DARKTRACE
Evolving threats call for evolved thinking

Detection and Response Across the Enterprise

CONTENTS

The Current Threat Landscape

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, endpoints, 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. Therefore, attention in the industry has shifted to focus on emerging threats and proactive measures. In other words, how can defenders be better equipped to detect and respond to emerging threats that are already inside the business and how can they handle these threats before they become a crisis. Business leaders and security teams, challenged with digital complexity, are increasingly turning to artificial intelligence to keep pace.

The cost of credential theft to organizations increased 65% from $2.79 million in 2020 to $4.6 million at present.

77 → 85 days

The time to contain an insider threat incident increased from 77 days to 85 days, leading organizations to spend the most on containment.

$17.19m

Incidents that took more than 90 days to contain cost organizations an average of $17.19 million on an annualized basis.

Figure 1: Research conducted independently by the Ponemon Institute

Darktrace AI

/ 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. Darktrace's AI 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 discovers subtle, previously unseen patterns and emerging threats that would otherwise go unnoticed.

Darktrace's core detection engine uses AI technology to build a dynamic understanding of 'normal' for each organization it safeguards. Rather than rely on rules, signatures, fixed baselines, or training data, Darktrace 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.

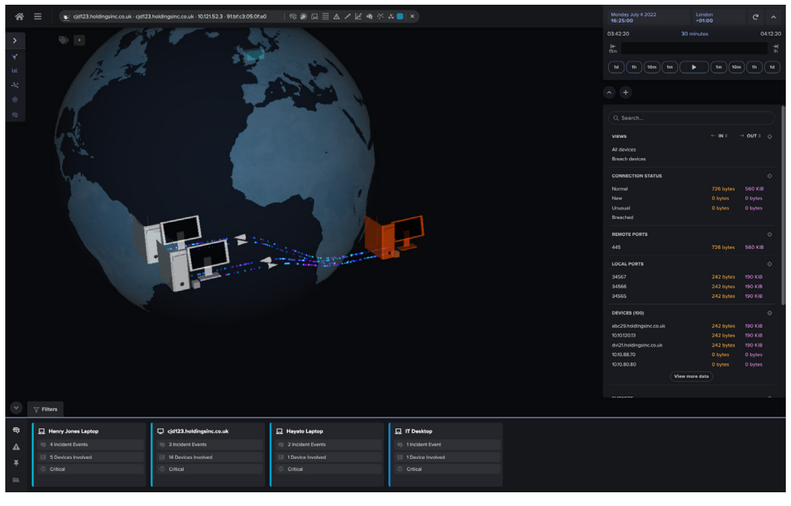

Figure 2: The Darktrace Threat Visualizer

/ Autonomous Response

Darktrace has also delivered the world's first proven Autonomous Response technology on the market. 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 Darktrace's rich understanding of the entire digital business, Darktrace 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, Darktrace RESPOND 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.

> "Darktrace is the cornerstone of our cyber security strategy. The rules-based approach we had before Darktrace could only detect known threats. We needed something that could secure our network from insider threat and 'unknown unknowns.' Now Darktrace notifies us to threats as they emerge, while its autonomous response technology stops attacks in real time—defending our network 24/7."
>
> / CIO, Livingston County

Beyond this tactical protection, Darktrace RESPOND 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, Darktrace can seamlessly plug into and enhance your existing ecosystem, informing firewalls and network devices about attacks that have gotten through.

When Darktrace DETECT and RESPOND are working in tandem, even the most complex and vulnerable organizations are transformed into a resilient, self-defending digital businesses.

/ Fully Customizable

Darktrace RESPOND runs fully autonomously, or can be set to act within guiderails decided by the security team. It can, for example, be set to operate only at certain times, on certain devices, or in response to certain events. As organizations build trust in the autonomous decision-making, many switch to fully autonomous mode within weeks.

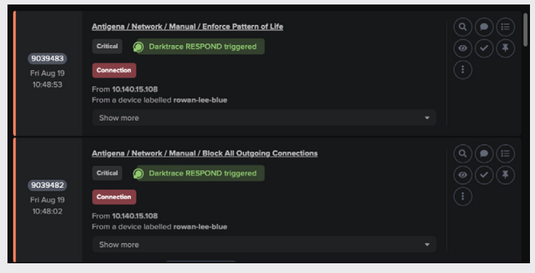

Figure 3: An example of Darktrace RESPOND actions

Cyber AI Analyst

While Darktrace DETECT and RESPOND 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 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 DETECT identifies, 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.

