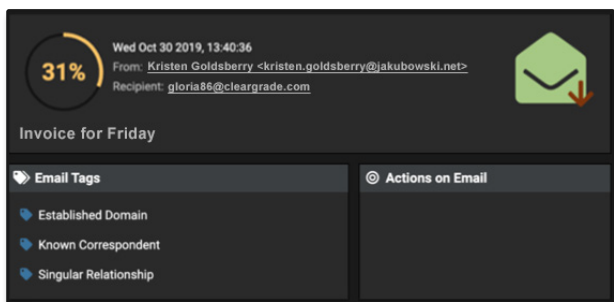


Figure 18: An example of a previous communication

2. A subsequent email was flagged as highly anomalous compared to the sender's previous communication patterns.

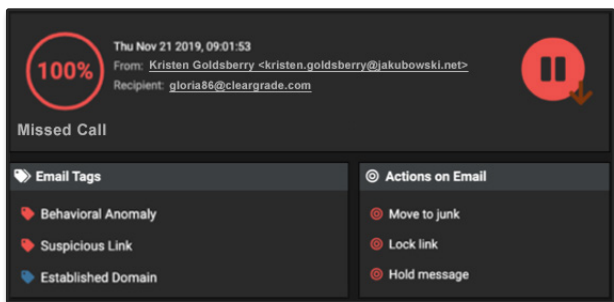


Figure 19: A later email tagged and three associated model breaches

3. As we can see, these emails were all tagged with the 'Behavioral Anomaly' model, and Antigena Email decided that the best action to take was to hold these messages back from the intended recipients.

4. Antigena Email identified multiple deviations from the normal 'pattern of life' of the external sender, including 'Anomalous Source Country' and 'Anomalous Source IP address'.

5. The malicious link in the email was also highly inconsistent with the company's 'patterns of life' across email and network traffic, and hence was locked by Antigena Email.


