Darktrace Immune System

Self-Learning Detection & Response

The Current Threat Landscape

Contents

Business leaders in the digital age face remarkably urgent risk factors in an era of automated and fast-moving cyber-threat – from the theft and manipulation of critical data, to the staggering losses caused by interruption to the business. These risks have heightened dramatically in recent years as threats develop and become more advanced, and as digital businesses continue to grow in complexity, diversity, and scale.

In the past, when threat actors were less advanced and when digital activity was more predictable, a traditional approach to security was often adequate to keep cyber-threats at bay. By configuring security tools with static rules and historical attack data, organizations have sought to detect threats by defining 'benign' or 'malicious' in advance – relying on representations of attacks that have either been conceived of in the form of a rule, or that have been observed 'in the wild' and reverse-engineered for future detection.

Yet the increasing frequency of novel external attacks and insider threats, together with the exploding complexity of the digital estate, have gradually disarmed security teams who still rely on traditional controls. These rigid defenses fail to detect the novel tactics and techniques of sophisticated cyber-criminals, who can now blend into the noise of the network and sweep through large and complex infrastructures within seconds.

Beyond the corporate IT network, security teams must also protect a diverse and fragmented patchwork of SaaS applications, cloud workloads, industrial machinery, and email platforms – all of which come with their own complex and incompatible controls. The interrelation of workforce behaviors across these different environments has rendered point solutions inoperable, as they lack the unified scope required to catch threats unfolding across the entire organization.

The fact is that targeted attacks will inevitably get inside, and so the industry's attention has shifted to the question of how defenders can be equipped to detect and respond to emerging threats that are already inside the business, but that can be handled before they become a crisis. And as in many other areas plagued by digital complexity, business leaders and security teams have ultimately turned to artificial intelligence to keep pace.

The Darktrace Immune System

A Self-Learning Approach

While traditional defenses continue to define the threat in advance, Darktrace focuses instead on learning the normal 'pattern of life' for individual businesses, and spotting subtle deviations indicative of a threat. Like the human immune system, the technology learns 'on the job', from the data and activity that it observes in situ. This means making billions of probability-based calculations in light of new evidence and continuously learning as the business evolves.

The threats that infiltrate your organization will typically not be historical attacks, but rather novel threats that have evaded existing defenses, or inappropriately behaving employees and third parties. By learning a sense of 'self' for your entire organization, Darktrace's immune system discovers subtle, previously unseen patterns and emerging threats that would otherwise go unnoticed.

Darktrace's core detection engine uses unsupervised machine learning to build a dynamic understanding of 'normal' for each organization it safeguards. Rather than rely on rules, signatures, fixed baselines, or training data, the immune system learns from your constantly changing digital environment – forming a bespoke and multi-dimensional understanding of every user, device, and all the complex relationships between them.

This unique self-learning approach enables Darktrace to detect advanced attacks at an early stage, and well before they have time to escalate into a crisis – from a novel strain of ransomware or an insider attack, to a coordinated spear phishing campaign or critical cloud misconfiguration.

Threat Visualizer

Darktrace's Threat Visualizer provides real-time visibility of your entire digital infrastructure, surfacing insights across email, cloud, and the corporate network in a single pane of glass. Cyber-threat visualization and investigation is simplified with this intuitive and easy-to-use graphical interface.

The Threat Visualizer allows the user to 'go back in time' to when an incident took place, and witness events as they unfold in real time. Only the most relevant threats are presented, allowing for incident prioritization, with the option to drill down into any single event in finer detail.
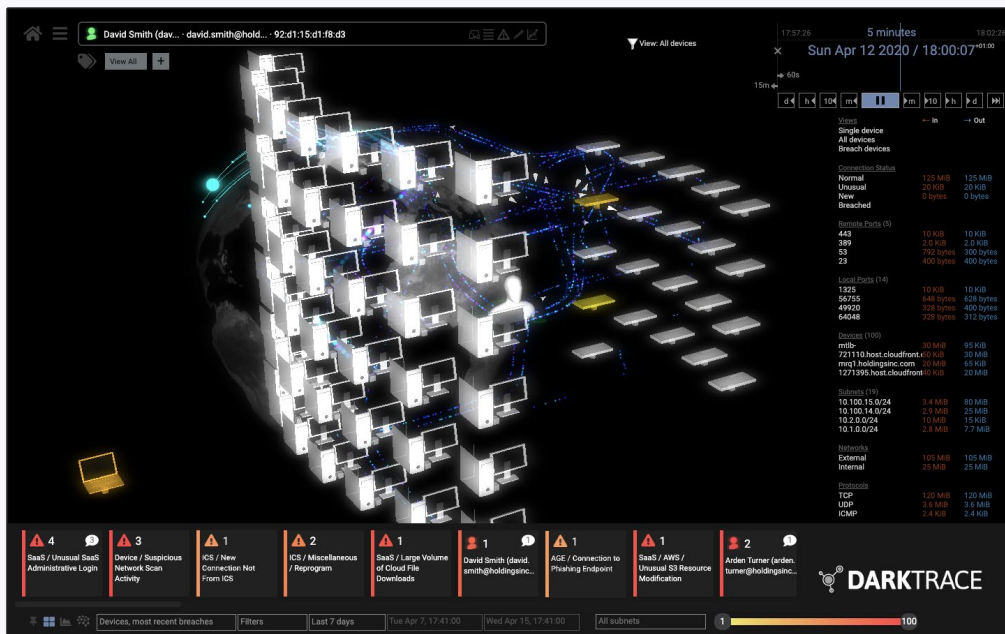

Figure 1: The threat tray at the bottom of the Threat Visualizer surfaces threats identified across the digital business

Autonomous Response

Darktrace has also delivered the world's first proven Autonomous Response technology on the market with Darktrace Antigena.

With this innovation, Darktrace has evolved to not only detect but also intelligently fight back against in-progress attacks before they can make an impact. Using the Enterprise Immune System's rich understanding of the entire digital business, Darktrace Antigena takes swift and targeted action to interrupt attacks with precision, even if the threat is targeted or entirely unknown.

Rather than generate broad-brushed quarantines that would only cause more disruption, Antigena works by surgically enforcing the normal 'pattern of life' of an infected device or compromised user, neutralizing the threat in seconds and sustaining normal operations by design. These self-directed actions are not only granular, but also dynamically adapt to the severity of the threat as it unfolds.