/ Use Case Example: Attack Exploiting Zero-Day Vulnerability

Cyber AI Analyst 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.

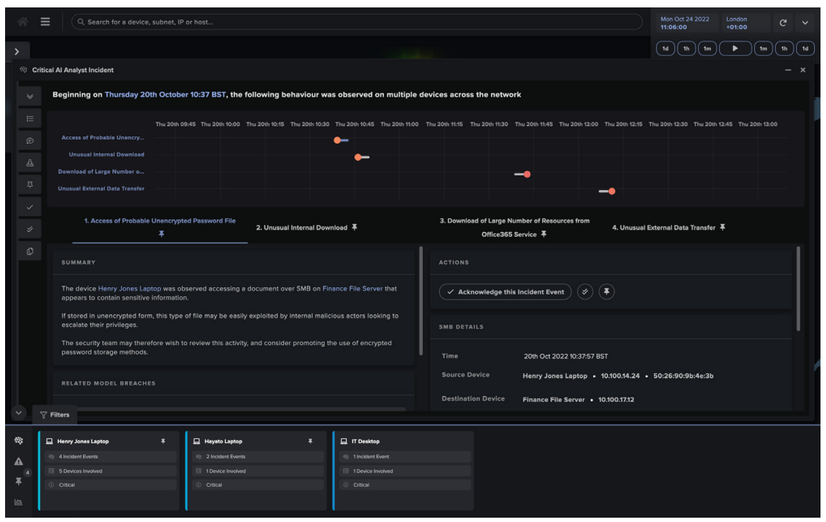

Figure 4: An incident summary generated by the Cyber AI Analyst

Protection Across the Enterprise

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 is specifically designed to cut across multiple stovepipes and enable unified detection and response, spanning across email, cloud, SaaS applications, industrial systems, endpoints, 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 understands that a user's normal patterns 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 AI identified attacks based on its unified under- standing of 'normal' across cloud, SaaS, email, industrial, and the corporate network.

Comprehensive Protection Wherever You Need It

Cloud | Apps | Email | Endpoint | Network | Zero Trust | OT

For every major partner and service including:

Google Cloud · Dropbox · slack · Windows · Microsoft 365 · box · aws · Microsoft Azure · Linux · zoom · salesforce

Darktrace/Cloud

The large-scale transition to cloud technology has reshaped digital businesses and traditional network paradigms.

While the cloud offers hybrid infrastructure and enables worker distribution, the increase in speed and efficiency is also a lure for attackers. These affordances of cloud technology thus introduce a new layer of complexity that most organizations are ill-equipped to address. However, Darktrace's AI technology is well prepared to detect and respond to any form of cyber threat in cloud environments.

Darktrace DETECT + RESPOND/Cloud work together to extract hundreds of metrics of raw data it receives from your cloud platforms and combines it with the overall context of the environment. While Darktrace's Autonomous AI can operate independently, IT professionals can select different modes of integration and design parameters for its decision making. Learn more about this in our discourse paper; "Managing Autonomous Decision Making."

Furthermore, Darktrace/Cloud grows with you, adapting to each new change in your business. Quickly configure and expand coverage of your environment with automation tools. Darktrace offers AWS Cloud formation Quickstart and Azure ARM Quickstart to provide you with autoscaling, and its osSensors are designed to scale easily as new instances are spun up.

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

Cybersecurity Insiders

/ 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 5: The sensitive files in Azure

Darktrace/Apps

Modern cyber-attacks traverse multiple fields of operation. Data from cloud applications is often just one piece of the puzzle, which is why Darktrace/Apps provides context derived from your entire digital ecosystem.

Darktrace/Apps learns every user's normal "pattern of life" to identify potential threats, based on millions of data points including login behavior, administration activity, file transfers, and more.

From unusual but benign activity to the more sophisticated cyber threats, Darktrace's AI will identify and respond to actions like increased access to data, compromised credentials, malicious insider behavior, and more by monitoring and locking out attackers, forcing logouts, and alerting security teams.

Darktrace Covers every major cloud application: Slack, Google, Salesforce, Dropbox, Box, and Zoom so that customers of all shapes and sizes –Enterprise, Government, Critical Infrastructure, and SMB, each with their own lineup of favorite & essential applications can feel secure. Furthermore, Darktrace/Apps integrates into your workflow including SIEMs, SOARs, and access via SSO.