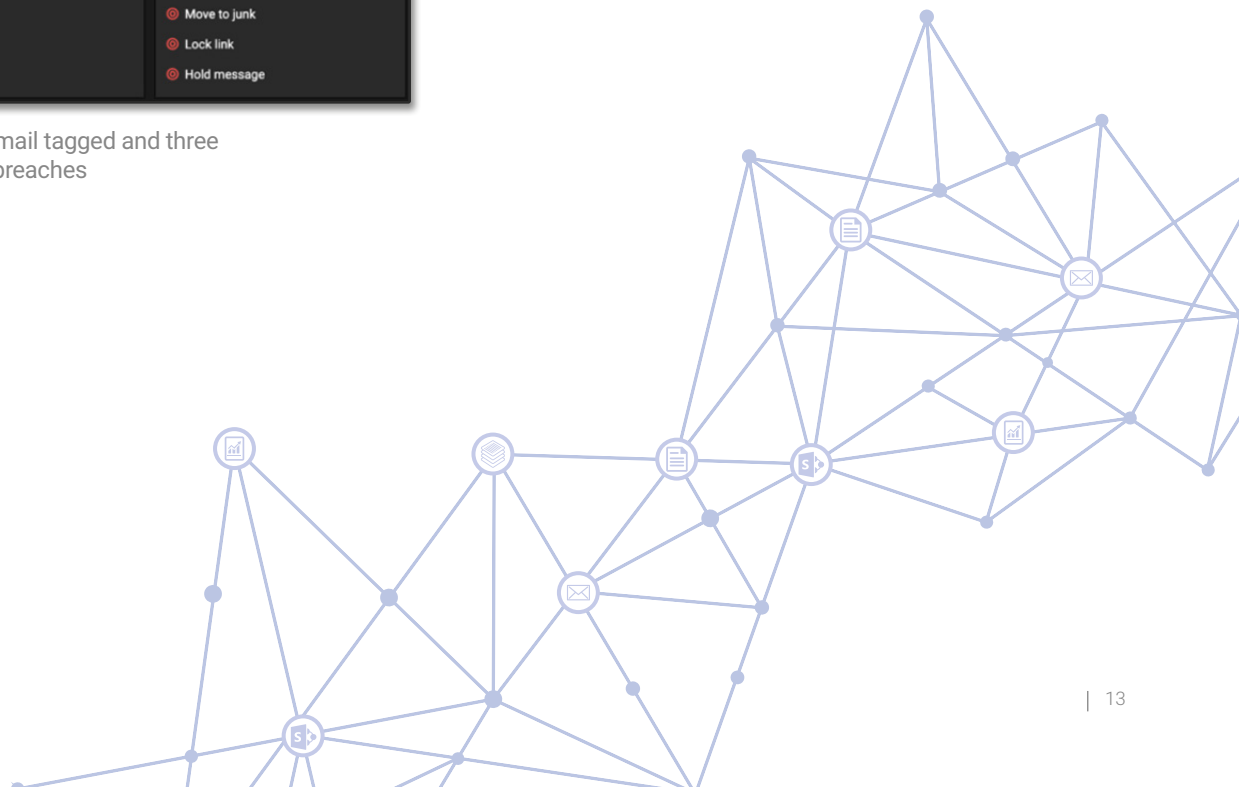


Figure 20: The malicious link identified

6. The link itself was hidden behind the display text 'Retrieve Message' and went to a OneDrive page. The use of file storage domains for hosting malicious content is difficult to catch using a traditional approach, as it is impossible to blacklist services such as SharePoint, and deciding whether a link such as this one is malicious or benign requires an understanding of the email in the context of the wider organization.



Social Engineering & Solicitation

“
We have Antigena Email deployed as well as legacy security tools. We were shocked by the things the traditional tools didn't catch that Antigena Email did.”

– CTO, Bunim Murray Productions

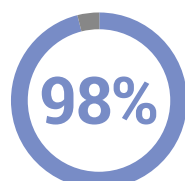
Social engineering and solicitation attacks typically involve a sophisticated attempt at impersonation, where disguised attackers urgently prompt a recipient to reply, take communications offline, or perform an offline transaction. Their goals range from wire fraud to corporate espionage and even IP theft. While organizations should of course invest in security training and educate their employees to look out for warning signs, no amount of guidance can guarantee complete immunity from these increasingly sophisticated attacks.

While traditional phishing campaigns generally include a malicious payload hidden behind a link or attachment, social engineering attempts often involve sending 'clean emails' that contain only text. These attacks easily bypass legacy security tools that rely on correlating links and attachments with blacklists and signatures. Moreover, this vector of attack generally involves registering new 'look-alike' domains, which not only trick the recipient but also bypass traditional defenses.

Antigena Email has a unified understanding of 'normal' across all email and network traffic that evolves with the business, allowing it to detect subtle cases of solicitation. Clean emails that bypass traditional defenses can be identified in seconds given a vast range of metrics, including suspicious similarities to known users, abnormal associations among internal recipients, and even anomalies in email content and subject matter.

More often than not, social engineering attacks aim to immediately take the conversation offline, which means that slow and reactive security measures tend to only intervene after the damage is done. Its powerful understanding of every user, device, and relationship in the organization allows Antigena Email to respond proactively and with high confidence the first time around, intervening at this crucial early stage.

Antigena is also unique in its ability to intelligently tailor responses to specific threat types. It understands that the 'dangerous' element in a solicitation attack will often be the email content itself, and the system will therefore prevent delivery before the intended recipient even has a chance to comply with the attacker's urgent request.



