

Darktrace Cyber AI Analyst

Augmenting Your Security Team with AI-Driven Investigations



Introduction

Contents

Cyber AI Analyst	2
Augmenting the Security Team with AI: Technical Overview	4
The Journey to AI-Powered Investigation	5
Applying Human Expertise at Machine-Speed and at Scale	7
Tracing Chinese Nation-State Actor APT41	8
Sensitive Sharepoint Files Accessed Via M365	8
Boosting Productivity and Saving Critical Time	9
Multi-Stage Network Vulnerability Attack	10
SaaS Account and Documents Compromised	10
Jumpstarting Remediation with AI-Driven Insights	11
Conficker Virus Evades Legacy Tools	12
Conclusion	13

The myriad of security tools used by businesses today creates massive quantities of data and surfaces too many alerts for analysts to effectively manage. As threats become increasingly sophisticated and the cyber security industry continues to face a skills shortage, over-worked and under-resourced teams urgently need augmentation.

Cyber AI Analyst, the product of a research initiative from Darktrace's R&D Center in Cambridge, was built to augment security teams and optimize threat investigation. It continuously examines every event that arises in Darktrace's Enterprise Immune System, emulating expert human thought processes for autonomous triaging and reporting.

The technology combines expert analyst intuition with the consistency, speed, and scalability of AI. It illuminates the highest priority threats at any one time and rapidly synthesizes all of the context around an attack into a human-readable report.

By applying a combination of supervised and unsupervised machine learning, as well as deep learning methods and advanced mathematics, Cyber AI Analyst can do much of the heavy lifting a human would otherwise have to do. It leverages insights collected from Darktrace's world-class experts over years of threat investigation to make highly accurate decisions and offers this wealth of knowledge to the public for the first time.

With Cyber AI Analyst, time-to-meaning and time-to-response are dramatically reduced – allowing your team the time to use their expertise where it really matters.

Cyber AI Analyst

Key Benefits:

- ✓ Automatically investigates every security event detected by the Enterprise Immune System, 24/7
- ✓ Highlights the most critical issues at any one time for advanced incident prioritization
- ✓ Pulls together related events and behaviors into an Incident Report that can be read in minutes and actioned even by non-technical users
- ✓ Reduces triage time by up to 92%, buying back time so teams can focus on strategic work

A key feature of the immune system approach and enabled by artificial intelligence, Cyber AI Analyst can sift through large volumes of data at a speed and scale, augmenting human teams and buying back time to focus on strategic work.

Its investigations are enterprise-wide, allowing the technology to piece together disparate anomalies before settling on a high-level conclusion about the nature, root cause, and extent of the wider security incident. This powerful analytical capability has been found to reduce triage time by up to 92%.

Cyber AI Analyst applies various forms of artificial intelligence, including deep learning, as well as supervised machine learning on the ever-growing data set that captures how Darktrace's human analysts investigate threats.

With this knowledge, Cyber AI Analyst is able to understand which threats are most crucial for investigation, which events constitute a connected incident, and how an attack should be managed.

This is a powerful capability considering that the Enterprise Immune System often surfaces advanced and novel threats that legacy tools cannot identify. Cyber AI Analyst delivers expert analysis of all types of cyber-threats, even those characterized by innovative attack techniques that would be impossible to detect and respond to with pre-defined playbooks.




Figure 1: Cyber AI Analyst is a key technology grounding Darktrace's Cyber AI Platform.

The technology can communicate valuable contextual information and response recommendations in the form of a concise, meaningful narrative that security experts can then apply their own insight to.