Beyond this tactical protection, Darktrace Antigena can also deliver strategic response by acting as the 'AI brain' of the entire security stack, leveraging high-confidence detections to hand off and integrate with inline defenses as a mechanism for response. Through active integrations, Antigena can seamlessly plug into and enhance your existing ecosystem, informing firewalls and network devices about attacks that have gotten through.

With Autonomous Response working in tandem with Darktrace's core immune system, Cyber AI gives control back to the defenders, transforming even the most complex and vulnerable organization into a resilient, self-defending digital business.
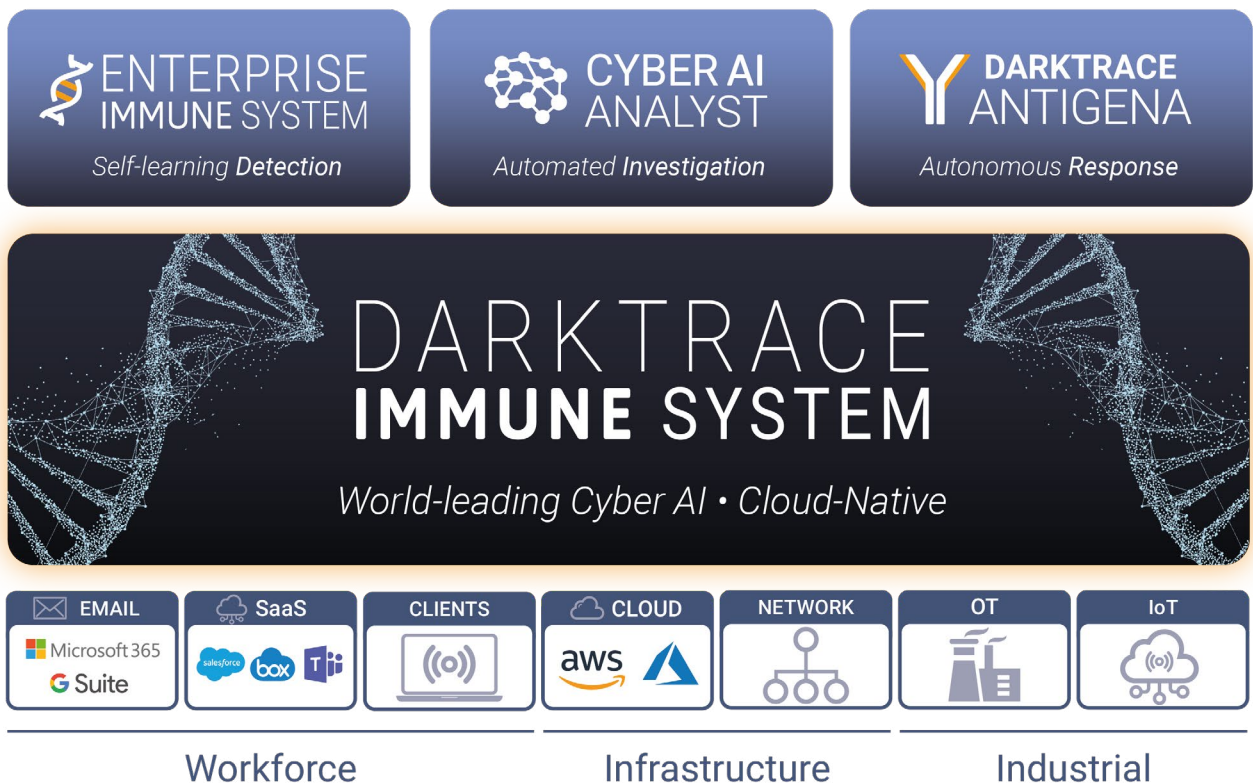

Figure 2: The Darktrace Immune System

Cyber AI Analyst

While Darktrace's Immune System and Antigena speed up 'time to detection and response', Darktrace's Cyber AI Analyst drastically reduces 'time to meaning' by fully automating threat investigations for the first time.

Human security analysts typically investigate threats by following leads, forming hypotheses, reaching conclusions, and sharing their findings with the rest of the business. These are labor-intensive steps that take time and require expertise, often against the backdrop of machine-speed threats that outpace the inherently limited reach of human responders. By combining expert analyst intuition with the speed and scalability of AI, Cyber AI Analyst transcends these limitations with AI-driven investigations that reduce time to triage by up to 92%.

When Darktrace's Immune System detects a pattern of suspicious behavior, Cyber AI Analyst launches into an enterprise-wide investigation, stitching together disparate anomalies before settling on a high-level conclusion about the nature and root cause of the wider security incident. Because the AI can operate everywhere at once, it can generate thousands of queries and follow hundreds of parallel threads simultaneously, rapidly illuminating the full scope of incidents in real time.

Critically, Cyber AI Analyst not only automates analyst workflows at speed and scale, but also preserves the inherent flexibility of human expertise. This means that the system can quickly interpret and report on security incidents characterized by innovative attack techniques that would be impossible to capture with pre-defined playbooks.

By continuously investigating 100% of the security events that Darktrace's Immune System detects, Cyber AI Analyst produces a dynamic situational dashboard as well as written reports that immediately put resource-strained security teams in a position to take action.

Darktrace Cyber AI Analyst Stops Zero-Day Vulnerability Attack

Cyber AI Analyst recently proved crucial when a number of Darktrace customers were hit by an attack targeting the Zoho ManageEngine zero-day vulnerability CVE-2020-10189. The intrusions were later attributed to Chinese threat actor APT41, and formed part of a wider campaign aiming to gain initial access to as many companies as possible during the window of opportunity presented by the vulnerability.