98% of attacks in user inboxes contained no malware

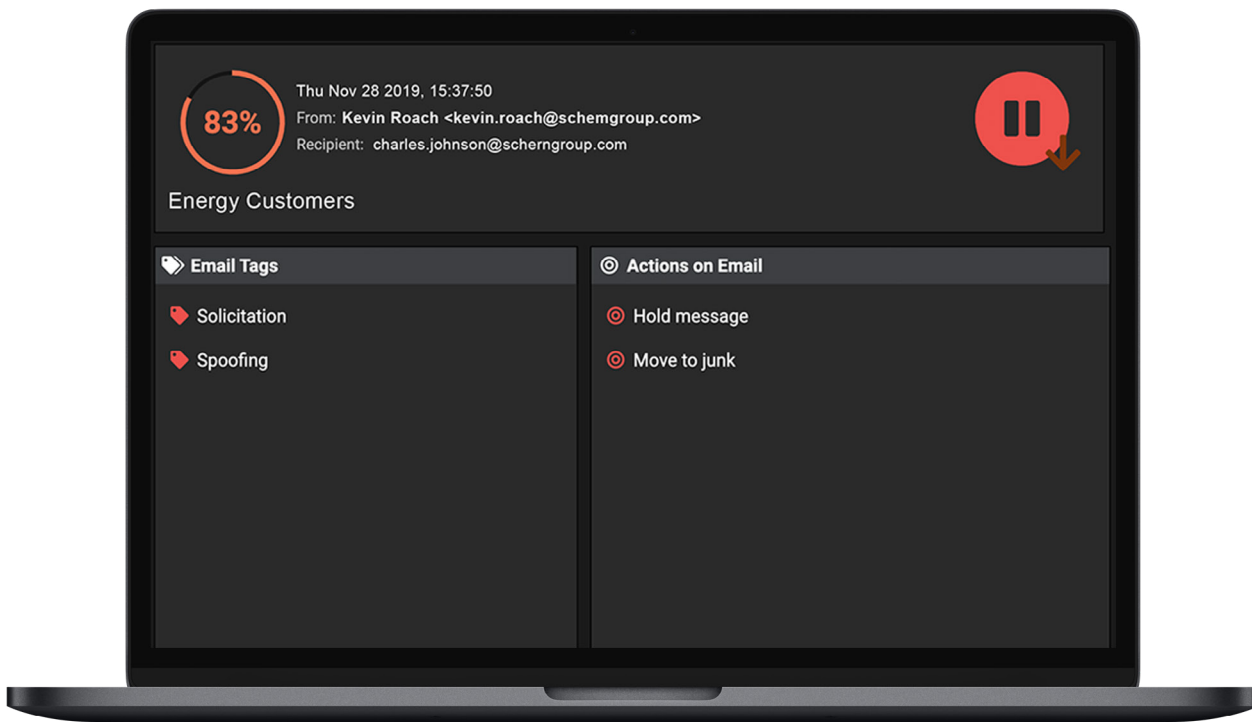
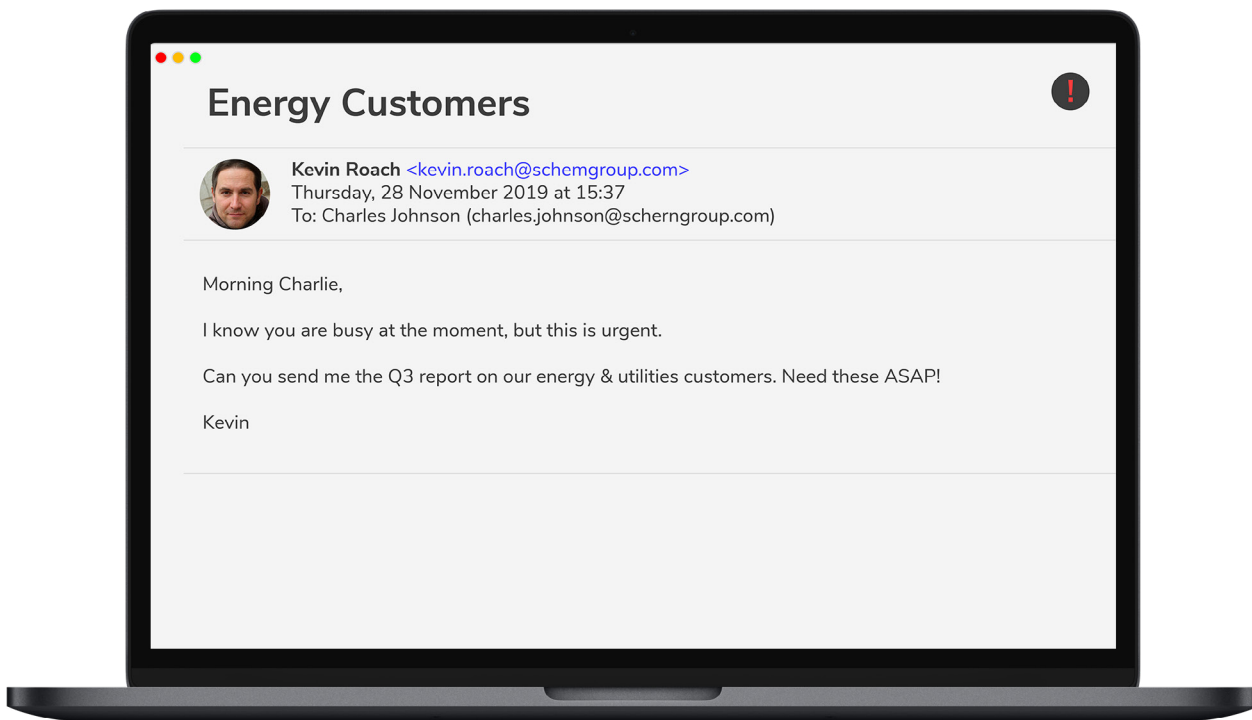


Figure 21: An attacker posing as an executive, seeking to leverage sensitive documents. Note the spoofed email address.