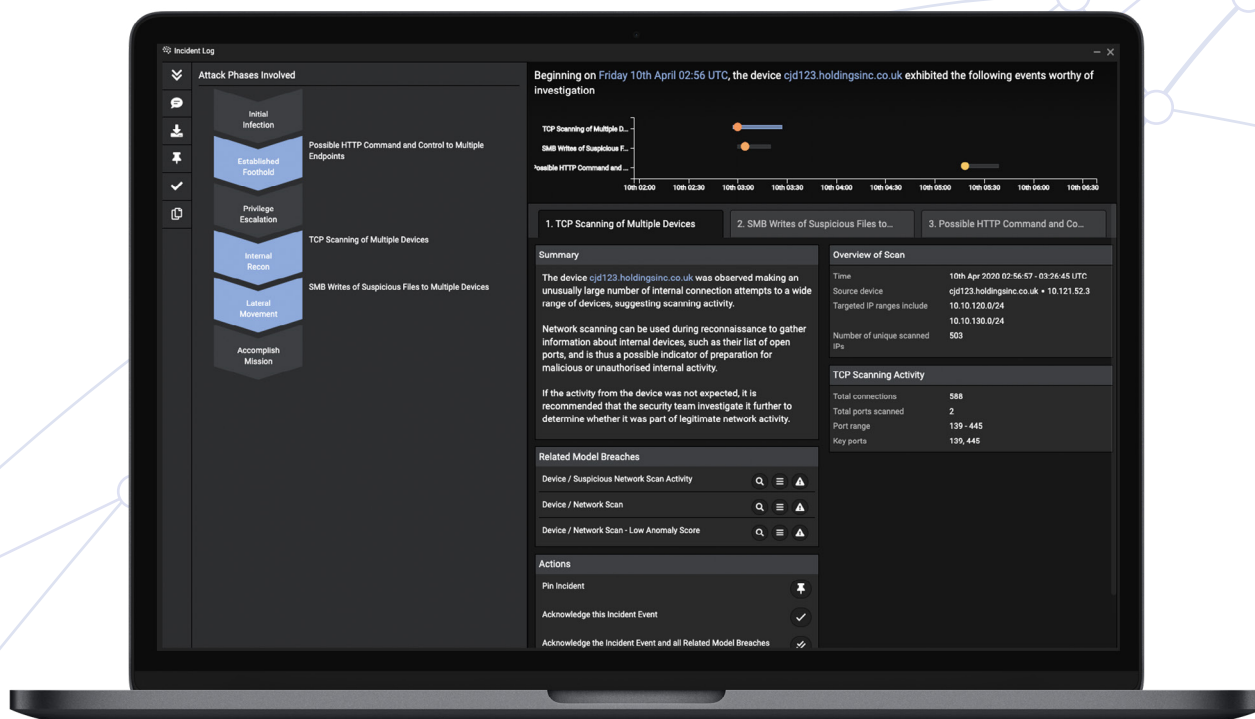
These “Incident Reports” are easily actioned by technical staff and executives alike and can be translated into any language at the click of a button. The reports take an average of three minutes to read, allowing even a non-technical responder to review and remediate sophisticated attacks in minutes.

“With Cyber AI Analyst we can see the whole picture. It’s all right there in one place, with all the context and details we’d need to take action.”

– Rick Bertoncin, Director of Technology & Security,
Gallagher-Kaiser Corporation

While Incident Reports are always created for the most critical threats at any one time, investigations can be applied on-demand to any event of interest.

Cyber AI Analyst technology can also be integrated with tools across your security stack, allowing investigations to be triggered based on data from third-party sources like CrowdStrike or Carbon Black. The rich context and insights of Incident Reports can additionally be exported to SIEM, SOAR, or ticketing systems to enhance your existing workflows.



Augmenting the Security Team with AI: Technical Overview

Cyber AI Analyst begins its analysis as soon as the Enterprise Immune System detects a Model Breach or pattern of anomalous activity, which serves as the 'lead' of the investigation.

Just like a human analyst, Cyber AI Analyst starts with this lead and then asks questions to generate a plausible hypothesis about the nature of the potential threat and the potential underlying cause.

It then pivots across the enterprise to query data that may confirm, deny, or refine its hypothesis. This process is repeated continuously until Cyber AI Analyst settles on a high-level description of the nature and root cause of the wider security incident.

Cyber AI Analyst's supervised machine learning does not use historical attack data, but rather learns on a dataset of human analyst behavior that has been accumulating over the past four years. The AI monitors Darktrace's own experts' behavior as they investigate threats from across our customer base, mining every action for implicit knowledge on how analysts triage threatening and suspicious activity.

“By automatically investigating security events, the AI Analyst helps reduce noise more than any other technology.”

— Chris Kissel, Research Director, IDC

Crucially, Cyber AI Analyst is able to adapt to new and unique situations on the fly, automating thoughtful examinations rather than pre-defined playbooks or encoded human knowledge.

Incident Reports can be downloaded in PDF format to be shared with relevant shareholders, enabling easy reporting for compliance or management requirements.