Darktrace automatically detected and investigated the attack in its earliest stages, enabling customers to contain the threat before it could make an impact. The reports generated by Cyber AI Analyst highlighted and delineated every aspect of the incident in the form of a meaningful security narrative. Even a junior responder could have reviewed this output and acted on this zero-day APT attack in under 5 minutes.
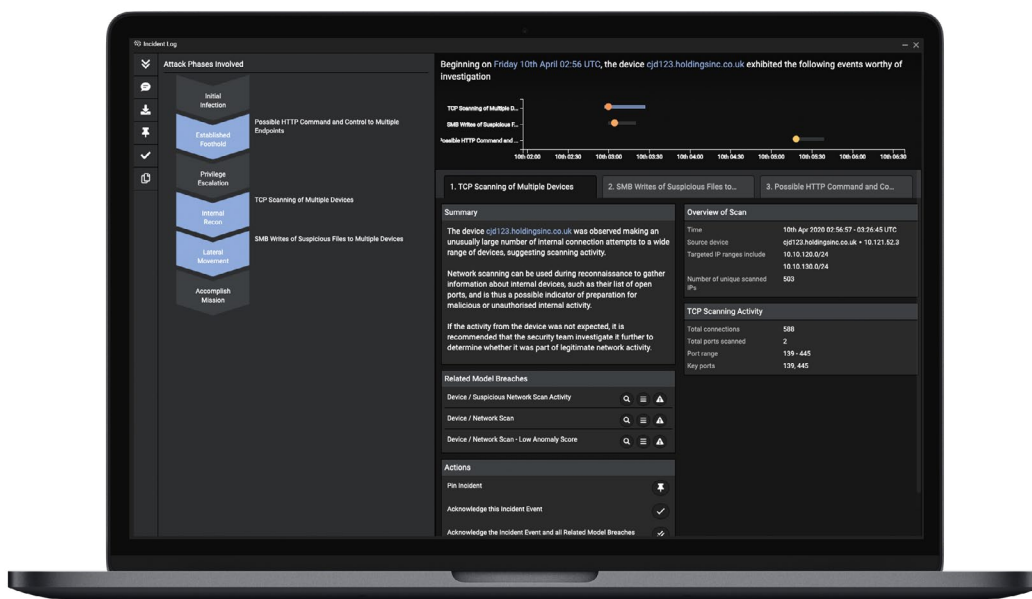

Figure 3: An incident overview generated by the Cyber AI Analyst, including the attack stages involved

Enterprise-Wide Protection

Increasingly, threat actors aren't limiting their attacks to one technology at a time, and as defenders it is essential that protections are unified across the entire digital business. Something as simple as a compromised password can result in an attack against multiple facilities at once. Darktrace's Immune System is specifically designed to cut across multiple stovepipes and enable unified detection and response, spanning across email, cloud, and the corporate network.

Insights across these diverse environments are not only surfaced in the same unified view, but also fused together and correlated by a single AI engine in the background. This design principle takes seriously the idea that the full scope of a device or user's normal patterns are made manifest in different parts of an organization, and that a single security incident typically includes related events and indicators that occur elsewhere in the digital environment. Being able to see this in real time is essential for meaningful incident management – it no longer makes sense to handle security on a per-technology basis.

As well as unifying detection and response, Darktrace believes strongly in enabling full visibility. For today's security teams, tooling must facilitate the ability to explore and illuminate multiple environments at will – rather than just simply generating security alerts.

In the real-world case studies that follow, Darktrace's Immune System identified attacks based on its unified understanding of 'normal' across cloud, SaaS, email, industrial, and the corporate network.
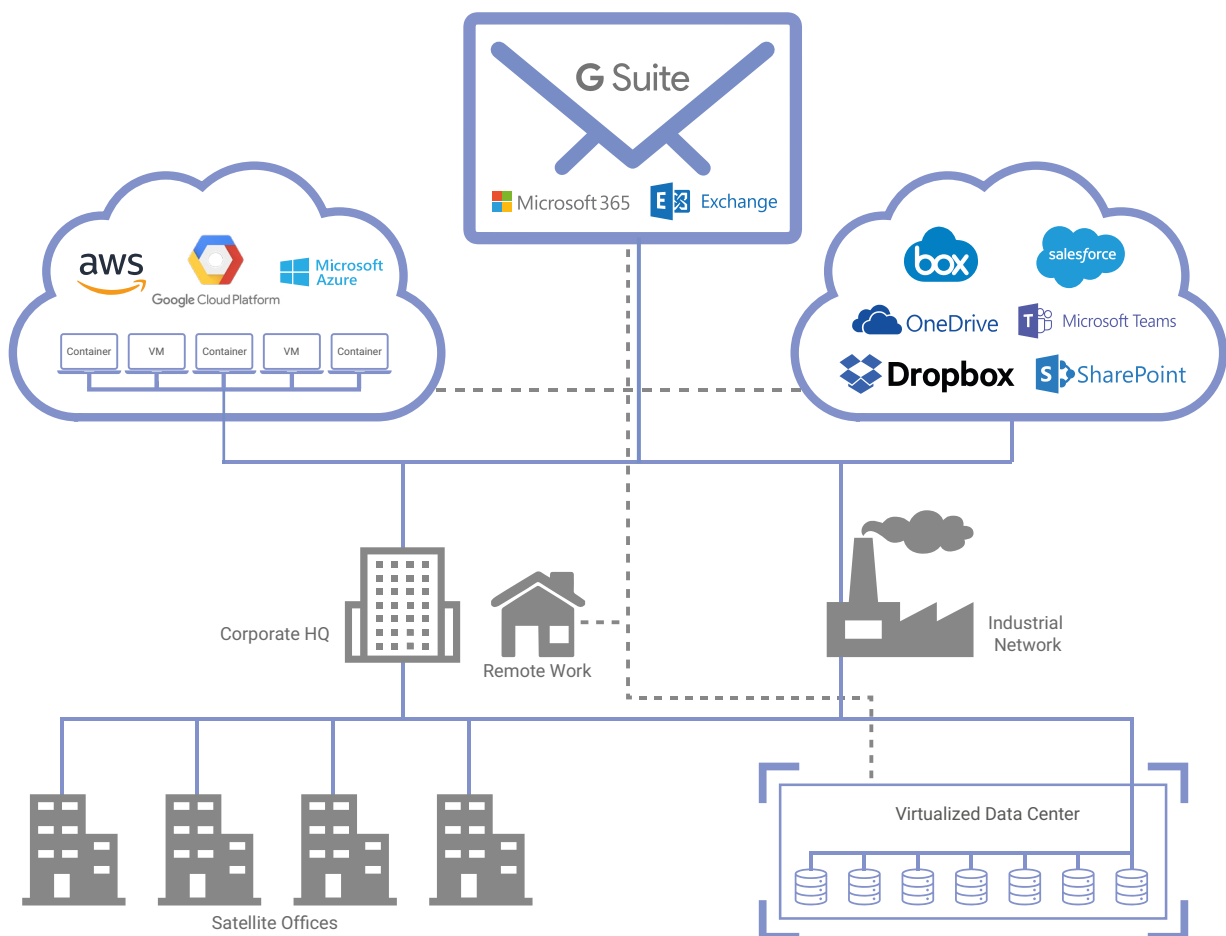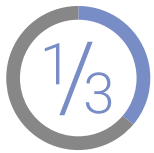

Figure 4: Darktrace's overview across the entire digital business

Cyber AI for Cloud & SaaS

1/3

Less than a third of businesses are monitoring abnormal behavior across their cloud enviroment

Source: Cybersecurity Insiders

The large-scale journey to the cloud has fundamentally reshaped the digital business and the traditional paradigm of the network perimeter. Hybrid infrastructure and distributed workers are now a part of the furniture of an increasingly diverse digital estate, with multi-cloud practices introducing a new layer of complexity that most organizations are ill-equipped to address.