Impersonation Attack

Antigena Email detected a targeted attack against 30 employees of a multinational technology company. Extensive research was clearly carried out, as for each targeted user, the attacker carefully impersonated the C-level executive with whom they were most likely to communicate. Antigena Email identified the social engineering attack, and as a result held back each email from the intended recipients.

1. The subject line of each email included the first name of the targeted employee, and came from a seemingly unrelated Gmail address. Despite the lack of a malicious payload (such as links or attachments), Antigena Email was still able to identify the emails as malicious.

2. Darktrace not only identified the impersonation attempts by recognizing the look-alike domain name, but also that the emails had breached the 'No Association' model, indicating that across its entire understanding of the company's email and network environment, it had seen no evidence of a relationship between this sender and the organization.

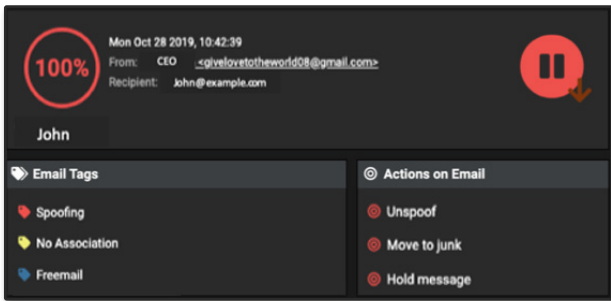


Figure 22: One of 30 emails, with a 100% anomaly score

3. Correlating multiple weak indicators, Antigena recognized these emails as components of one coordinated attack, causing it to hold them in a buffer for the organization's security team to review.

4. Antigena Email not only identified the three C-level executives who were being impersonated, but also recognized that the attacker was using a spoof of their CEO's legitimate external personal address as well.

Header From Personal	Count
CEO	18
CTO	11
CFO	1

Figure 23: Three C-Level executives identified

5. In addition, the exposure score of the impersonated users was high, indicating they were high-profile targets, and hence breaching the 'Whale Spoof' model. Understanding that key internal users had been targeted allowed Darktrace's AI to prioritize this attack, initiating a proportionate response in real time.

CEO Payroll Request

At an electricity distributor, Darktrace's AI detected a convincing spoof attempt discovered in an Office 365 email account. Allegedly from the company's CEO, the email was sent to a member of the payroll department requesting that the employee update the CEO's direct deposit information.

Since the email successfully mimicked the CEO's typical writing style, it could have easily succeeded if Darktrace's AI hadn't been analyzing the firm's mail flow in connection with the rest of the business.

1. By learning the normal 'pattern of life' of the employee, the CEO, and the wider organization across cloud and network traffic, Darktrace was able to immediately flag a number of subtle anomalies in the email, including the forged sender address.

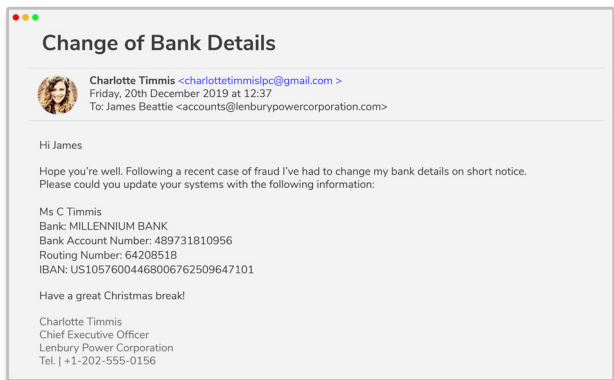


Figure 24: Screenshot of email impersonating the CEO

2. Among other weak indicators, Darktrace's AI automatically calculated the anomalous proximity of the domain to those of internal employees and trusted contacts.

3. The AI responded immediately, locking the email's links and clearly marking it as a spoof before it could reach the payroll department. Darktrace's rich understanding of cloud and network traffic allowed it to neutralize a high-severity threat that signature-based tools would have missed.

'Finance VP' Spoofing Attack

This incident involved the impersonation of a Finance VP at a well-known financial institution. The threat actors sent 11 similar emails to the organization, but Antigena Email took action to hold all of them given its multi-dimensional understanding of 'normal' across network, cloud, and email traffic. Analyzing the unrelated, clearly anomalous email address in connection with the content of the emails, Darktrace recognized this spoofing attempt, while the company's legacy gateway let all 11 emails through.