Generally, an Incident Report includes:

- A high-level narrative summary of the incident
- A detailed timeline pulling together all events related to the incident
- A list of related breaches and devices
- Attack phases involved in the incident (e.g. initial infection, established foothold, privilege escalation)
- Details around connections, endpoints, files, beaconing activity, and other relevant data

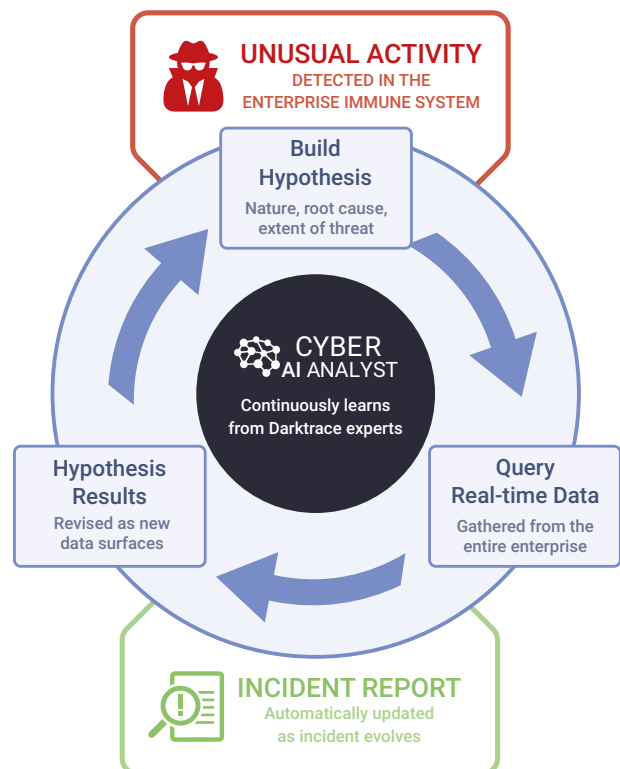


Figure 2: Cyber AI Analyst continuously builds and tests hypotheses, reasoning to conclusions at machine speed and scale.

The Journey to AI-Powered Investigation

Darktrace has always been the leader in AI-powered cyber defense. Where legacy security approaches rely on rules and signature-based detection, Darktrace invented self-learning Cyber AI and Autonomous Response technology to revolutionize cyber defense.

With Cyber AI Analyst, our R&D team has taken another major step forward, building AI that has learnt from our expert analysts in order to automate cyber-threat investigation and triage processes.

“

Cyber AI Analyst is a game changer because it takes people from watching a monitor to really starting to work through the trade craft, and reduces the time it takes to triage issues.”

– Laura Tibodeau, CIO, Americas Styrenics

Detecting Threats by Searching for “Bad”

The legacy approach to cyber defense focuses on programming security tools to search for threats using the most comprehensive definition of “bad” possible. This approach relies on knowing what previous threats have looked like by using known lists of Indicators of Compromise or other cyber-threat Techniques, Tools & Procedures.

While the rule and signature-based approach can be effective for the most clear-cut threats, it cannot detect incidents that are not identical to previously seen attacks.

Despite its shortcomings, the legacy approach does have one advantage: once a threat is identified, it can be quickly labeled, and threat intelligence can be easily searched. Yet any threat that varies from what has been seen before will inevitably slip past these traditional defenses.

One method for enhancing this legacy approach uses supervised machine learning to learn a greater volume and slightly more complex set of “bad” behaviors, allowing for stronger threat intelligence applications and often more accurate threat searches.

However, this method depends on training machine learning algorithms on retrospective data sets of known threats. While these tools may be more accurate than simpler legacy tools, they possess the same disadvantage: they are limited to searching for historical threats, and threats that are novel or highly targeted can still get through.

Detecting Threats by Learning Normal

A crucial pivot from the legacy approach has emerged, a paradigm shift that acknowledges that threats will inevitably evade those tools limited to historical knowledge.

The only way to thwart all cyber-threat is to learn what normal looks like for each unique digital environment, and to use that knowledge to identify the subtle behavioral changes that indicate compromise.

Some legacy solutions have responded to this shift by implementing basic anomaly detection. These tools may use some form of supervised machine learning to understand patterns for specific categories of devices and users, but they still rely on retrospective data sets, leaving security teams struggling to manage many false positives and allowing sophisticated attacks to go undetected.

Darktrace's immune system approach, which identifies threats by using unsupervised machine learning to understand normal behavior in real-time, is entirely distinct. Darktrace's Cyber AI builds a bespoke knowledge of organizations and their workforces, continuously analyzing unique "patterns of life" across digital infrastructure in order to identify behaviors that do not belong in a specific context.

This approach allows even the most advanced and novel attacks to be detected and neutralized, and shows that a technology doesn't need to know exactly what it's found to identify a threat.

Now, Darktrace has found that Cyber AI can be applied not only for advanced threat detection, but also for augmenting the investigation process – until now something that human analysts have been solely responsible for.