Deployed passively and connecting to all apps in your digital ecosystem, Darktrace DETECT + RESPOND/Apps interacts directly with the SaaS vendor to understand activity within that cloud service and analyze every user. Using AI technology to contextualize user activity, the perfect counter response is developed for any conceivable threat.

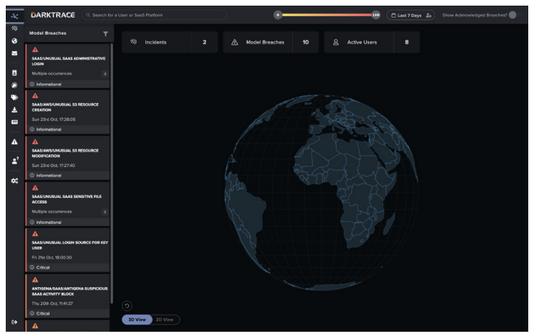

Figure 6: Darktrace's dedicated SaaS Console gives an overview of anomalous behavior in SaaS applications and displays the geographical locations of the activity

/ M365 Compromise and SharePoint Infiltration

At a US-based insurance company, Darktrace Self-Learning 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, Self-Learning 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.

The 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, Darktrace's Self-Learning 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.

/ 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 Darktrace's bespoke knowledge of 'self' for every member of the organization's workforce, the technology was able to identify the threat immediately. Darktrace DETECT revealed 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, Self-Learning 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.

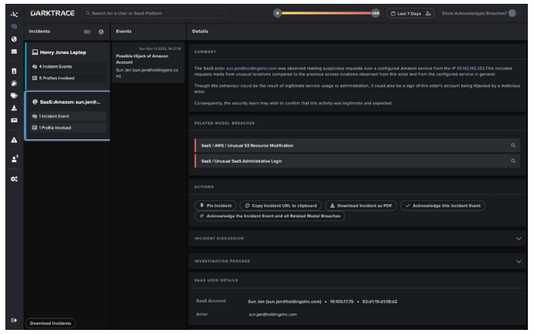

Figure 7: Darktrace showing the location of the unusual IP address

/ 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-crimi- nals 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 has been crucial in shining a light on account takeovers and other malicious activity across our cloud applications. This has been especially important to us in the era of remote and hybrid working patterns: having an extra layer of visibility across these applications gives us the confidence that we have all bases covered."
>
> / Head of Global Infrastructure, Network and Security, Boardriders

/ Compromise Across Microsoft 365 and Teams

Darktrace/Email alerted on numerous 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 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.

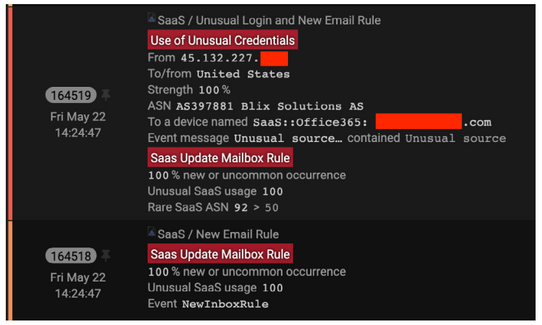

Figure 9: Darktrace noted a 100% rare IP logging into the user's Microsoft 365 account and the creation of a new mailbox rules

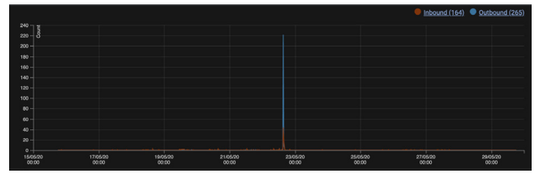

Figure 10: A recreation of the email sent by the attacker, containing the malicious attachment

The unusual login behavior detected by Darktrace/Apps could be connected to the anomalous outbound email behavior flagged by Darktrace/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.

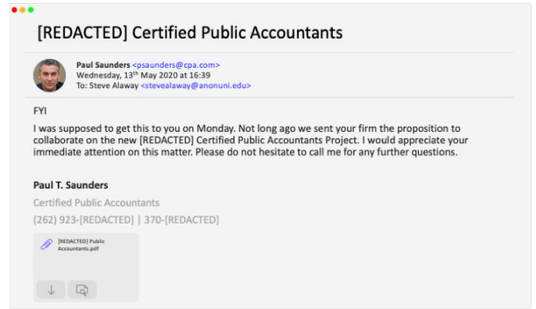

Figure 11: A recreation of the email sent by the attacker, containing the malicious attachment

Darktrace/Email

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.