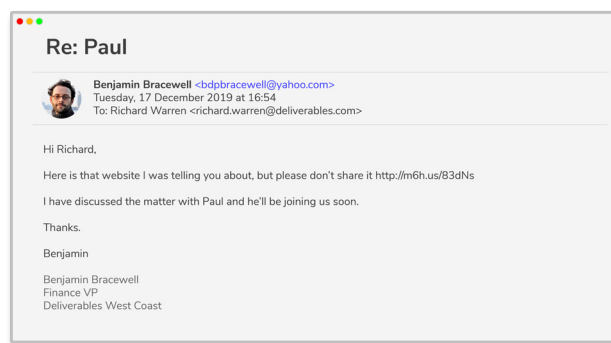


Figure 25: Screenshot of email sharing suspicious link

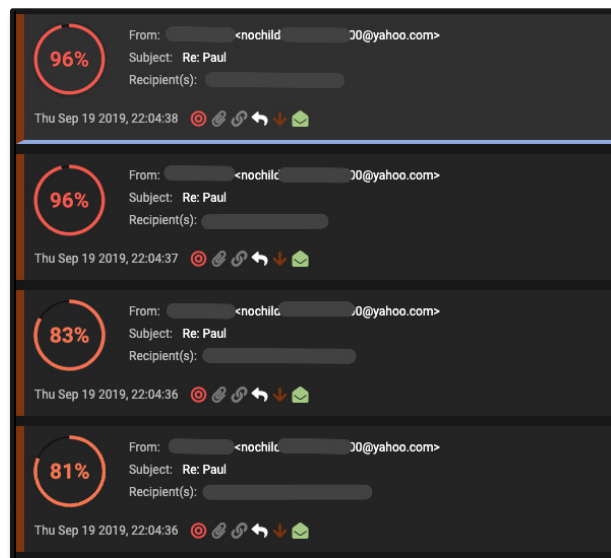




Figure 26: Four of the 11 emails, showing the high anomaly score and associated Antigena Email action



Compromised Employee Credentials

Credential compromise has increased 280% between 2016 and 2019



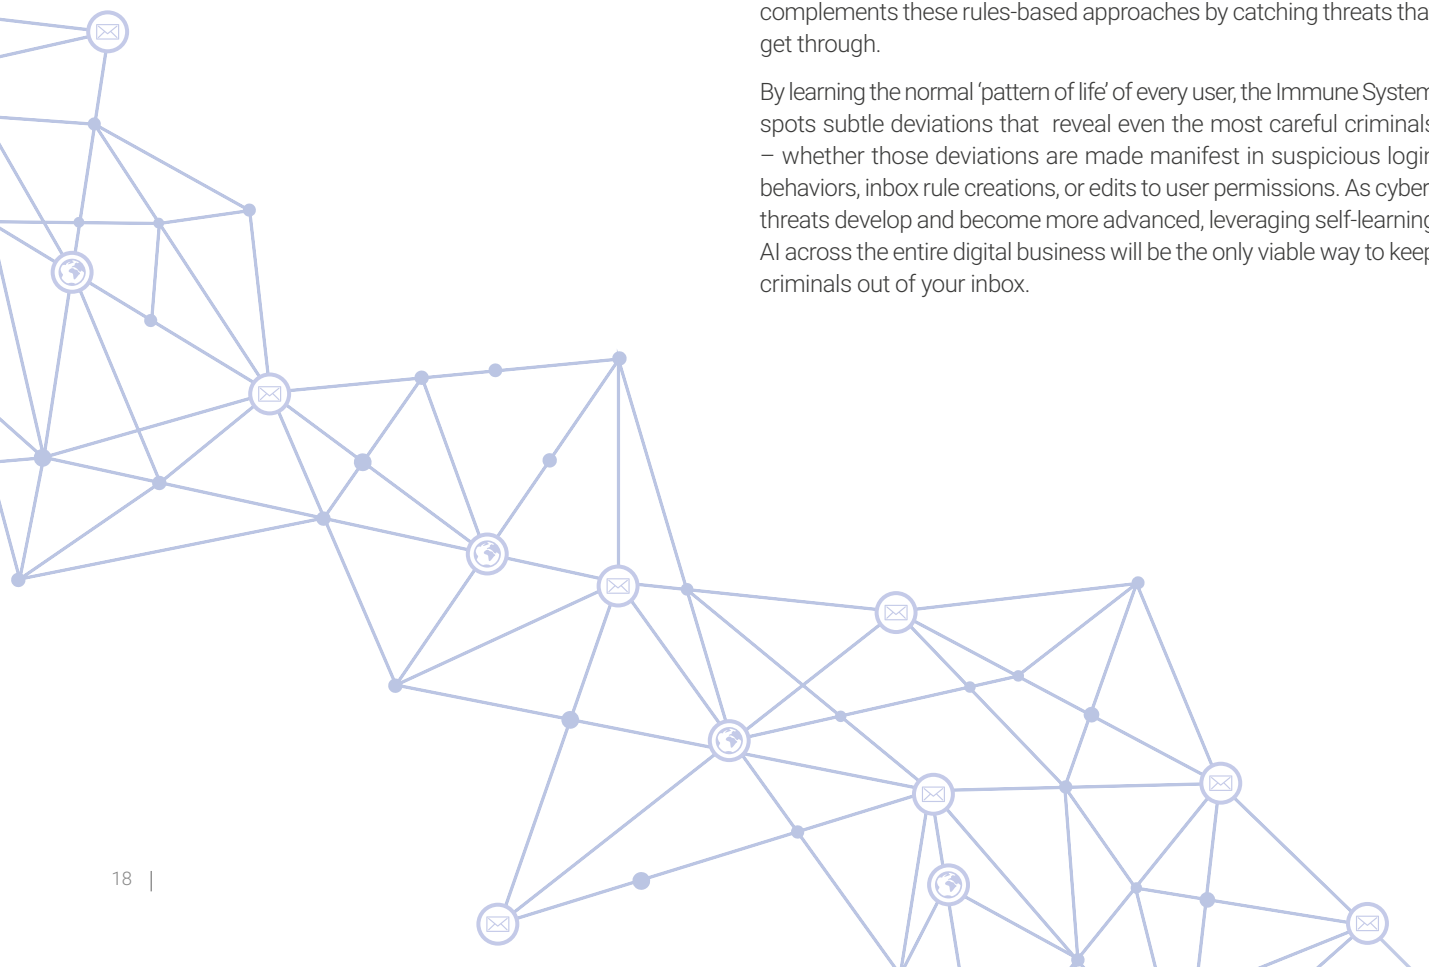
Business leaders rarely consider how valuable a corporate inbox can be until it falls into the wrong hands. Yet once inside, threat actors enjoy a wide range of attack options and pivot points from which to choose. The ease with which attackers can gain access – whether through phishing campaigns, brute force attempts, or exchanges on the Dark Web – should be cause for alarm.