Applying AI to Threat Investigation

In discussions around humans and artificial intelligence, we always find that humans have superior decision-making skills and strategic impulses based on our incredibly nuanced understanding of various contexts.

However, the neural network "brain" formed by artificial intelligence algorithms does have some advantages over the human mind – namely, that it can intake massive amounts of data and work at machine-speed.

Darktrace's Cyber AI Analyst combines the power of machine learning with expert human analyst intuition to create a technology that can make the leap from advanced threat detection, to automated threat investigation.

Human analysts investigate threats by finding patterns, forming hypotheses, reaching conclusions, and sharing their findings with the rest of the business. These are labor-intensive steps that require expertise and extensive time commitments.

Yet by learning how expert humans investigate threats, Cyber AI Analyst is able to form hypotheses and conclusions at a speed and scale that humans cannot replicate. While supervised learning is not a sufficient approach for threat detection, it is a gamechanger when it comes to understanding how analysts triage and respond to threats that have been surfaced.

With Cyber AI Analyst, the initial investigation steps can be conducted at machine-speed across the entire enterprise and provides AI-powered insights, serving as a critical force-multiplier, and jumpstarting incident remediation.

Applying Human Expertise at Machine-Speed and at Scale

“

Darktrace's Cyber AI Analyst can automatically tell you which events are linked and should be treated as a single incident rather than three or four separate events. That's completely unique.”

– Chris Sprague, Security Engineer,
TruWest Credit Union

Cyber AI Analyst does not try to fit events into any predefined pattern, nor does it simply look at how previous, similar event types have been categorized and responded to. Instead, it analyzes threats based on every connected detail, and considers how humans have responded in similar contexts.

The technology can replicate the decision-making of thousands of world-class analysts with the machine power of Cyber AI, meaning it can look at incredible amounts of data and correlate patterns across all parts of the digital environment.

In understanding the context as a human expert would, Cyber AI Analyst understands just how serious a threat may be and connect the dots between the related security events that make up a single threat incident.

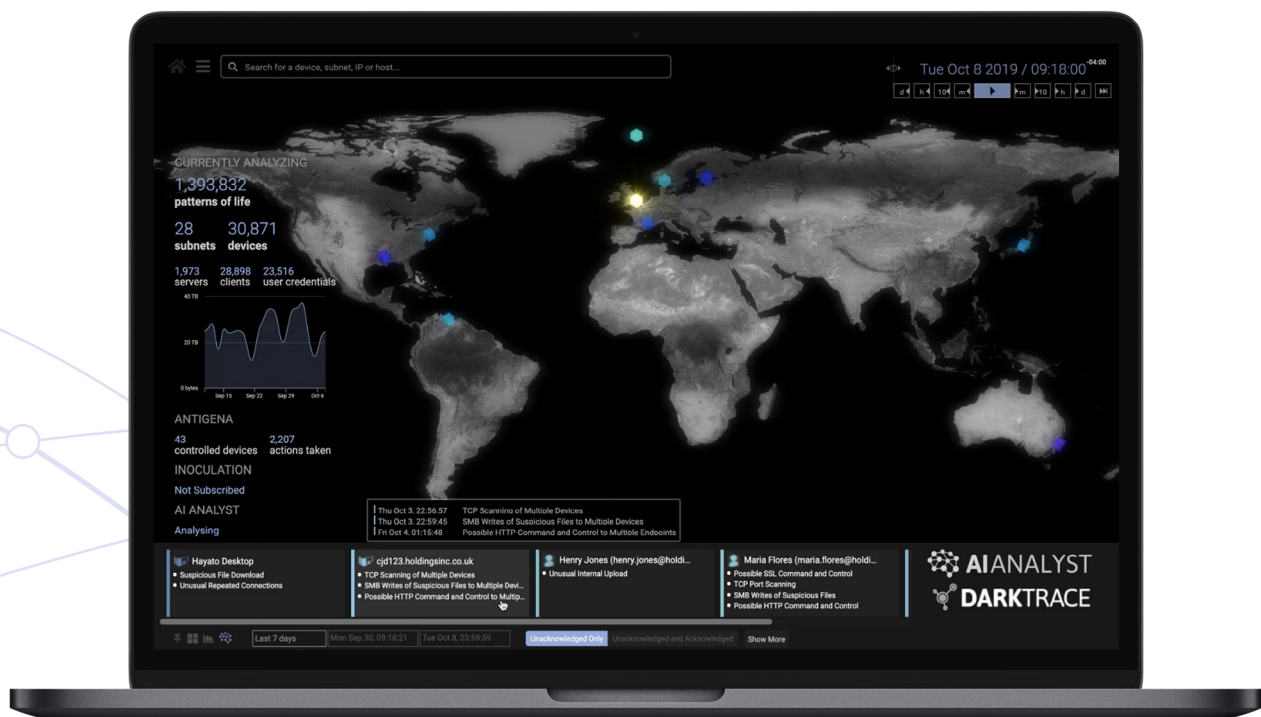


Figure 3: Cyber AI Analyst highlights the most significant incidents at any time, offering a cohesive narrative for each one.