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

Data Breach Investigations Report

Yet by analyzing the normal 'pattern of life' for every user and correspondent, Darktrace/Email understands the unique behaviors of the 'human' within email communications. Powered by Self-Learning 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.

Darktrace/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, Darktrace/Email I can spot subtle deviations from 'the norm' that reveal seemingly benign emails to be unmistakably malicious.

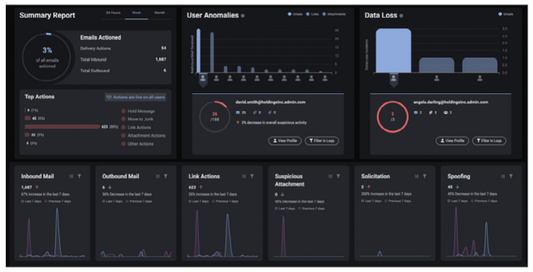

Figure 12: Darktrace/Email's interface displaying an overview of alerts

/ 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.

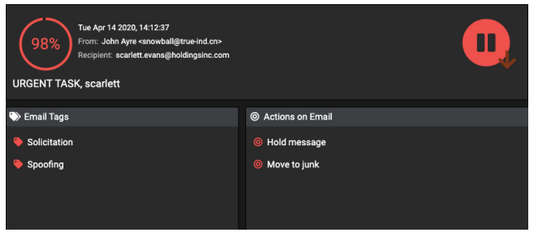

Figure 13: One of the 30 emails, with a 98% anomaly score

2. **No Association.** Across Darktrace's entire understanding of the company's email and network environment, Darktrace/RESPOND 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, Darktrace DETECT 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/Email identified a supply chain attack, recognizing that the sender was well known to the company, with several 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 Darktrace/Email held all 39 emails back and double- locked the payloads, based on the following anomalies:

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.

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.

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.

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.

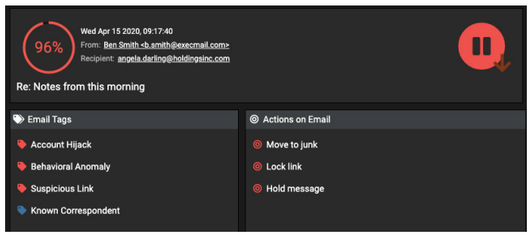

Figure 14: Darktrace detected the account takeover and held the emails back

Darktrace/Endpoint

As hybrid working practices continue to pull employees away from the relative safety of firewalled company networks, endpoints have come to represent an easy point-of-entry into organizations. They complexify the digital estate for security teams, but massively broaden the attack surface for attackers. Insider threat – both inadvertent and intentional – also becomes a bigger risk, with employees misusing company devices, and opening up avenues into the wider network.

Darktrace/Endpoint applies Self-Learning AI to every device in a company's digital estate in order to spot sophisticated and novel threats. Incorporating network, email, SaaS and cloud contexts, Darktrace/Endpoint investigates the full scope and pathway of every threat, and maintains the security of the entire digital estate by detecting attacks the moment they emerge.

With Darktrace/Endpoint, these attacks are not only detected in seconds, but neutralized with Darktrace's unique Autonomous Response technology. These actions are based on the AI's understanding of each user and device in the environment, consolidating hundreds of data points in order to stop threats without disrupting business. As a breached device is protected by this response, the device's user is able to continue with their normal business operations undisturbed.

The actions taken by Darktrace/Endpoint may include temporarily blocking anomalous connections, enforcing a 'pattern of life' or preventing anomalous data uploads and downloads. The speed of the AI's detection and response is such that most attacks will be stopped on the first device, before it can spread into the rest of the organization.

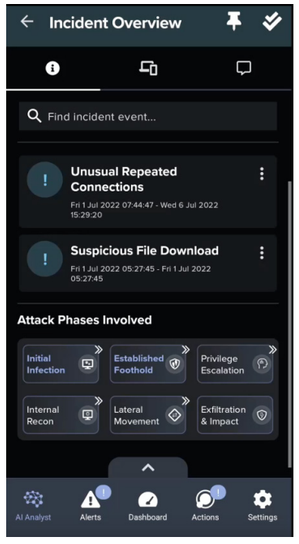

Figure 15: Notifications of Darktrace's findings and actions can be delivered to the Mobile App

/ Potential Government Exfil Event

In early 2022, at a governmental organization in North America, Darktrace/Endpoint detected a government laptop transferring data to a new, external endpoint with an unusual string. It is suspected that this string was created by a domain generation algorithm (DGA), which can be used by malware in order to bypass rules-based security measures. The unusual data transfer has, therefore, been considered a likely malicious exfiltration attempt. As this was a governmental organization, it held lots of highly sensitive data, the loss of which could have had serious consequences.