In many cases, attackers will pillage your inbox for the valuable data it contains. Personal information from private chats to billing details can be leveraged for fraud or blackmail, while old email threads may contain highly confidential company information. Customer lists, pricing documents, and even roadmap and IP details are often just a few search terms away from being discovered.

In other cases, criminals will use the account as a launching point for the next stages of an attack. They may sit quietly in the background to gather intelligence about high-value executives or partners, reviewing documents, reading conversations, and learning how to blend in when they inevitably strike. As with supply chain account takeovers, the ability to read an ongoing email thread and follow up with a plausible reply is often the most effective way to achieve an attack mission without triggering suspicion.

While the possibilities for attackers are nearly endless, the options for defenders are limited. Corporate account takeovers are typically monitored for by simple and static defenses, including ‘impossible travel’ rules that rarely catch attackers who know how to hide. Thanks to its enterprise-wide view, however, Darktrace’s Immune System Platform complements these rules-based approaches by catching threats that get through.

By learning the normal ‘pattern of life’ of every user, the Immune System spots subtle deviations that reveal even the most careful criminals – whether those deviations are made manifest in suspicious login behaviors, inbox rule creations, or edits to user permissions. As cyber-threats develop and become more advanced, leveraging self-learning AI across the entire digital business will be the only viable way to keep criminals out of your inbox.