Tracing Chinese Nation-State Actor APT41

When the Enterprise Immune System detected highly targeted attacks on customers in Europe and the US, Cyber AI Analyst was able to piece together the narrative and understand the severity of the threat.

Two weeks later, the attacks were attributed to a Chinese state-sponsored threat actor APT41. Yet before attribution and any threat intelligence was available, Cyber AI Analyst tied together the APT41 behavior as it emerged and highlighted the extent of the attack with striking accuracy and granularity.

The attacks targeted public-facing servers and exploited recent high-impact vulnerabilities. In one representative customer environment, Cyber AI Analyst first detected the initial intrusion, in which the APT41 actor exploited the Zoho ManageEngine zero-day vulnerability CVE-2020-10189. Then, the technology saw the Microsoft BITSAdmin command line tool used to fetch and install a malicious Batch file.

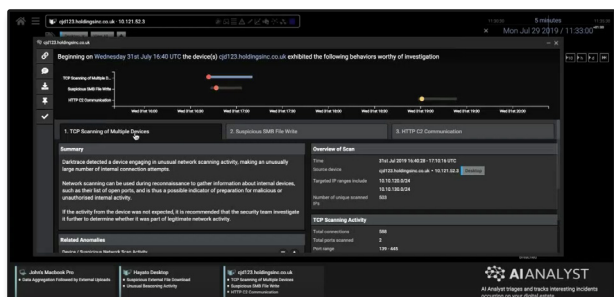


Figure 5: Cyber AI Analyst provides an attack summary in this sample UI image.

Teamviewer activity and the download of Notepad++ was also observed at the same time as command and control traffic was starting, indicating that APT41 may have been trying to use familiar tools instead of completely living-off-the-land.

Cyber AI Analyst's Incident Reports, grounded as they are in self-learning detection and expert analyst intuition, highlighted and delineated every aspect of this incident in the form of a meaningful security narrative.

Even a non-technical responder could have reviewed the Incident Report and remediated this zero-day APT41 attack in minutes.

Sensitive Sharepoint Files Accessed Via M365

When a threat actor infiltrated an employee's Microsoft 365 account to access sensitive financial documents hosted in SharePoint, Cyber AI Analyst immediately connected the dots between anomalous activities to paint a comprehensive narrative of the attack.

After comprising the employee's business email, the attacker sought out important materials in the user's SharePoint account, including pay slip and banking details. They then made configuration changes to the hacked inbox, deleting items and making updates that may have allowed them to cover their tracks.

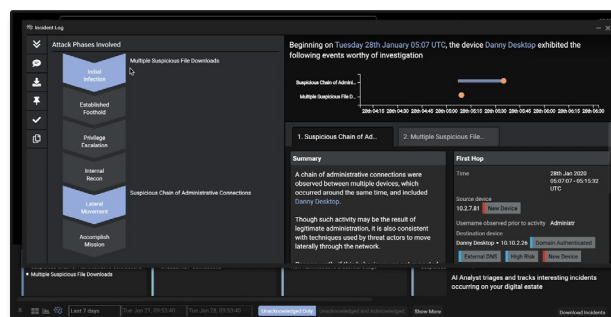


Figure 4: Sample UI image of Cyber AI Analyst highlighting attack phases, a timeline of events, and the incident's source.

Darktrace's Cyber AI initially observed the actor logging in from a range of unusual IP addresses that this user had never logged in from before, nor had anyone in his peer group at the organization. This in itself was not necessarily indicative of malicious activity – after all, employees often change locations.

However, the unusual login location was accompanied by an unusual login time and a new User-Agent. As it intelligently connected all of these anomalies, Cyber AI Analyst decided to trigger a deeper investigation.

The technology used its ability to build meaningful and concise attack narratives to illuminate the likely account compromise as a high-priority incident, and compile relevant data into a human-readable report.