However, Darktrace/Endpoint immediately identified the anomalous behaviour as a potential attack. Because Darktrace's detections are not based on pre-defined attack data, the novelty of the DGA-generated string didn't prevent it from spotting the threat. In this case, Darktrace/Endpoint had been set up in human confirmation mode, and was therefore unable to take action to autonomously stop the attack, but it indicated that it would have blocked the data transfer on the relevant port for an hour, halting the exfiltration attempt.

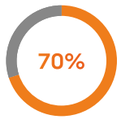

Up to 70% of successful breaches originate on the endpoint, according to an IDC report.

/ External Configuration Request Indicates Possible Insider Threat

At a UK technology company, Darktrace/Endpoint detected an internal device requesting a wpad[.]dat file from a location which was both rare and external to the organization. Wpad is a protocol used to discover the location of configuration files, and is therefore normally only found in internal communications. With an external destination, this protocol could have returned malicious configuration settings. If, for instance, the device had been instructed to route traffic through a threat actor's infrastructure, the whole organization could have been opened up to severe attacks.

The cause of this anomalous behaviour may have been a misconfiguration, a pre-existing compromise, or even a more malicious insider threat. To ensure the safety of the organization, Darktrace/Endpoint blocked the connection for one hour via the relevant port, allowing the behaviour to be verified before any changes could be implemented within the wider digital environment.

Darktrace/OT

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.

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

Ponemon

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. Organizations operating critical infrastructure must often comply with legislation like the recent US Cyber Incident Reporting for Critical Infrastructure Act, requiring prompt cyber incident reporting. AI-generated natural language summaries accelerate this process, making it considerably easier for organizations to hit government deadlines.

The Industrial Internet of Things (IIoT) creates further risks. In most cases, security teams can't run standard anti-virus software on these devices given a lack of disk space, a CPU, or a traditional operating system, and most organizations lack fundamental visibility over their IP-based IoT devices. With Self-Learning AI, organizations can monitor 100% of their devices, wherever they are on the network, while Darktrace AI spots the full range of attacks targeting IIoT devices.

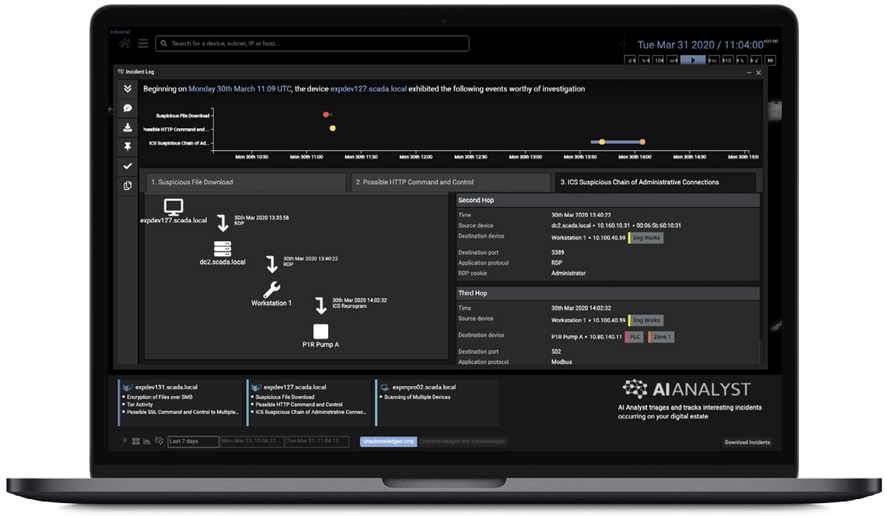

Figure 16: The Cyber AI Analyst surfacing all remote desktop 'hops' in a Triton-style attack targeting IT and OT

/ 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's Self-Learning 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.

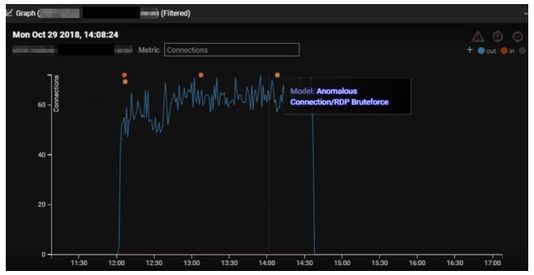

Figure 17: A clear plateau in increased internal connections can be seen. Every colored dot on top represents an RDP brute force detection

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.