In the cloud, security teams not only struggle with a lack of visibility and control, but also diverse and incompatible defenses that often lead to overly relaxed permissions and simple mistakes. This traditional 'stove-pipe' approach to security is rarely robust and unified enough to provide sufficient coverage, relying on static and siloed methods that fail to detect compromised credentials, insider threats, and critical misconfigurations.

Darktrace's Immune System fills these gaps with self-learning AI that understands 'normal' at every layer, dynamically analyzing the dispersed and unpredictable behaviors that show up in email, cloud, and the corporate network. This unified scope allows the system to spot subtle deviations indicative of a threat – from an unusual resource creation or open S3 bucket in AWS, to suspicious data movement in Salesforce, to a new inbox rule or strange login location in Microsoft 365.

Unlike policy-based controls, the immune system understands the human behind every trusted account in the cloud, providing a unified detection engine that can correlate the weak and subtle signals of an advanced attack.

Figure 5: Darktrace's dedicated SaaS Console gives an overview of anomalous behavior in SaaS applications and displays the geographical locations of the activity

DARKTRACE

M365 Compromise and SharePoint Infiltration

At a US-based insurance company, Darktrace Cyber AI's bespoke knowledge of 'normal' and visibility across SaaS platforms was crucial for stopping an attack that started with a compromised Microsoft 365 account.

When a threat actor successfully logged in to one of the client's Microsoft 365 accounts from an IP address located in the United Arab Emirates, Cyber AI identified the behavior as anomalous, as no other M365 accounts had ever been observed logging in from this IP address. Four days later, another rare IP located in the UAE was seen accessing the same compromised account. This time, the threat actor set up a new email rule, and used their illegitimate access to read and write to files on the user's personal SharePoint account.

Darktrace Cyber AI had not previously seen any other user accounts communicating with UAE-based IPs from the particular network identified in these incidents, indicating that the observed behavior was highly unusual for the customer and the result of compromise.

While the customer's legacy tools only allowed them to see the threat when changes were made to the compromised account, Cyber AI detected the anomalous behavior as soon as it occurred and clearly illuminated the attacker's movement between SaaS services. Darktrace was able to alert the security team immediately of the earliest stages of the attack, shining a light on every detail and ensuring the threat was neutralized before serious damage could occur.

Cloud Misconfiguration

A leading manufacturing company in Europe was using a Microsoft Azure server to store files containing product details and sales projections. Whilst the files on the server and the root IP were gated with a username and password, this sensitive data was then left unencrypted. Anomalous activity was detected when a device downloaded a ZIP file from a rare external IP address that Darktrace deemed highly anomalous.

It was later discovered that the external IP was a newly configured Microsoft Azure server and the ZIP file was accessible to anyone who knew the URL, which could have been obtained by simply intercepting network traffic, either internally or externally. More dedicated attackers could have even brute-forced the file 'key' parameter of the URL.

The loss or leakage of the sensitive files in question could have placed an entire product line at risk, but in reporting this incident as soon as it was detected, Darktrace helped to prevent the loss of valuable intellectual property, and proceeded to assist the security team in revising their data storage practices in the cloud in order to better protect their product information moving forward.
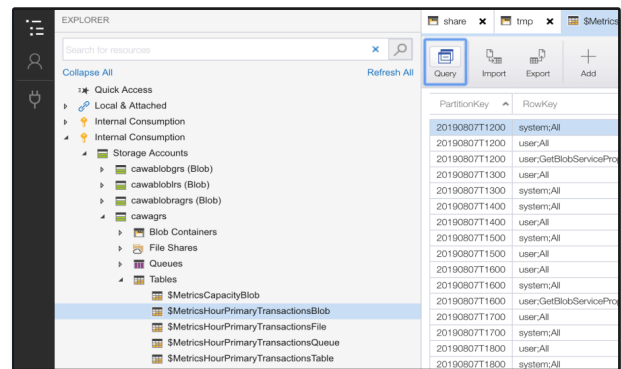

Figure 6: The sensitive files in Azure

Suspicious Box File Download

At a global produce supplier, several suspicious requests within the company's Box platform suggested that a user account had been compromised.

The actor behind the account logged in to Box successfully, and then proceeded to download expense reports, invoices, and other financial documents. The potential threat actor also went on to unlock a file containing a list of sensitive passwords.

With Cyber AI's bespoke knowledge of 'self' for every member of the organization's workforce, the technology was able to identify the threat immediately. Darktrace's Immune System detected that the activity occurred at a highly unusual time for the legitimate user, and that the location of the actor's IP address was also anomalous compared to the employee's previous access locations for this particular SaaS service.

While accessing these documents may have been normal for the employee in another context, Darktrace Cyber AI's deep understanding of user behavior and granular visibility within Box allowed it to spot the subtle signs of account compromise. When Cyber AI Analyst autonomously investigated, it was able to illuminate the wider narrative, understanding that each unauthorized file exposure was part of a connected incident and highlighted the breach as a key concern for the security team.
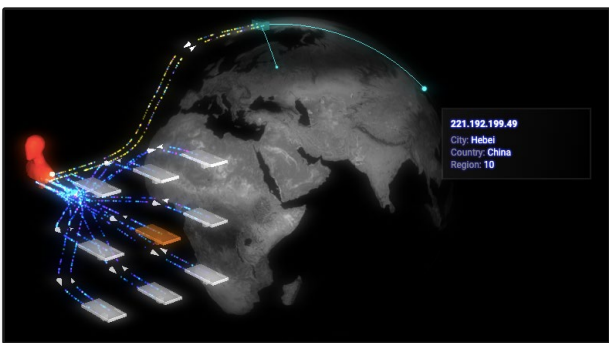

Figure 7: Darktrace showing the location of the unusual IP address

Attack Evades 'Impossible Travel' Rule in Microsoft 365

In one international non-profit, Darktrace detected an account takeover in Microsoft 365 that bypassed Azure AD's 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.