Boosting Productivity and Saving Critical Time

“

It will be crucial to use AI as a force multiplier and keep pace with attacker enhancements as they upscale their attack capabilities and efforts.”

– Forrester Research

Instead of wasting valuable security staff time on manually sifting through immense amounts of data and alerts, Cyber AI Analyst lets you automate the initial investigation, applying the insight of Darktrace's expert human analysts to allow your team to focus on strategic decision-making.

Cyber AI Analyst ensures that the most significant threats are always brought to attention first, helping prioritize limited security resources and mitigating business impact.

In a recent report with Forrester, Darktrace found that 83% of cyber security decision-makers believe that organizations require AI augmentation that can adapt to shifting threat vectors, as well as reduce time-to-detection and time-to-meaning in understanding the full scope of an incident.

Cyber AI Analyst responds to this need perfectly, slashing the time-to-response and time-to-meaning when a threat is detected, and reducing triage time by up to 92%.

The technology also helps reduce the rate of false positives, mitigating alert fatigue for already over-worked teams and making their jobs more fulfilling.

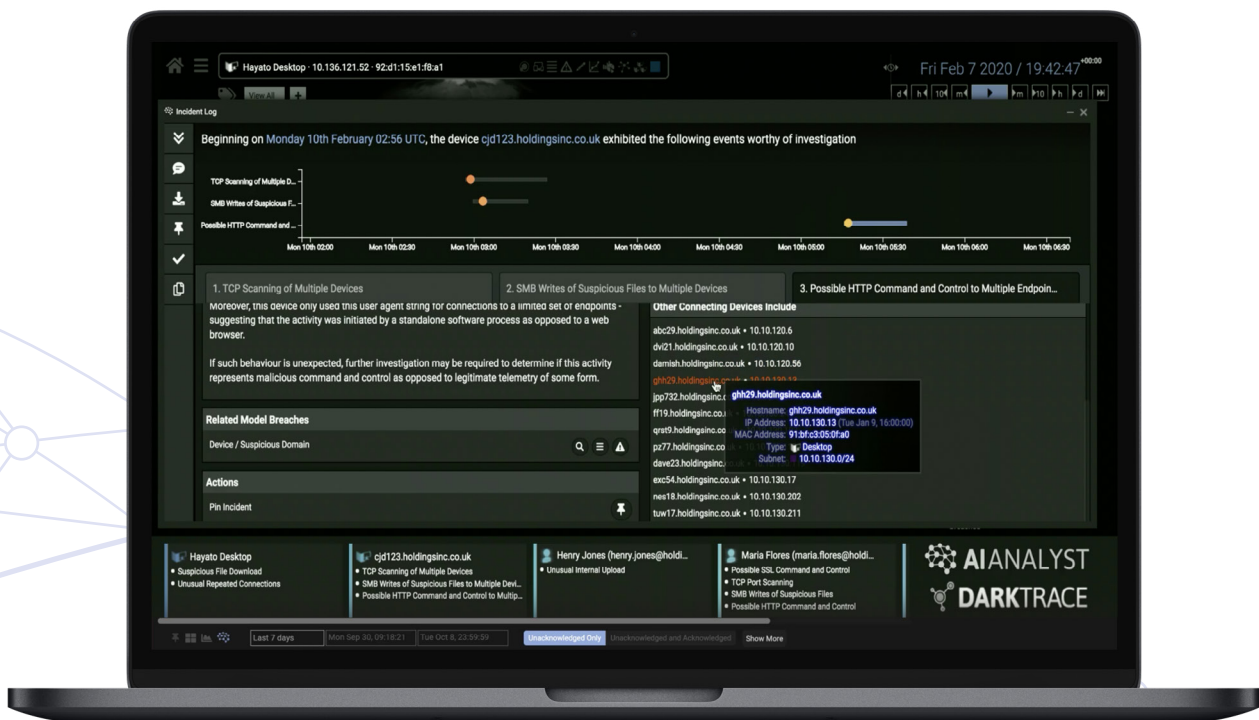


Figure 6: Cyber AI Analyst offers a high-level incident timeline and summary, as well as granular details.

Multi-Stage Network Vulnerability Attack

When a threat-actor took advantage of vulnerabilities on multiple devices at a US pharmaceuticals company, Cyber AI Analyst proved crucial in illuminating every step the attacker took and responding to the threat.

Darktrace's Cyber AI first saw a company desktop making a new connection to a second internal device using the Remote Desktop communication protocol and administrative credentials. While this connection was ongoing, the second device made a connection to an internal server using the SMB file sharing protocol. With this connection, the second device wrote to a hidden file share on the internal server, an activity that had never been seen before in this context.