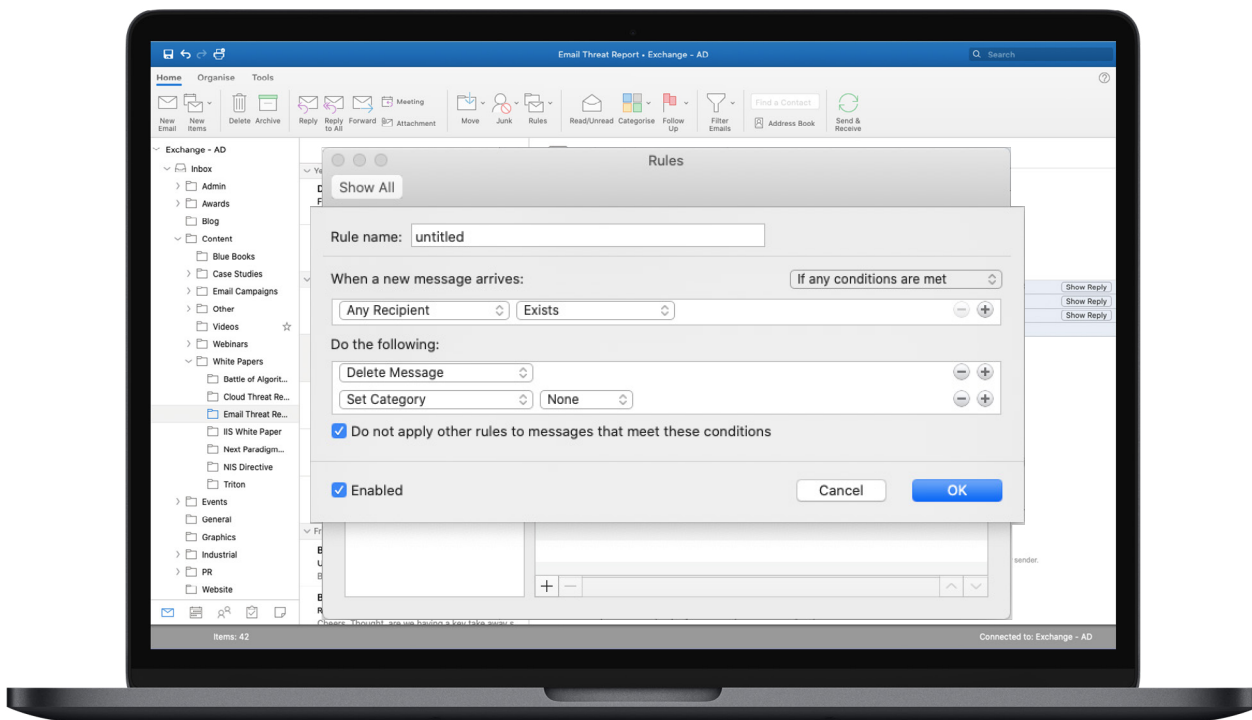


Figure 27: An email processing rule being set up on a compromised account, and the Threat Visualizer displaying the geographical login locations.

Unusual Login at Panamanian Bank

One Office 365 account was used in a brute force attack against a well-known bank in Panama, with logins originating from a country that deviated from the normal ‘patterns of life’ of the company’s operations.

Darktrace identified 885 logins over a period of 7 days. While the majority of authentications originated from IP addresses in Panama, 15% of the authentications originated from an IP address that was 100% rare and located in India. A further analysis revealed that this external endpoint was included in multiple spam blacklists, and that it had recently been associated with abusive behavior online – possibly unauthorized Internet scanning or hacking.

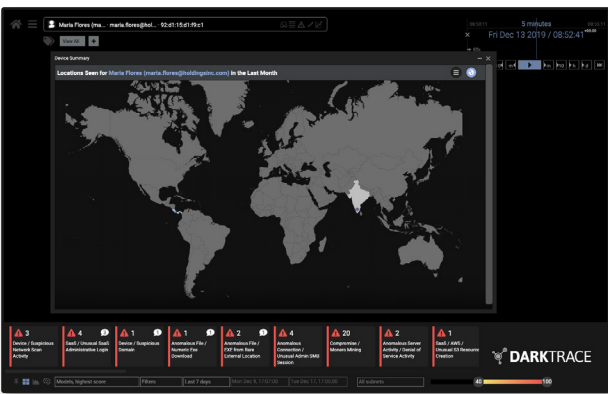


Figure 28: The user interface showing login locations

Darktrace then witnessed what appeared to be an abuse of the password reset function, as the user in India was observed changing account privileges in a highly unusual manner. What marked the activity as particularly suspicious was that after the password reset, failed log-in attempts from an IP normally associated with the organization were observed, suggesting the legitimate user was locked out.

03/12 20:45:39	SaaS:Admin	Regular	UpdateUser
03/12 20:45:39	SaaS:Admin	Regular	ChangeUserLicense
03/12 20:26:43	SaaS:Login	Regular	UserLoggedIn
03/12 20:26:43	SaaS:FailedLogin	Regular	UserLoginFailed
03/12 20:26:36	SaaS:FailedLogin	Regular	UserLoginFailed
03/12 18:31:31	SaaS:Login	Regular	UserLoggedIn
03/12 17:57:46	SaaS:Admin	Regular	ChangeUserLicense
03/12 17:57:46	SaaS:Admin	Regular	UpdateUser
03/12 17:06:57	SaaS:Admin	Regular	UpdateUser

Figure 29: The activity associated with the SaaS account, highlighting the changed credentials

Attempted Access from Rural Japan

At a financial services corporation based in Europe, an Office 365 credential was observed logging in from an unusual IP address linked to a location in rural Japan.