> **Darktrace detected an account takeover in Microsoft 365 that bypassed Azure AD's static 'impossible travel' rule.**

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.
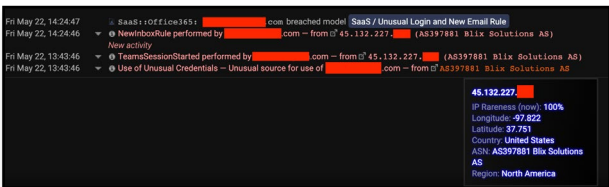

Figure 8: Just after the new email rule was created, a Microsoft Teams 100% rare IP login occurred

'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.
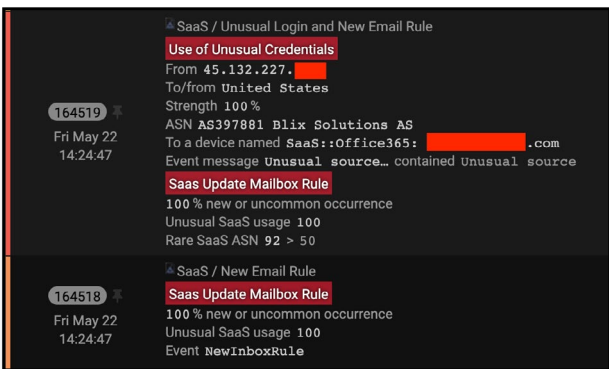

Figure 9: Darktrace's SaaS Module noted a 100% rare IP logging into the user's Microsoft 365 account and the creation of a new mailbox rules. All factors indicated 100% unusual SaaS activity

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.
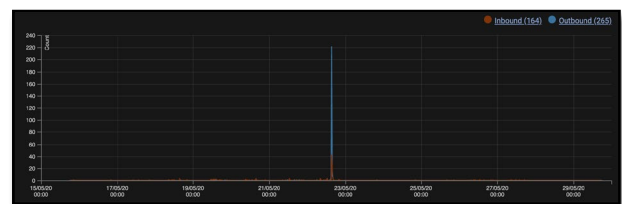

Figure 10: Antigena Email detected a surge in recipients that indicated a serious breach of normal behavior for this user

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.
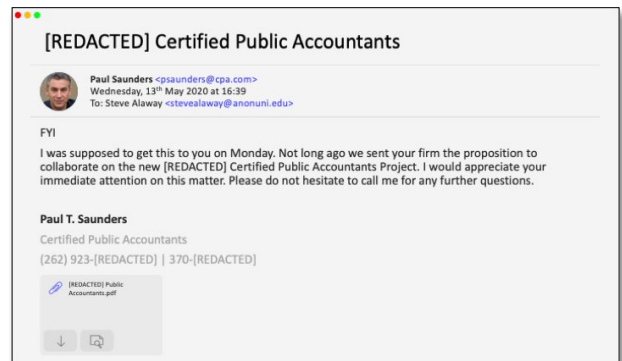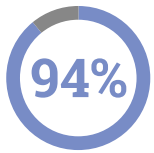

Figure 11: A recreation of the email sent by the attacker, containing the malicious attachment

Cyber AI for Email

94%

94% of cyber-threats enter an organization through the inbox

Source: Data Breach Investigations Report

By spoofing an email or hijacking a trusted account, cyber-criminals can trick users into wiring millions out of the business or triggering a ransomware attack with a single click. Whether native or third-party, traditional email controls work by analyzing emails in isolation and at a single point in time, correlating them against blacklists, signatures, and pre-definitions of bad. While this approach can often catch basic spam and similarly indiscriminate 'drive-by' campaigns, it routinely fails to spot the weak indicators of an advanced social engineering attack or stealthy spear phishing campaign.

Yet by analyzing the normal 'pattern of life' for every user and correspondent, Darktrace's Antigena Email can uniquely develop an evolving understanding of the 'human' within email communications. Powered by Cyber AI, it is the only technology 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 multi-dimensional understanding enables the system to make highly accurate decisions and neutralize the full range of advanced email attacks, from 'clean' spoofing emails to supply chain account takeover.

Antigena Email works by learning the dynamic patterns of every internal and external user, analyzing both inbound and outbound email together with lateral, internal-to-internal communications. By treating recipients as dynamic individuals and peers, Antigena Email can spot subtle deviations from 'the norm' that reveal seemingly benign emails to be unmistakably malicious.
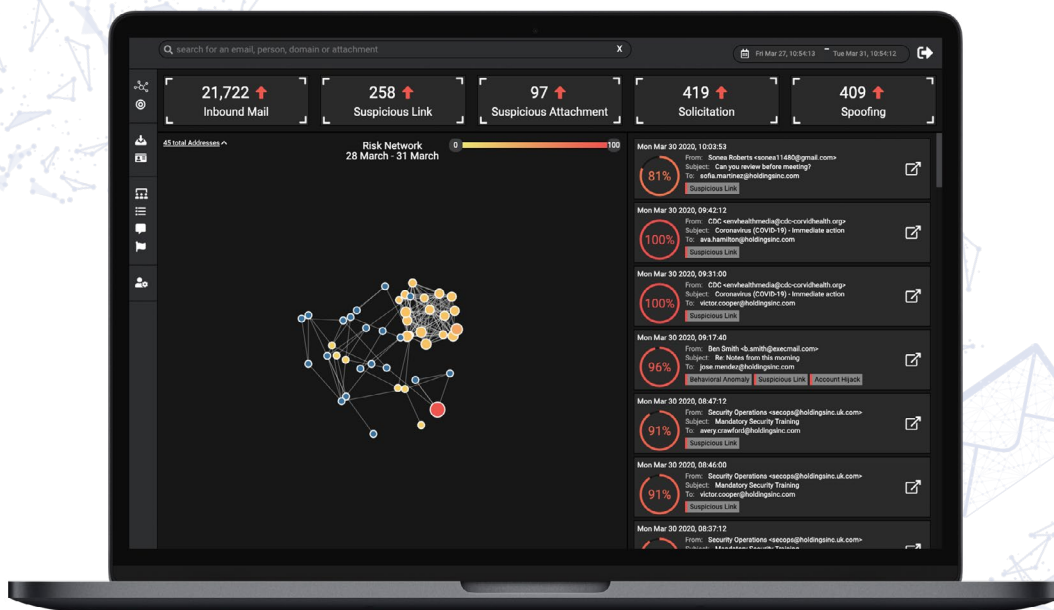

Figure 12: Antigena Email's interface displaying an overview of alerts

Coordinated Spoofing Attack

Darktrace detected a highly targeted social engineering attack impersonating C-level executives at a US technology company, when a threat actor apparently sent a number of 'clean' emails in an effort to garner trust and establish offline communications, preemptive of a request for payment. While the legacy email defenses in place were unable to detect the attack given their static analysis and limited scope, Darktrace held back every email from the intended recipients based on the following observations.