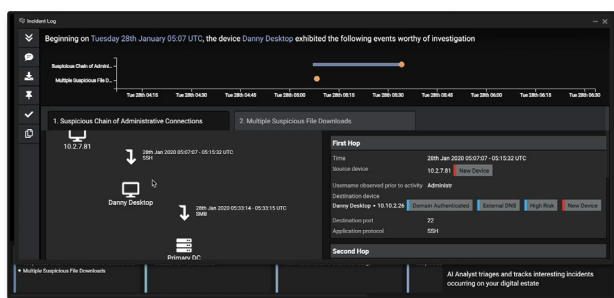


Figure 7: Cyber AI Analyst visually represents a multi-stage attack moving through devices in this sample UI image.

While the initial new connection was not necessarily suspicious, Cyber AI Analyst was able to correlate all of this activity to understand how each subtle anomaly was part of the wider threat narrative.

In understanding that the activity may have indicated an attacker working their way through vulnerable network devices to gain access to critical infrastructure and resources, the technology was able to alert the security team to the emerging threat immediately and provide crucial contextual information, along with response recommendations.

With Cyber AI Analyst augmenting the initial investigation steps, the security team was easily able to ensure that a serious, multi-stage incident was neutralized as soon as possible.

SaaS Account and Documents Compromised

At a global produce supplier, several suspicious requests within the company's Box platform suggested that an account had been compromised. A user logged in to the Box account from a country that the organization typically operates in, but from an unusual IP space and ASN.

With this suspicious and unusual detail, Cyber AI Analyst automatically actioned further investigation into the user's activity.

The actor behind the account proceeded to download expense reports, invoices, and other financial documents. It soon became evident that the user was reading files that were highly unusual for the account to access. Darktrace Cyber AI recognized that neither the account itself, nor its peer group, usually accessed the file called 'PASSWORD SHEET.xlsx'.

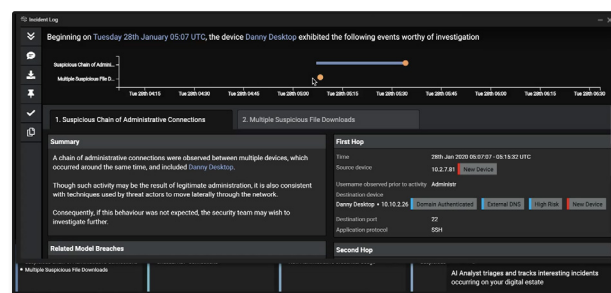


Figure 8: Sample UI image shows Cyber AI Analyst identifying unusual administrative activity and file downloads.

Cyber AI Analyst's automated investigation was able to shine a light on the greater narrative, understanding that each unauthorized file exposure was part of a connected incident. The technology highlighted the breach as a key concern for the security team, and summarized all related behavior in an actionable Incident Report.

With Cyber AI Analyst illuminating the initial, unusual behavior as part of a more significant threat and making the relevant contextual data easily accessible, the security team had the power to respond more quickly and efficiently to an attack that might have otherwise gone unnoticed.

Jumpstarting Remediation with AI-Driven Insights

“

With a small IT team, we need an AI Analyst to help us figure out the root cause of an incident and make it easier to respond to an infection.”

– Bobby Garrett, IT Director, Gray, Gray & Gray

Cyber AI Analyst was designed to help teams jumpstart remediation with its concise Incident Reports. These reports bring the deep visibility and comprehensive situational awareness of Cyber AI into a cohesive narrative, which is continuously revised as incidents evolve.

New activity and data prompt Cyber AI Analyst to automate further investigations and test new hypotheses, so that your team always has the most up-to-date insights grounding their response decisions.

The reports can be used to understand the root cause of an incident, the extent of an attack, and may even reveal unexpected infrastructural vulnerabilities. In this way, the technology can be leveraged to evolve your security posture.

Cyber AI Analyst technology also makes reporting a simple, low-friction process – even when it's a complex incident that needs to be shared with executives, board members, or even legal authorities.

