

Email Security Threat Report 2020

4 key trends from spear phishing to credentials theft



The Email Threat Landscape

Contents


The Email Threat Landscape	1
Antigena Email: The Self-Defending Inbox	2
Spear Phishing & Payload Delivery	3
WeTransfer Attack	5
Malware Hidden in Fake Invoices	6
Municipality Address Book Compromised	6
Supply Chain Account Takeover	7
Consecutive Supply Chain Attacks	9
Malicious File Hidden in OneDrive Page	12
Social Engineering & Solicitation	13
Impersonation Attack	15
CEO Payroll Request	16
'Finance VP' aiming to initiate trusted internal relationship	16
Compromised Employee Credentials	17
Compromise across Microsoft 365 and Teams	19
'Change of bank details' sent from Accounts Department	20
Account Takeover at Panamanian Bank	21
Unusual External Source	21
Office 365 Account Compromised and Sabotaged	22

Email and collaboration platforms are the connective tissue of most businesses and organizations, where information is shared, plans are hatched, and alliances formed. Yet as a human-driven medium, email often represents the 'weakest link' in an organization's security strategy. Indeed, 94% of cyber-threats originate in the email environment.

While traditional gateway tools seek to filter out malicious emails on entry, their reliance on lists of 'known-bad' IPs, domains, and file hashes to determine an email's threat level is extremely limiting. A rule-based approach can often identify known spam and other low-hanging fruit, but it fails to keep pace with attacker innovations.

Spear phishing, impersonation attacks, and account takeovers, in particular, remain fruitful ways that cyber-criminals can infiltrate an organization. Increasingly targeted email attacks of this kind, which overcome the limitations of traditional defenses, are a significant challenge for security teams today.

As Peter Firstbrook, VP Analyst at Gartner, puts it: "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."



Antigena Email: The Self-Defending Inbox

Antigena Email is the world's first Cyber AI solution for the inbox. By learning the normal 'pattern of life' for every user and correspondent, the technology builds an evolving understanding of the 'human' within email communications.

While traditional defenses ask whether elements of an email have been observed in historical attacks, Antigena Email is the only solution that can reliably ask whether it would be unusual for a recipient to interact with a given email, in the context of their normal 'pattern of life', as well as that of their peers and the wider organization.

This contextual knowledge enables the AI to make highly accurate decisions and neutralize the full range of email attacks, from 'clean' spoofing emails that seek to wire a fraudulent payment, to sophisticated spear phishing attempts.

Inspired by the human immune system, Antigena Email uses Darktrace's core artificial intelligence to learn a sense of 'self' for every internal and external user, analyzing both inbound and outbound communications together with lateral, internal-to-internal communications. By treating recipients as dynamic individuals and peers, Antigena Email uniquely spots subtle deviations from 'the norm' that reveal seemingly benign emails to be unmistakably malicious.

In the case studies that follow, Darktrace's evolving sense of the 'self' of email users and their peers has enabled it to detect and stop the email threat that – in every case – traditional security tools let through. These email attacks fall into one of four highly sophisticated attack categories that routinely bypass your organization's 'protective skin':

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




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 the ‘patterns of life’ for the individual and their wider peer group are taken into account. This basic truth is precisely why a dynamic analysis that takes into account hundreds of metrics based on user and group behaviors is so critical.



Powered by cyber AI, Antigena Email can analyze links, attachments, domains, content, and other elements of an email alongside ‘patterns of life’ across the organization, correlating a rich constellation of data points that reveal seemingly benign emails to be unmistakably malicious.

Unlike any other solution, Antigena Email understands the human behind email interactions, identifying anomalous sending patterns, whether the location of a link in an email is strange, the topics of discussion and content are unusual, or even whether patterns in the URL pathway are suspicious.

This unique approach means that Darktrace can take highly proportionate and targeted actions to neutralize phishing attacks within seconds.