1. Abnormal Subject and Sender. The emails had the first name of the targeted employee as the subject line, and further were sent from a seemingly unrelated Gmail address.
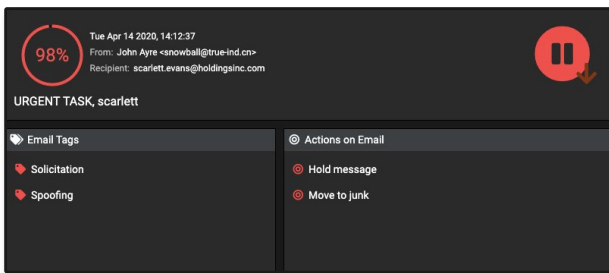

Figure 13: One of the 30 emails, with a 98% anomaly score

2. No Association. Across Darktrace's entire understanding of the company's email and network environment, Antigena had seen no evidence of a relationship between this sender and the organization.

3. Exposing the Whale Spoof. Darktrace 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. In addition, the exposure score of the impersonated users was high, indicating that they were high-profile targets subject to a 'Whale Spoof' attack.

Correlating these multiple weak indicators, Antigena recognized the emails as components of one systematic attack, causing it to hold them back in a buffer for the organization's security professionals to review – preventing the targeted recipients from engaging with the contents of the email and establishing offline communications.

Supply Chain Takeover Email Attack

At a multinational energy corporation, Darktrace identified a supply chain attack, recognizing that the sender was well known to the company, with a number of internal users having previously corresponded with them. Less than two hours after a routine exchange, emails were sent rapidly to 39 users, each containing a phishing link. Variation in the subject lines and links suggested highly targeted emails from a well-prepared attacker, but Antigena held all 39 emails back and double-locked the payloads, based on the following anomalies:

1. Unusual Login Location. Extracting the geo-locatable IP address revealed 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 links were all hosted on the Microsoft Azure developer platform – likely to skirt reputation checks on the host domain, but highly inconsistent for the sender based on previous correspondence history, as well as 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, Darktrace deduced that this recipient group was 100% anomalous by just the third email.

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, Darktrace recognized that these recipients do not typically receive emails about business proposals using this style of phrasing.
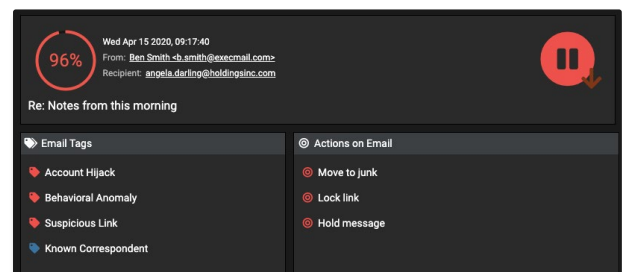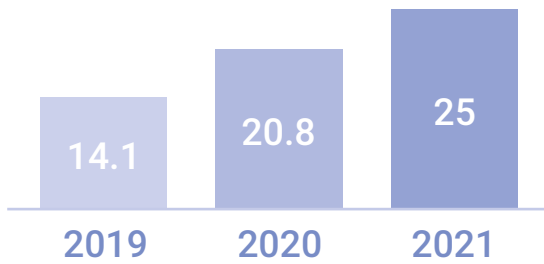

Figure 14: Darktrace detected the account takeover and held the emails back

Cyber AI for the Internet of Things

Number of connected IoT devices (billions)

2019	2020	2021
14.1	20.8	25

Source: Gartner

From smart coffee machines to internet-connected CCTV cameras, the advent of IoT has introduced an entirely new threat vector in the enterprise. For all their convenience and appeal, most IoT devices were not created with security in mind, often providing an easy inroad or a surreptitious avenue of exfiltration.

For devices that can support it, traditional endpoint security is useful for stopping known threats, but companies need a much broader strategy for dealing with the unpredictability of IoT. In most cases security teams can't run standard anti-virus software on smart devices given a lack of disk space, a CPU, or a traditional operating system. Still worse, in order to install an endpoint solution on a smart device, one would need to know that it exists in the first place. Yet most organizations lack the network visibility required to provide an accurate number of their core workstations and servers, let alone of their IP-based IoT devices.

To address IoT security, we need to take a bigger picture view – not only looking at vulnerabilities or managed devices, but also complex behaviors that show up across the digital business. With Cyber AI, organizations can monitor 100% of their devices, wherever they are on the network. Learning a normal 'pattern of life' for every device, Darktrace can spot the full range of attacks targeting the Internet of Things – from smart fish tanks to autonomous vehicles. Darktrace Antigena then responds in real time, containing the threat and mitigating risk in every corner of your organization.

Figure 15: A compromised smart printer and anomalous connections represented with yellow lines

Corporate Espionage Through CCTV Hack

At a Japanese investment consultancy, Darktrace discovered that an internet-connected CCTV system had been infiltrated by unknown attackers. The perpetrators had used the device to gain a foothold into the network, and could watch all of the camera's video recordings from there. Installed to monitor the entire office space, from the CEO's office to the boardroom, the camera became the security risk.