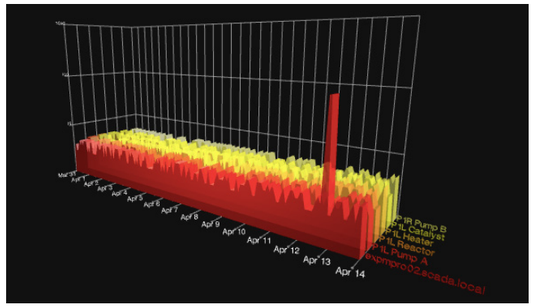

Figure 18: The anomalous connections from the SCADA device are clearly visible

Darktrace/Network

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, endpoint, 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.

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

Capgemini Research Institute

In turn, Darktrace/Network 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 platform'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.

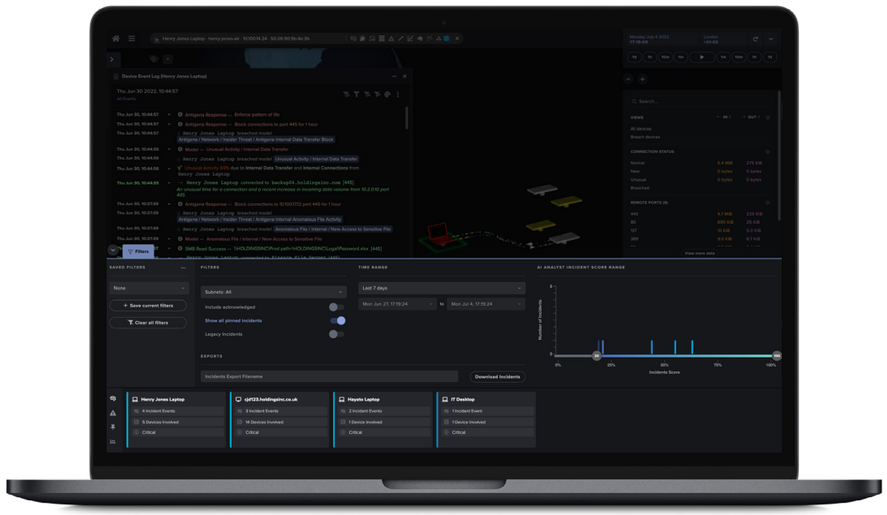

Figure 19: Customizable filters within the Darktrace user interface

/ 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.

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

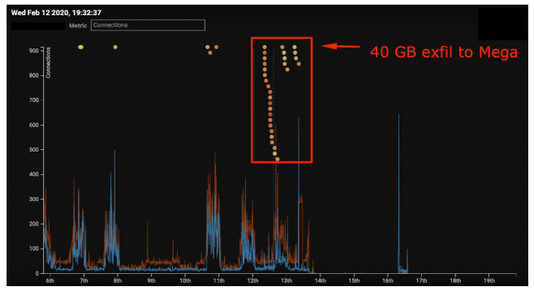

Figure 20: Connections to the domain controller

/ 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.

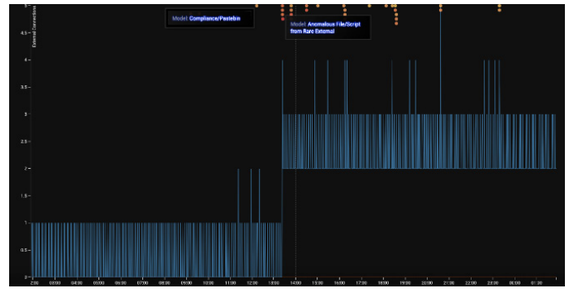

Figure 21: Graphical representation of the sudden increase in external connections and related model breaches

About Darktrace

Darktrace (DARK.L), a global leader in cyber security artificial intelligence, delivers complete AI-powered solutions in its mission to free the world of cyber disruption. Its technology continuously learns and updates its knowledge of 'you' for an organization and applies that understanding to achieve an optimal state of cyber security. It is delivering the first ever Cyber AI Loop, fueling a continuous end-to-end security capability that can autonomously prevent, detect, and respond to novel, in-progress threats in real time. Darktrace employs over 2,200 people around the world and protects over 8,100 organizations globally from advanced cyber-threats.

Scan to LEARN MORE

DARKTRACE
Evolving threats call for evolved thinking

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