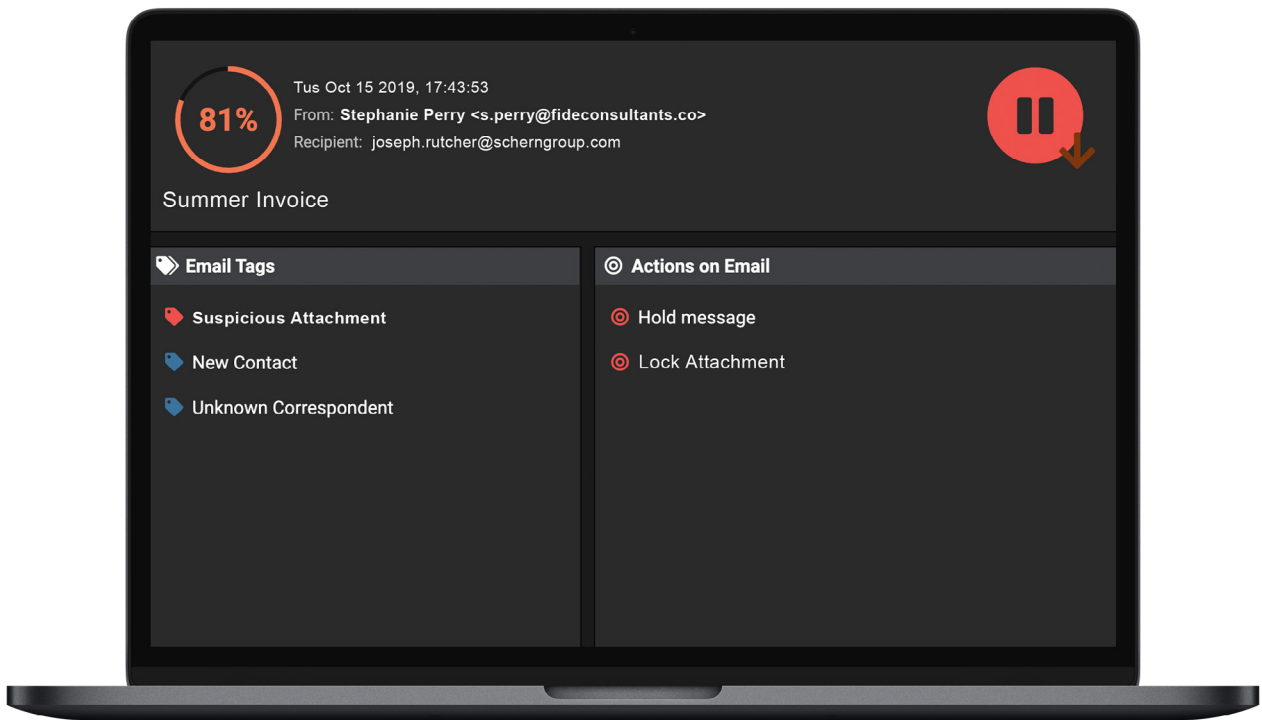
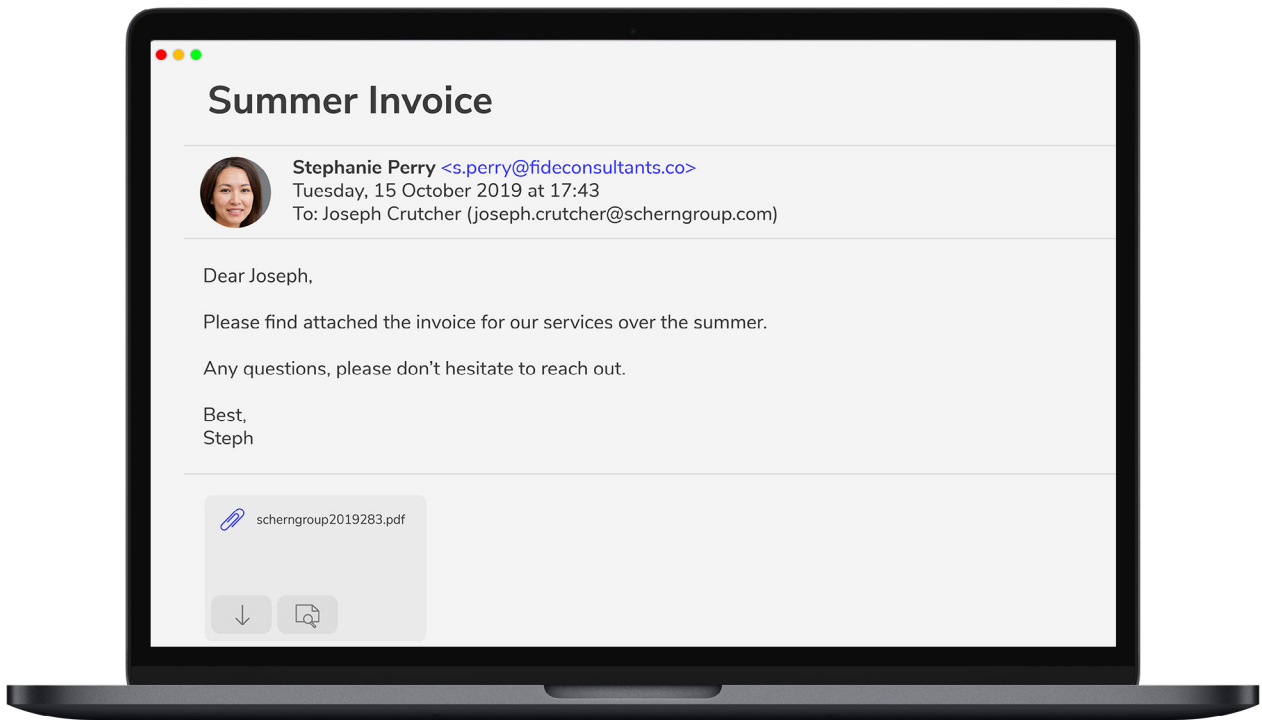


Figure 1: An email coaxing an employee to click on an attachment containing a malicious payload, and the corresponding view within Darktrace's user interface, showing the anomaly tags and actions taken

REAL-WORLD CASE STUDY

WeTransfer Attack

Antigena Email neutralized a sophisticated email attack targeting five high-profile users at an academic institution. The emails were well-written and plausible, attempting to coax the recipients into clicking on a malicious link.

These emails were assigned a 100% anomaly score and Antigena Email 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.

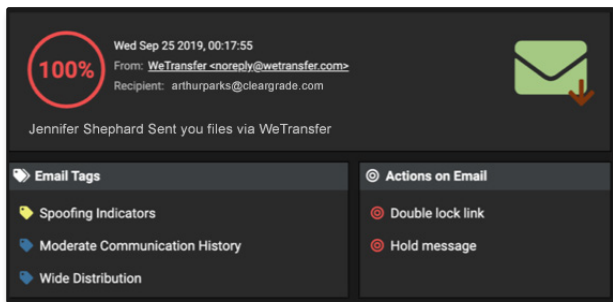


Figure 2: The user interface showing the model breaches and actions

1. From the connection data, there were no clear signs in the headers that this email was not in fact coming from WeTransfer, so the impersonation attempt would have likely fooled the recipient. The 'Width' and 'Depth' indicate that this email address has communicated with many people in the organization, across multiple days.




Figure 3: The connection data of the relevant emails

2. However, Antigena Email was able to detect a range of subtle anomalies given its understanding of 'normal' for the user and wider environment.

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 Antigena Email 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 the technology to identify it as the malicious payload in the email.




Figure 4: 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'.

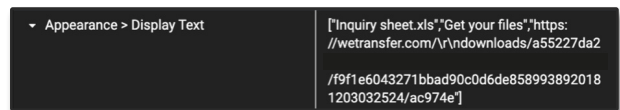


Figure 5: Antigena was able to determine where the link appeared within the email

This incident demonstrates how a self-learning approach allows Antigena Email to identify advanced phishing attacks that leverage the familiarity of a trusted website in order to deliver a harmful link and gain multiple footholds in the organization.



REAL-WORLD CASE STUDY

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 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.

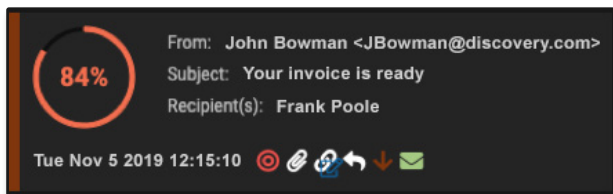


Figure 6: Header of the malicious emails, showing suggested action

REAL-WORLD CASE STUDY

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, Antigena Email immediately recognized that the domain was unusual for the organization. The system also recognized that the way the links were hidden behind each button was highly suspicious. It raised a high-confidence alert, and responded by locking each link. While the municipality's legacy email defenses finally woke up to the attack at the letter 'R', Antigena Email neutralized it at 'A', as soon as the first email came through.