Although access from remote locations is possible if a user travels or uses a proxy service, this could also be a strong indicator of compromised credentials and malicious access by an unauthorized user. Given that the access point was substantially different from the usual accessing IPs, Darktrace flagged this as anomalous and immediately suggested further investigation.

The security team was able to remotely lock the Office 365 account and reset the credentials, preventing the malicious actor from further activity. Had this activity gone unnoticed, the threat actor could have used their access privileges to deploy malware in the organization or solicit a fraudulent payment.

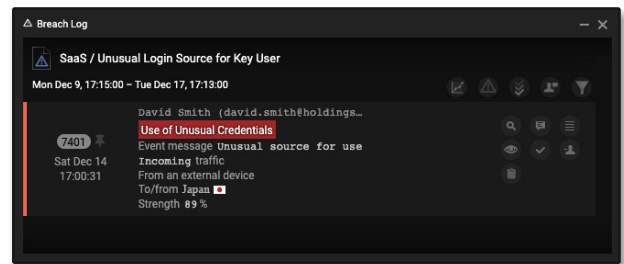


Figure 30: The login from Japan breached several models



Office 365 Account Compromised and Sabotaged

In one international non-profit, Darktrace detected an account takeover in Office 365 that bypassed Azure's AD static 'impossible travel' rule. While the organization had offices in every corner of the globe, Darktrace's self-learning AI identified a login from an IP address that was historically unusual for that user and her peer group and immediately alerted the security team.

Darktrace then alerted to the fact that a new email processing rule, which deletes inbound and outbound emails, had been set up on the account. This indicated a clear sign of compromise and the security team was able to lock the account before the attacker could do damage.

With this new email processing rule in place, the attacker could have initiated numerous exchanges with other employees in the business, without the legitimate user ever knowing. This is a common strategy used by cyber-criminals seeking to gain persistent access and leverage multiple footholds within an organization, potentially in preparation for a large-scale attack.

Analyzing the rare IP address in conjunction with the out-of-character behavior of the apparent user, Darktrace confidently identified this as a case of account takeover, preventing serious damage to the business.

Automated Brute Force Attack

Darktrace detected several failed login events on an Office 365 account using the same credential, every day over the course of a week. Each batch of login attempts were performed at precisely 6.04pm on six days. The consistency in both the time of day and the number of login attempts was indicative of an automated brute force attack, which is programmed to discontinue after a certain number of failed attempts in order to avoid lockouts.

Darktrace considered this pattern of failed attempts highly anomalous and so alerted the security team. Were it not for Darktrace correlating multiple weak indicators and fleshing out the subtle signs of emerging threat, this automated attack could have continued for weeks or months, making educated guesses at the users' password based on other information it had already gathered.

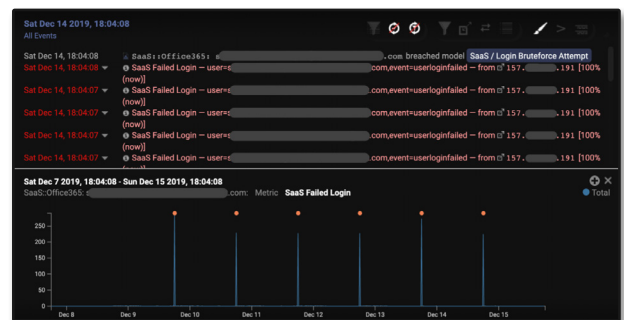



Figure 31: A graph illustrating the repeated login attempts





About Darktrace

Darktrace is the world's leading cyber AI company and the creator of Autonomous Response technology. Its self-learning AI is modeled on the human immune system and used by over 3,000 organizations to protect against threats to the cloud, email, IoT, networks and industrial systems.

The company has over 1,000 employees and headquarters in San Francisco and Cambridge, UK. Every 3 seconds, Darktrace AI fights back against a cyber-threat, preventing it from causing damage.

Contact Us


North America: +1 (415) 229 9100

Europe: +44 (0) 1223 394 100

Asia-Pacific: +65 6804 5010

Latin America: +55 11 97242 2011

info@darktrace.com | darktrace.com

 [@darktrace](https://twitter.com/darktrace)