Darktrace's AI quickly detected that something was amiss. Massive volumes of data were observed moving to and from the unencrypted CCTV server, as the attacker gathered data in preparation to exfiltrate sensitive information. At the point when the attacker tried to exfiltrate the data, Antigena took rapid and precise defensive action.
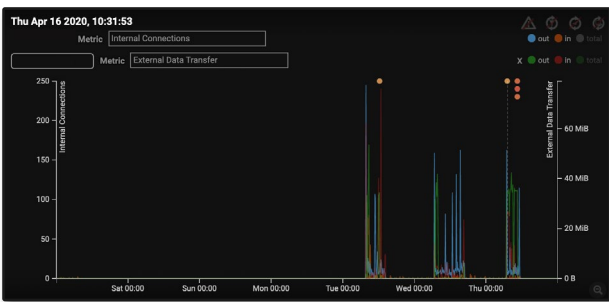

Figure 16: Internal and external connections pertaining to and from the IoT device

The system decided to surgically block data movement from the device to an external server – while still allowing the CCTV to operate in its intended capacity. The AI fought back at machine-speed, preventing a serious breach of market-sensitive information.

By taking proportionate action to contain the attack at an early stage, Antigena gave the security team vital time to investigate and remediate the threat before any damage was done.

Sensitive Data Exfiltrated Through Smart Locker

An amusement park in North America came under attack when a threat actor attempted to steal sensitive customer data via a vulnerable internet-connected smart locker.

As part of its default setting, the smart locker regularly established contact with the supplier's third-party online platform. The threat actor identified the source of this automated process, and hijacked it to compromise the device.

Darktrace's AI spotted the attack shortly after the locker started sending an unusual quantity of unencrypted data to a rare external site. The connections were timed in accordance with the device's regular communications with the supplier's platform, suggesting that this was a 'low and slow' attack specifically designed to evade rules-based security defenses.

Figure 17: The anomalous number of connections to the smart locker

By continuously analyzing the communications in relation to the locker's prior behavior and that of its peers, Darktrace's AI determined that an AI cyber response was required. Within seconds, Darktrace Antigena took action, intelligently blocking all outgoing connections from the compromised device, giving the security team time to remediate the threat and prevent any exfiltration.

By learning 'on the job', and continuously revising its understanding in light of new evidence, Darktrace's Immune System detects subtle threats that other tools miss, and generates autonomous actions that adapt to the threat as it unfolds.

Cyber AI for Industrial Networks

90%

90% of OT security teams suffered at least one damaging cyber-attack in the last two years

Source: Ponemon

Traditionally isolated from the Internet, Industrial Control Systems (ICS) have been increasingly converged with the corporate IT network, in order to meet new business objectives and efficiency measures. Unfortunately, from a security perspective, this introduces an array of new challenges in the security of operational technology.

Decades-old devices, built without security in mind, are now exposed to cyber-criminals scanning an organization's perimeter for any vulnerability. Exposed machinery is often used as a gateway for a more pernicious attack on the network, and attacks that start in the IT network can result in collateral damage to physical operations, causing catastrophic losses to production.

With industrial environments growing in size and scope, organizations are turning toward AI for a more in-depth and effective response to these cyber-physical attacks. Darktrace's unified insights and analysis across OT and IT allows the technology to spot a threat as soon as it enters the organization, wherever it enters. Here as elsewhere, customers have also found the insights of Cyber AI Analyst invaluable for the technical translation work and high-level summaries of incidents presented at machine speed.
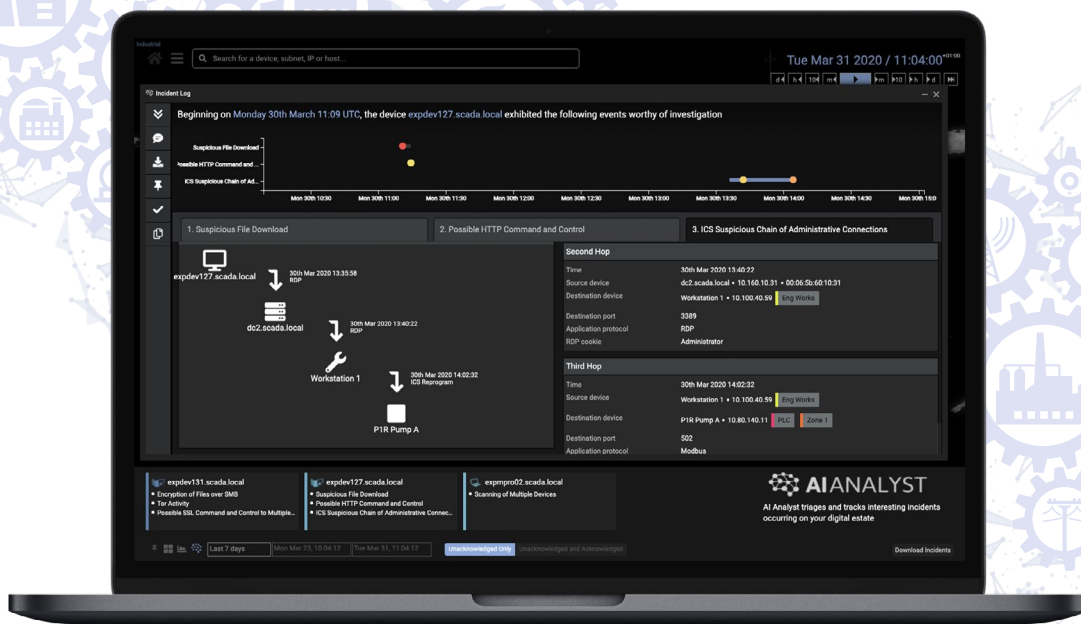
Figure 18: The Cyber AI Analyst surfacing all remote desktop 'hops' in a Triton-style attack targeting IT and OT

Shamoon Virus Detected

The Shamoon malware wipes compromised hard drives and overwrites key system processes, intending to render infected machines unusable. Darktrace detected this notorious cyber-attack directly targeting Industrial Control Systems during a trial period at a global energy company.

Darktrace observed a Shamoon-powered cyber-attack when several Middle Eastern firms were impacted by a new variant of the malware.

Darktrace Cyber AI detected unusual network scans on remote port 445 conducted by dozens of infected devices simultaneously, as well as unusual Remote Powershell usage. Remote PowerShell is quite often abused in intrusions during lateral movement. The devices involved did not classify as traditional administrative devices, making their use of WinRM even more suspicious.

Figure 19: A clear plateau in increased internal connections can be seen. Every colored dot on top represents an RDP brute force detection

Darktrace later identified another cluster of activity likely to represent unusual credential usage. Correlating these insights together with the abnormal use of certain protocols allowed Darktrace to identify a number of related anomalies that were highly unusual for the organization's environment as a whole, and identify this as attackers moving laterally in the network.