Figure 7: 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 of a trusted contact in your supply chain, threat actors can easily gain the trust of a user and coax them into clicking a malicious link, or even wiring money out of the business. Legacy email defenses assume trust, which means that sophisticated account takeovers often go unnoticed.

Compromised accounts have been responsible for several high-profile attacks on large organizations in recent years. Cyber-criminals are increasingly leveraging supply chains – suppliers, partners, contractors – to infiltrate their ultimate target 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 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.

Analyzing patterns of communication with the full context of all inbound, outbound and lateral mail flow, Antigena Email uses a range of metrics to identify cases of account takeover, something that is impossible to detect without a detailed understanding of ‘normal’ behavior for the entire organization.

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 evolving understanding of ‘normal’ to determine the likelihood that an email from a trusted supplier is 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.



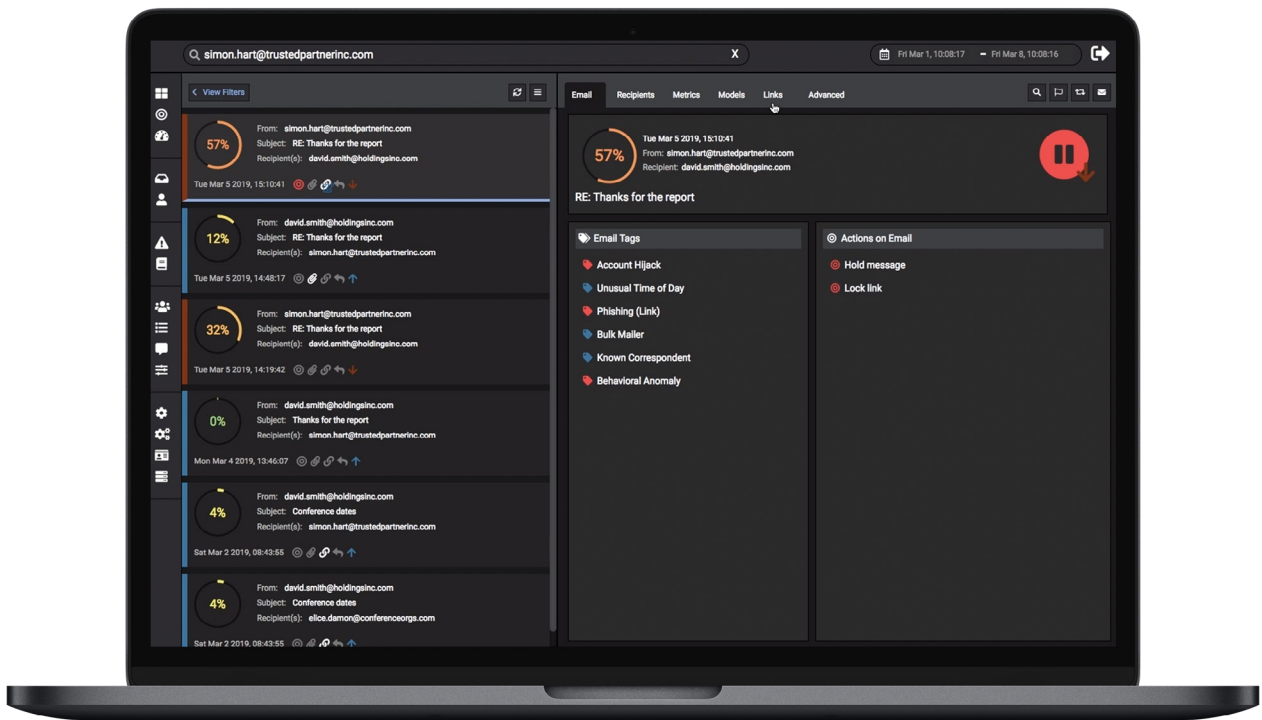
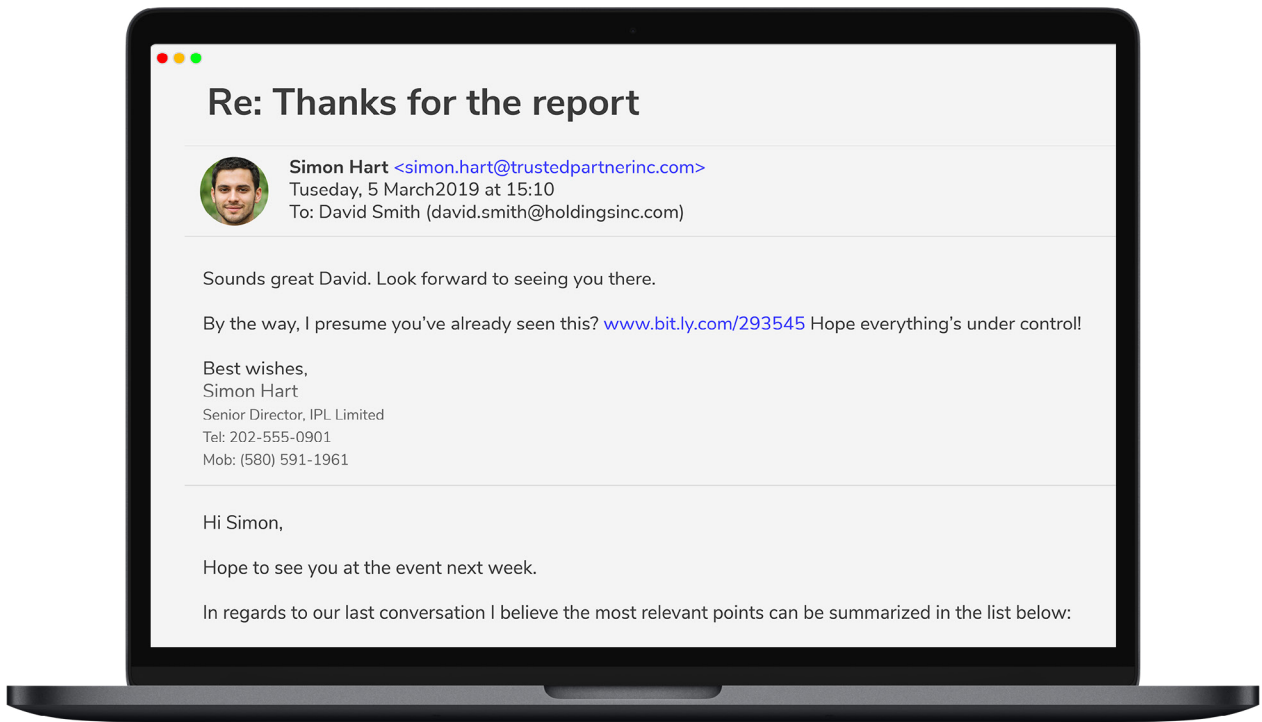


Figure 8: A plausible reply sent from a trusted supplier's compromised account following a thread of email correspondence. The link contained a malicious payload

REAL-WORLD CASE STUDY

Consecutive Supply Chain Attacks

A Darktrace customer 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.

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 was a UK-based environmental consultancy firm.

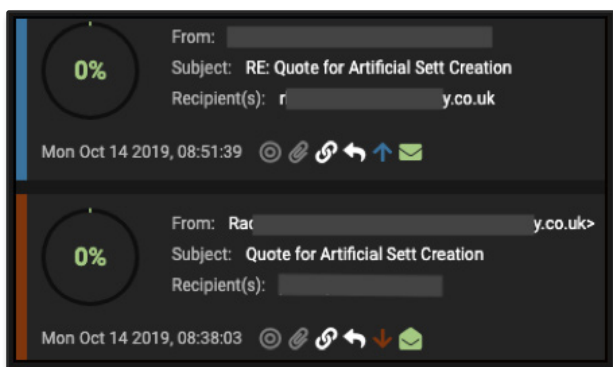


Figure 9: Earlier ‘normal’ correspondence with the sender – with a 0% anomaly score

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.

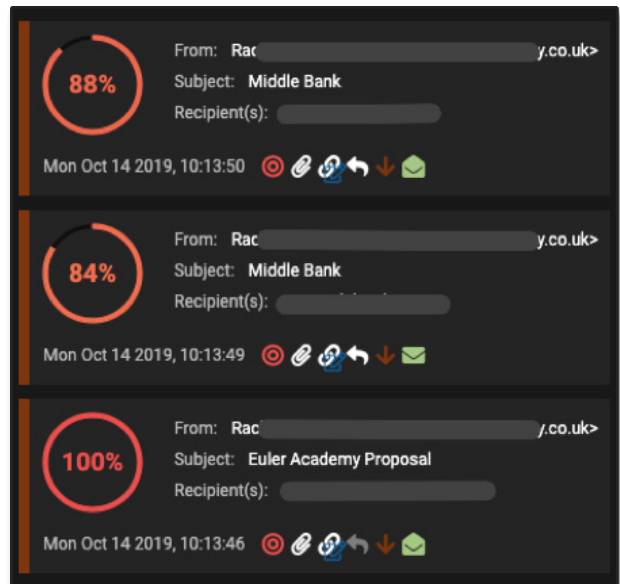
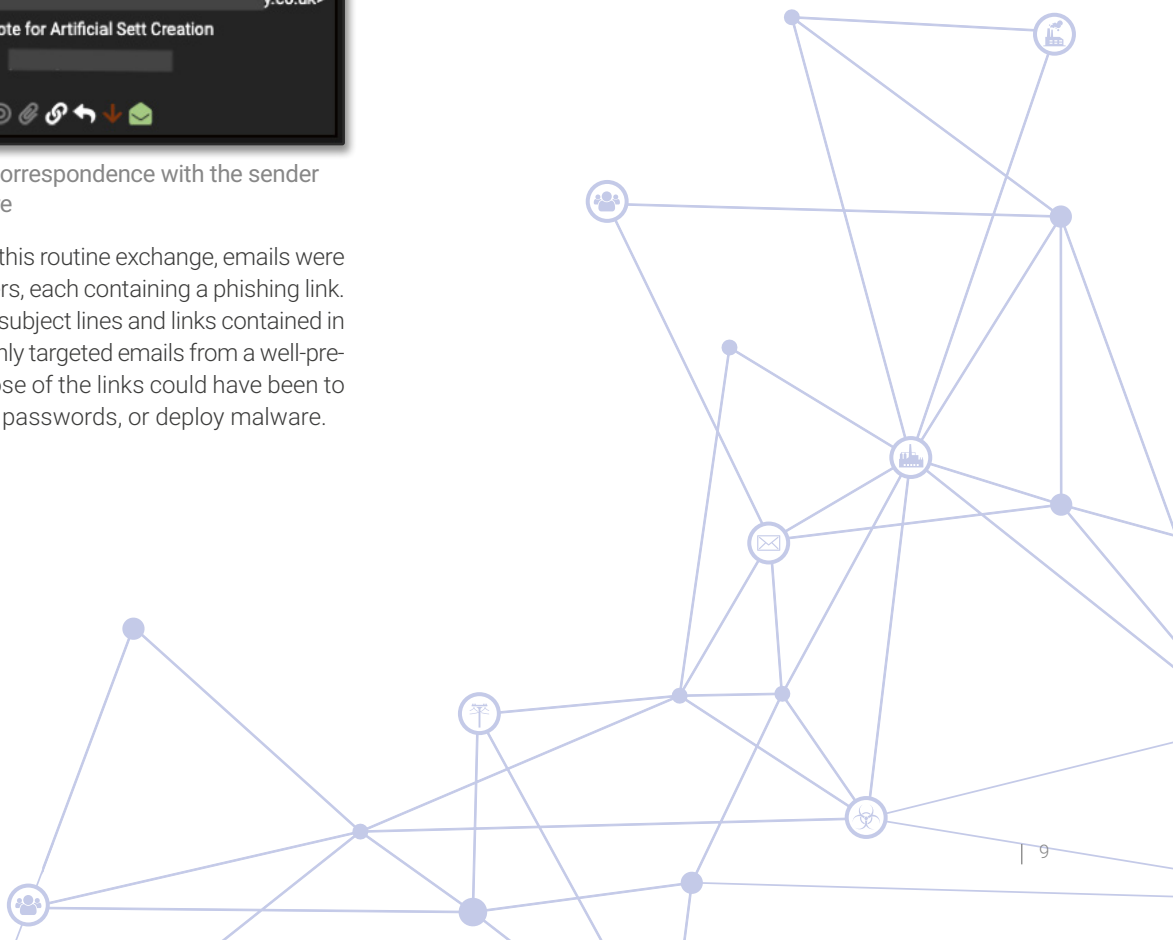


Figure 10: Emails sent later the same day containing malicious attachments



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.