Scanning Tools Targeting ICS

ICS systems often introduce blind spots in an organization's traditional cyber security defenses. Darktrace illustrated this whilst being trialed in a utilities organization, when an air conditioning control system was observed receiving a large number of connections over an unusual communication channel from multiple devices outside of the network, and in fact outside of the country where this network was located.

Upon closer investigation, Darktrace had seen connection requests specifically related to vulnerability scanning using a reconnaissance tool, suggesting an attempt to gain illicit access to the device. Furthermore, the external device had then requested to read data from the control unit, indicating access to potentially sensitive ICS information by an external party. This incident illustrates Darktrace's capacity not only for widespread visibility over both IT and OT networks, but also for fine-grained investigation of anomalous connection attempts to internal devices from outside the network.
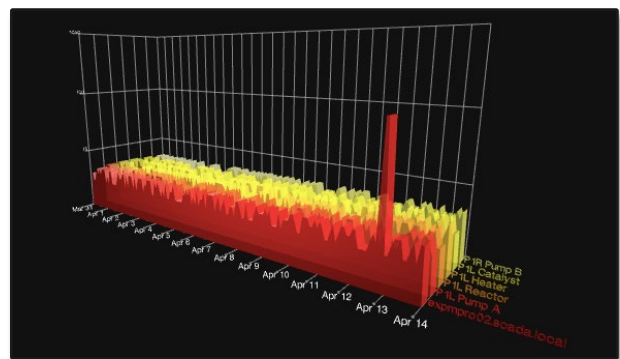

Figure 20: The anomalous connections from the SCADA device are clearly visible

Cyber AI for the Network

69%

69% of organizations think AI is necessary to respond to cyber-attacks

Source: Capgemini Research Institute

Darktrace's self-learning AI is designed to protect the dynamic systems and workers in your organization – no matter where they operate, or the nature of their applications. Unlike legacy on-prem defenses, Darktrace's understanding of normal behavior in the network is augmented by behaviors in cloud, SaaS, and email services as well. This additional context enables Darktrace to detect the full range of cyber-threats in the network, from 'low and slow' data theft and compromised credentials, to machine-speed ransomware.

In turn, Darktrace Antigena surgically interrupts emerging threats in the network at machine speed, giving security teams time to catch up before critical data can be lost or encrypted. This dynamic protection is as intelligent and surgical as it is far-reaching, automatically neutralizing ransomware, crypto-mining operations, and insider threat via self-directed actions and active integrations with inline defenses.

Equally, real-time insights from the corporate network also inform the immune system's decision-making on data points in other areas of the business. If, for example, a device becomes infected after an employee clicks a malicious link in an email, Darktrace can interrupt the infection in the network and automatically identify and neutralize any other emails that are part of the same campaign.

In every case, detections in the network serve as launching points for Darktrace's Cyber AI Analyst to investigate the full scope of the incident at speed and scale. By automatically generating a detailed incident report that can be consumed and actioned in minutes, Darktrace harnesses the full power of artificial intelligence to not only stop threats in seconds, but also allow human teams to focus on more strategic work.

Figure 21: Darktrace detecting a laptop carrying out a network scan

Sodinokibi Ransomware Infects Financial Services Firm

Darktrace detected a targeted Sodinokibi ransomware attack targeting a mid-sized US service company. This 'double-threat' runs targeted attacks using ransomware while simultaneously attempting to exfiltrate its victims' data, enabling the attackers to threaten to make data publicly available if the ransom is not paid.

Darktrace identified the initial compromise when an external-facing RDP server began to make anomalous connections to a rare external IP address in Ukraine. The AI then detected a download of 300MB data from file sharing platform Megaupload, recognizing that nobody in the organization regularly used this service, and therefore instantly flagging it as unusual.

Three minutes later, Darktrace detected a network scan, and then persistent command-and-control traffic, as the infected RDP server started making highly anomalous connections to external destinations. Finally, the AI detected an upload of around 40GB of data, followed by unusual files being accessed on internal SMB shares, which appeared to be ransom notes.
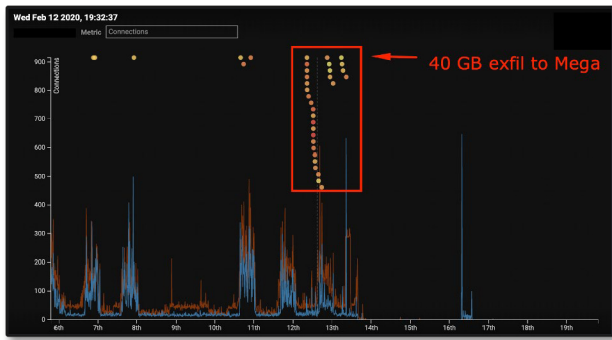

Figure 22: Connections to the domain controller

Over twenty Darktrace models were triggered in the final stage of the attack alone. Had Darktrace Antigena been active, it would have responded to neutralize the threat within seconds.

Bitcoin Mining Under the Hood

An acclaimed 500-person law firm had traditional security controls that scanned for known threats, and yet was unaware that bitcoin mining had been taking place within their network for a period of 5 months.

After installing Darktrace, it transpired that a summer intern had installed bitcoin mining malware on the company's infrastructure, co-opting more than 75 computers. As well as slowing down the network and therefore negatively impacting the firm's productivity, this crypto-mining operation exposed the company to significant reputational risk.

Had the AI not caught this anomalous behavior, the operation could have continued for many months – long after the internship had ended.
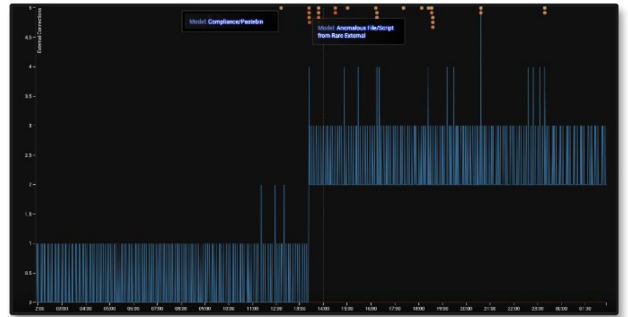

Figure 23: Graphical representation of the sudden increase in external connections and related model breaches

Discover Cyber AI in your own environment.

Contact Gene Fourney at TechnologyWest today!

970.579.1585
gene@technologywest.io

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.