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.

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

Figure 11: 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 12: Antigena Email rapidly detected that this group of recipients was not closely related




Figure 13: 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 moved these emails to Junk and locked their links.

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.

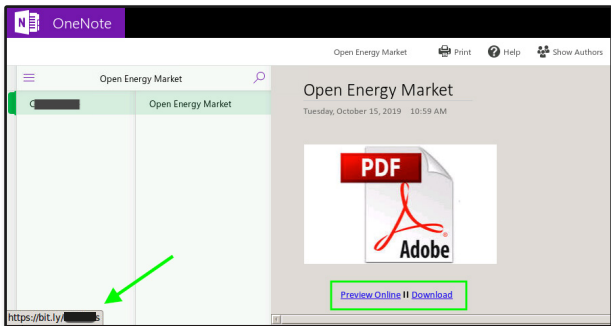


Figure 14: Screenshot exposing a hidden link




Figure 15: 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.

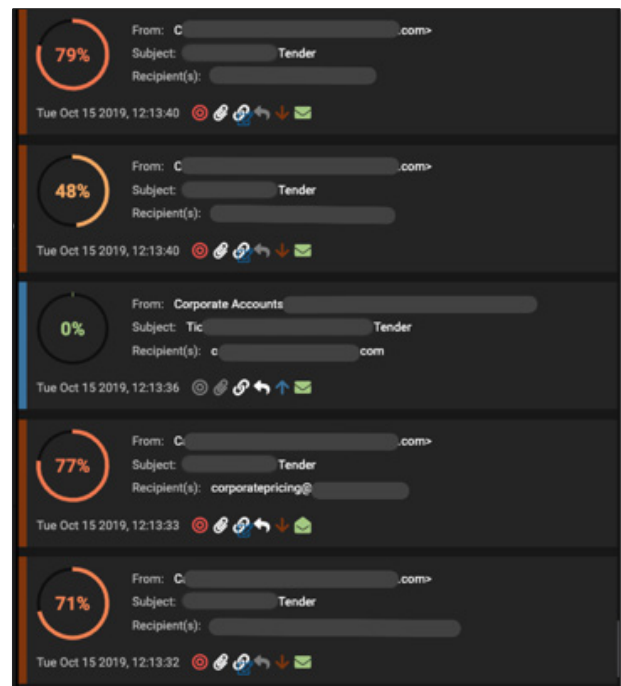


Figure 16: Emails from Incident 2 as they appear in the Antigena Email console, including those that were sent outbound in response. It reveals the ‘corporate accounts’ user acknowledged the email by opening a ticket

REAL-WORLD CASE STUDY

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.

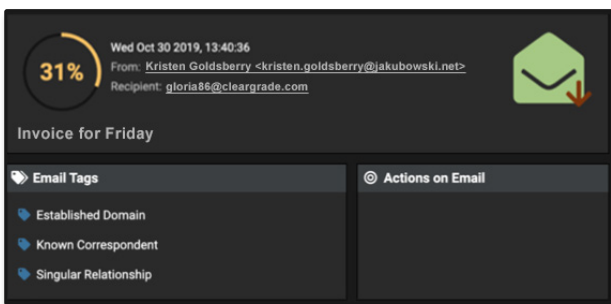


Figure 17: An example of a previous communication

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

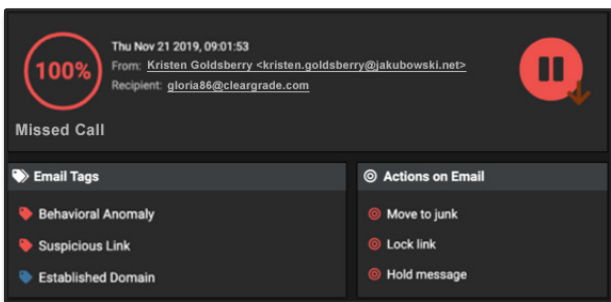


Figure 18: A later email tagged and three associated model breaches

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 traffic, and hence was locked by Antigena Email.





Figure 19: 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 are advised to invest in employee security training, no amount of guidance can guarantee immunity from 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. This vector of attack generally involves registering new 'lookalike' domains, which not only trick the recipient but also bypass traditional defenses.

Antigena Email has a unique understanding of the human behind the email address that evolves over time, 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, and slow, reactive security measures only intervene after the damage is done. Its powerful understanding of every user, device, and their relationships within the organization allows Antigena Email to respond proactively and with high confidence the first time around, intervening at an early stage.

Antigena Email is also able to tailor responses to specific threat types. It understands that the dangerous element in a solicitation attack is often the email content itself, and the system can prevent delivery, thus preemptively protecting the intended recipient.

98% of attacks in user inboxes contained no malware

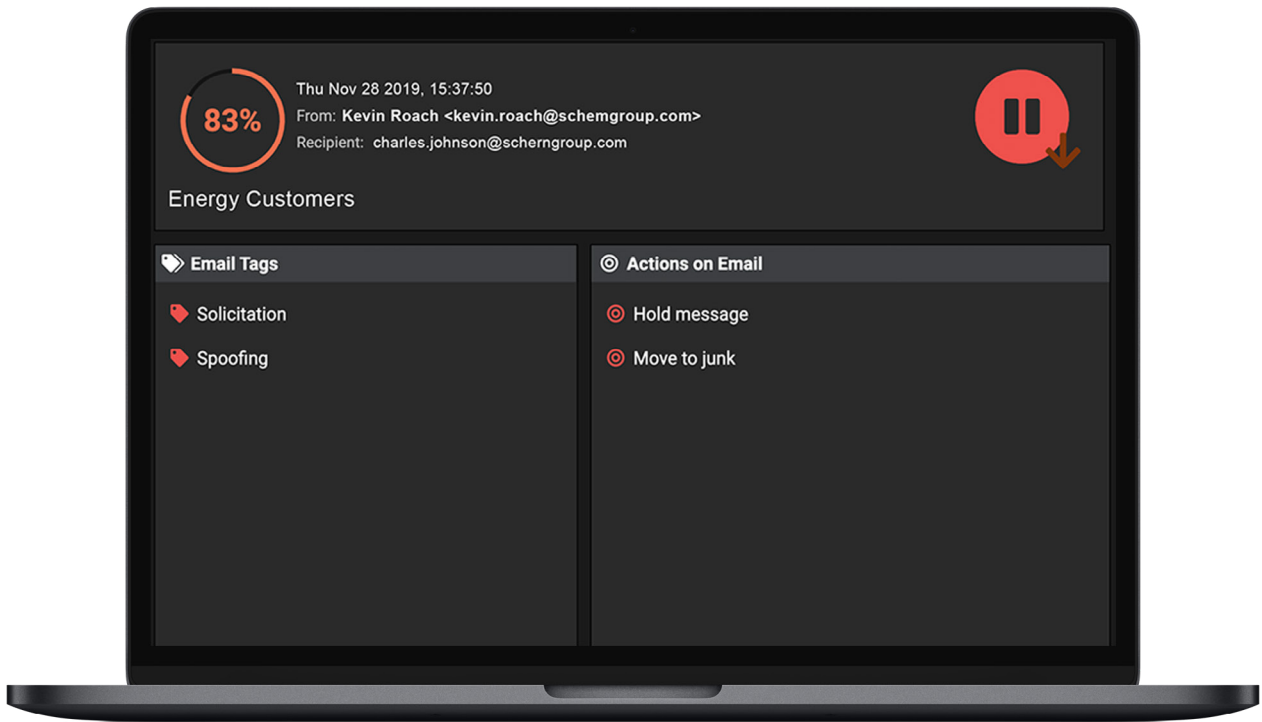
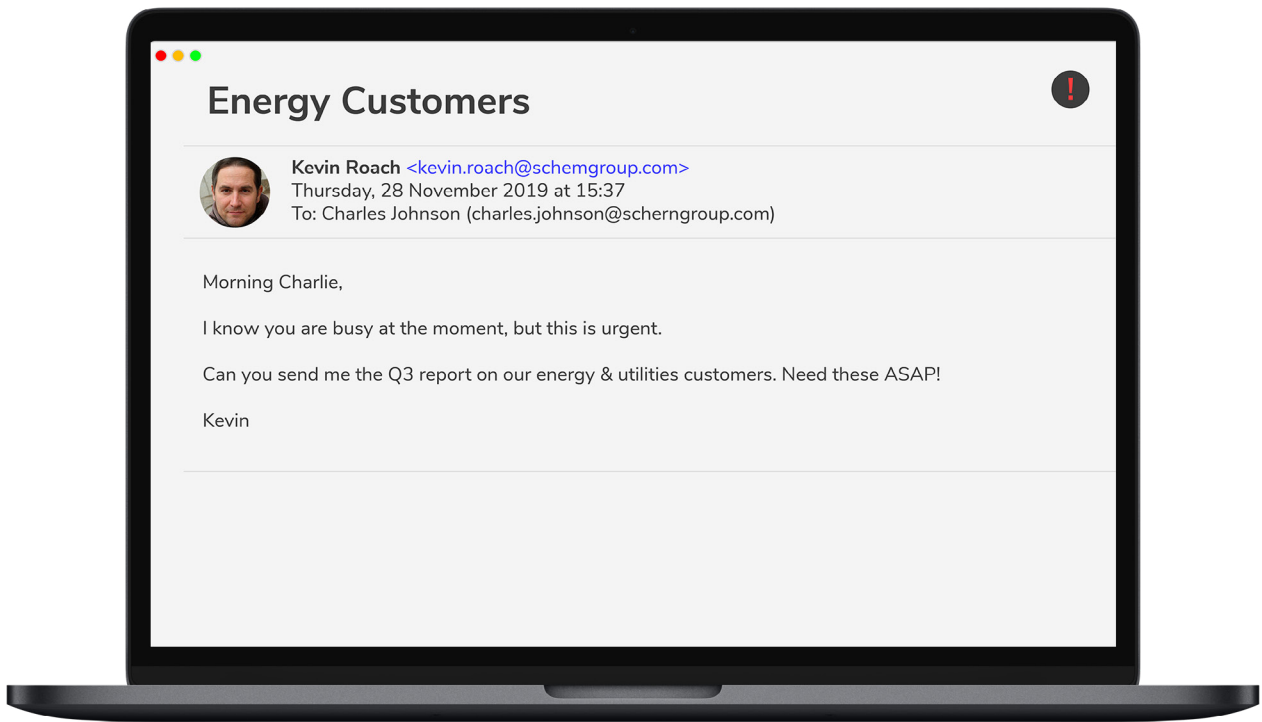


Figure 20: An attacker posing as an executive, seeking to leverage sensitive documents. Note the spoofed email address

REAL-WORLD CASE STUDY

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 user that was targeted, the attacker had impersonated the C-level executive with whom they were most likely to communicate. Antigena Email identified the social engineering attack, and 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 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, it had seen no evidence of a relationship between this sender and the organization.

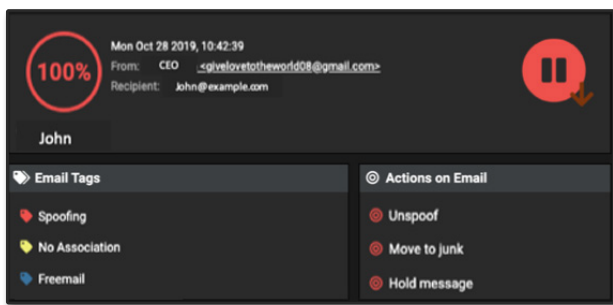


Figure 21: One of 30 emails, with a 100% anomaly score

3. Correlating multiple weak indicators, Antigena Email 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 22: Three C-Level executives identified

5. In addition, the exposure score of the impersonated users was high, indicating that 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.

These emails invited the recipients to take the conversation off email altogether, asking the recipients to send some highly sensitive information out of the organization via a SaaS application. Email tools focused solely on detecting incidents of unexpected data loss therefore would have failed to spot this attack.



REAL-WORLD CASE STUDY CEO Payroll Request

At an electricity distributor, Antigena Email 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.

The attacker had accurately mimicked the CEO's writing style, and the email would most likely have succeeded in its objectives, if Darktrace's AI hadn't identified its subtly anomalous characteristics.

1. By learning the normal 'pattern of life' of the employee, the CEO, and their peer groups, Darktrace was able to immediately flag a number of subtle anomalies in the email, including the forged sender address.

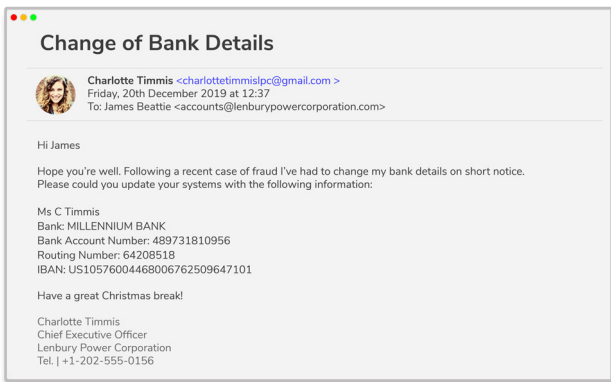


Figure 23: Screenshot of email impersonating the CEO

2. Among other weak indicators, Antigena Email automatically calculated the anomalous proximity of the domain to those of internal employees and trusted contacts.

3. Antigena Email 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 the targeted user and her peer group allowed it to neutralize a high-severity threat that signature-based tools would have missed.

REAL-WORLD CASE STUDY 'Finance VP' aiming to initiate trusted internal relationship

This incident involved the impersonation of a Vice President at a well-known financial services company. 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 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.

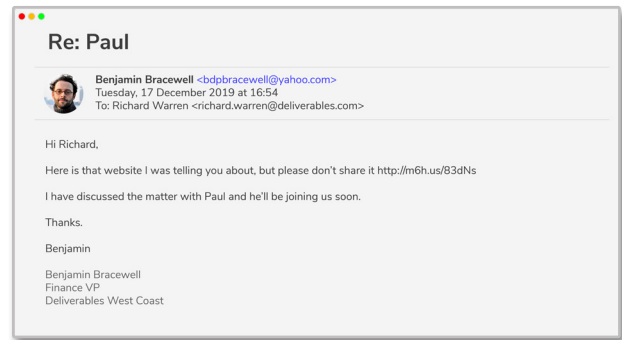


Figure 24: Screenshot of email sharing suspicious link

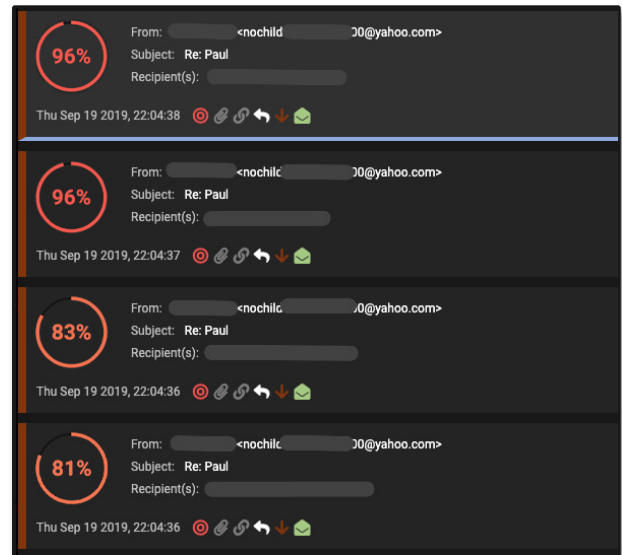



Figure 25: Four of the 11 emails, showing the high anomaly score and associated Antigena Email action

Compromised Employee Credentials

Credential compromise has increased 280% between 2016 and 2019




Attackers may steal email credentials using a variety of methods, from using software that records keystrokes on a compromised machine to stealing databases. 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. Account takeovers are typically monitored for by simple and static defenses, including ‘impossible travel’ rules that rarely catch attackers who know how to hide.

By learning the normal ‘pattern of life’ of every user, self-learning AI 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.



REAL-WORLD CASE STUDY

Compromise across Microsoft 365 and Teams

A Microsoft 365 account was recently compromised at a public accounting firm based in the United States. Darktrace initially picked up on several anomalies, including a sudden surge in outbound email traffic as well as the unusual login location – while the company and nearly all of its users were located in Wisconsin, an IP address located in Kansas was used to log in to the Microsoft 365 account. Along with the unusual login, a login to Microsoft Teams from the same Kansas IP address was detected.

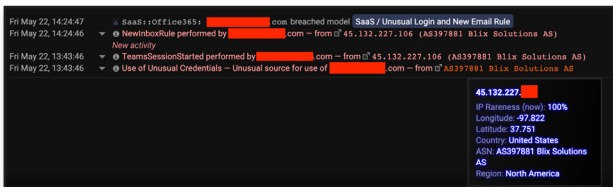


Figure 27: Just after the new email rule was created, a Microsoft Teams 100% rare IP login occurred

‘Impossible travel’ rules alone would have missed these anomalies, but an understanding of activity and behavior across different SaaS applications allowed Darktrace’s AI to recognize these events as one systematic case of credential theft. When the threat-actor subsequently created a new email rule, Darktrace was able to connect this event with the other anomalous behavior and understand its potentially malicious nature.

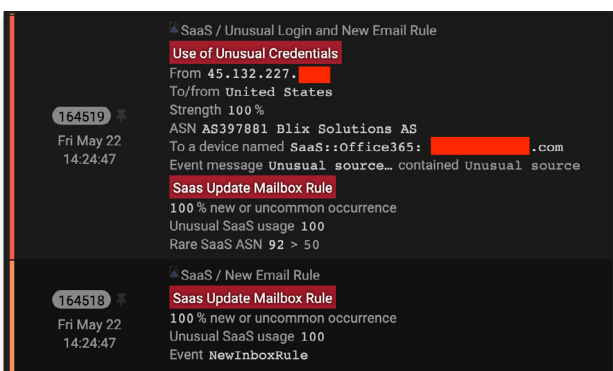


Figure 28: Darktrace’s SaaS Module noted a 100% rare IP logging into the user’s Microsoft 365 account and the creation of a new mailbox rules. All factors indicated 100% unusual SaaS activity

Five minutes later, Antigena Email alerted on a large number of outbound emails containing a generic subject line and an attached PDF. The technology also detected that there was a clear spike in outbound emails from this user and flagged each of these emails with the “Out of Character” tag, which in this case denoted a change from normal behavior with the surge in recipients, and likely internal compromise.




Figure 29: Antigena Email detected a surge in recipients that indicated a serious breach of normal behavior for this user

The unusual login behavior detected by Darktrace’s SaaS Module could be connected to the anomalous outbound email behavior flagged by Antigena Email, allowing the security team to see the extent of the attack and neutralize it as it emerged. It was clear that the account was being used to engage in malicious activity, as each of the 220 outbound emails used a generic subject line and contained a suspicious attachment. The security team therefore immediately disabled the compromised account.

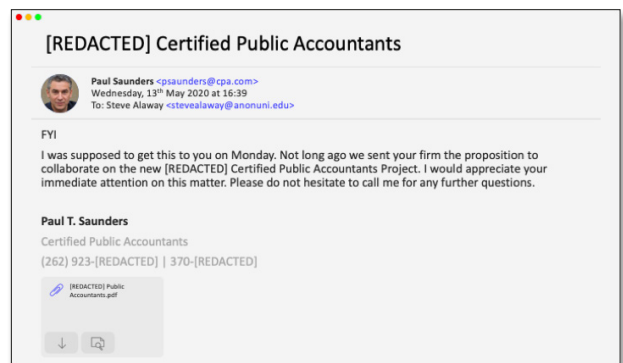


Figure 30: A recreation of the email sent by the attacker, containing the malicious attachment



REAL-WORLD CASE STUDY

'Change of bank details' sent from Accounts Department

When an Accounts Department's Microsoft 365 account was compromised and used to send targeted phishing emails, Darktrace was able to track the attacker's movement within the inbox, tying together information from Darktrace's SaaS Module with Antigena Email's alerts to understand the full picture of the threat and stop the attack.

The SaaS account appears to have been compromised via an inbound spear phishing attack, or some other form of attack that occurred before Darktrace began monitoring the organization. While Darktrace Cyber AI had no oversight of the initial compromise, it was still able to distinguish later attacker behavior as malicious, based on its actively evolving understanding of the organization and its workforce.

When the account user logged in from a 100% rare French IP address, Darktrace's SaaS Module picked up on the anomaly immediately, and further detected a series of activities carried out after the unusual login. At the same time, Antigena Email noted an email being sent.



Figure 31: The login from a French IP was noted as 100% rare for this user and SaaS account

Darktrace then identified more activity occurring from a second rare login location, a Swiss IP address. Very little email activity occurred when the account was logged in from this IP. Instead, Cyber AI saw the threat-actor using their illegitimate SaaS access to view information on the legitimate account user and files related to banking, invoices, and payments.


Then, Antigena Email identified a series of email communications that, when seen in the context of the SaaS account compromise, pointed to a clear threat. There were no obvious malicious attachments or links in the emails. However, the subject of the final reply from the destination is 'Change of Bank Details,' strongly implying that the malicious actor had sent emails instructing the destination to change payment details in order to route money to the attacker, instead of the company.

It seems the attackers went through the banking and invoicing files in order to find a customer with a big bill to pay, then used the compromised email account to launch an outbound phishing attack, changing the billing details. With Darktrace AI correlating information within the SaaS platform and insights from Antigena Email, this targeted phishing attack could be contained before further compromise or damage could occur.

The below screenshot also indicates a series of inbox processing rules made on the compromised account, showing actions that are typical of an account takeover.

Timestamp	Account	Action	Created Name	File Name	File Size	File Type	File Hash	File Content	File Location	File Size
2023-03-11 09:00:00	accounts	rule	SendData							0.100.100.0
2023-03-11 09:00:00	accounts	rule	MoveToDeletedItems							0.100.100.0
2023-03-11 09:00:00	accounts	rule	Copy							0.100.100.0
2023-03-11 09:00:00	accounts	rule	SendData							0.100.100.0
2023-03-11 09:00:00	accounts	rule	Copy							0.100.100.0
2023-03-11 09:00:00	accounts	rule	SendData							0.100.100.0
2023-03-11 09:00:00	accounts	rule	Copy							0.100.100.0
2023-03-11 09:00:00	accounts	rule	SendData							0.100.100.0
2023-03-11 09:00:00	accounts	rule	SendData							0.100.100.0
2023-03-11 09:00:00	accounts	rule	SendData							0.100.100.0

Figure 32: Darktrace's records of new inbox rules being set up on the compromised SaaS account



REAL-WORLD CASE STUDY

Account Takeover 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 33: 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 34: The activity associated with the SaaS account, highlighting the changed credentials

REAL-WORLD CASE STUDY

Unusual External Source

At a financial services corporation based in Europe, an Office 365 user 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.

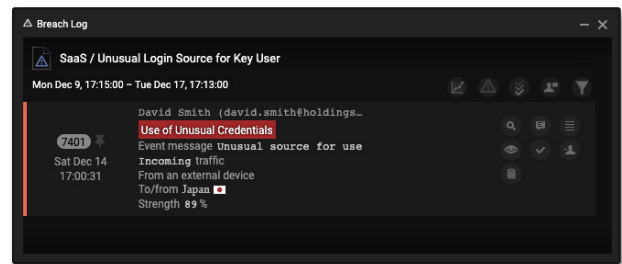


Figure 35: The login from Japan breached several models



REAL-WORLD CASE STUDY

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.

The Cyber AI Platform

Antigena Email can be enriched with additional data sources when deployed as part of Darktrace's wider Cyber AI Platform. This allows Antigena Email to incorporate intelligence from other parts of the business and augment its capabilities even further, extending the AI's understanding of your users' normal 'patterns of life' in the email environment, to their behavior and activity on cloud platforms and SaaS, as well as within a traditional network, and beyond.

This not only enhances Antigena Email's decision-making at the email layer, but enables self-learning protection across your entire digital infrastructure, leaving attackers with nowhere to hide. The AI's continuous analysis in particular is enhanced, as new evidence of a threat that becomes evident in the wider network can be used to retrospectively pull back a delivered email.

No other email security solution can integrate real-time information from a range of sources across the rest of the enterprise. When deployed with the rest of the Cyber AI Platform, Antigena Email completes a truly enterprise-wide security strategy, providing unparalleled visibility of your systems and unmatched protection against even the most sophisticated of cyber-threats.

Discover Antigena Email in your own environment.

[Click here to sign up for a free trial.](#)

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 3,500 organizations to protect against threats to the cloud, email, IoT, networks and industrial systems.

The company has over 1,200 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

[!\[\]\(b58c23cb5aab1cd63092eda333892cb9_img.jpg\) @darktrace](#)