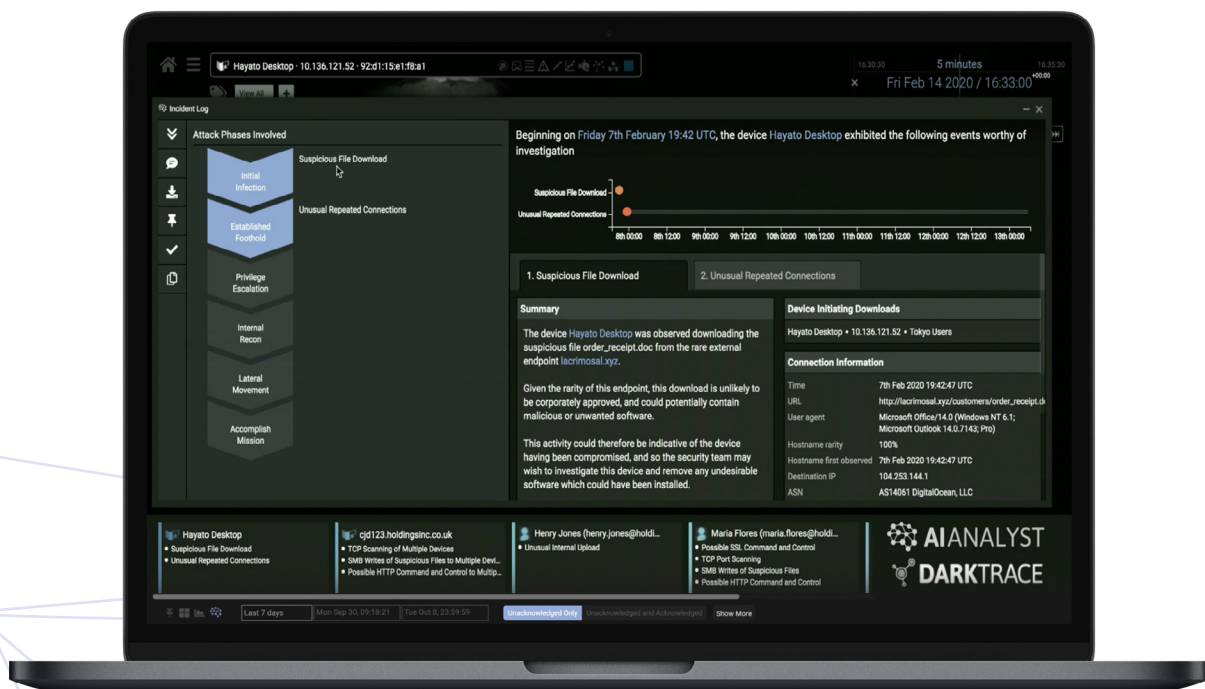


Figure 9: Cyber AI Analyst clearly presents the attack status and recommendations for next steps.

Conficker Virus Evades Legacy Tools

When Darktrace's Cyber AI saw a device making large volumes of DNS requests for domains that appeared to be algorithmically generated, Cyber AI Analyst immediately understood the gravity of the situation and presented an Incident Report to the security team.

Based on Cyber AI Analyst's clear narrative of the event, the human analysts were able to see that the device was likely infected with the Conficker worming virus.

Subsequent investigation revealed that this was likely to be a highly targeted attack, as the company's antivirus provider had rules configured to detect Conficker – yet this device appeared to have been configured incorrectly and was therefore able to be penetrated.

Without Cyber AI Analyst highlighting the root cause of this incident, the security team likely would not have known that the device was lacking adequate security configuration, and the device in question certainly could have been seriously compromised and damaged. However, with Cyber AI Analyst's ability to illuminate the context and extent of the incident, the security team was instead able to contain the device and mitigate further business impacts.

Cyber AI Analyst's investigation and report empowered the security team to quickly engage in comprehensive remediation efforts, with a new understanding of their existing security coverage and potential vulnerabilities.





Figure 10: In this sample image, Cyber AI Analyst offers details on a breached device and lists similar devices – giving security teams critical context for remediation.



83% of cyber security leaders agree that firms require AI augmentation that can reduce time-to-detection and time-to-meaning.

Source: Forrester Research

Conclusion

Amid the flood of information and alerts that comes with today's security technologies, Cyber AI Analyst gives teams the power to see what really matters and to respond as effectively as possible.

Cyber AI Analyst answers the key questions that you first ask when an incident arises: from identifying patient zero and the infection source, to showing which devices or files were affected, and uncovering C2 traffic or lateral movement.

The technology automates thoughtful and intelligent actions that no playbook can rival, expertly dissecting novel and highly sophisticated threats that the Enterprise Immune System is uniquely capable of detecting.

When a significant incident occurs, Cyber AI Analyst is always there – stitching together disparate security events, communicating the full scope of the incident, and weaving together the most important information you need to share with any relevant stakeholders.

Cyber AI Analyst's ability to replicate human thought processes at speed and scale will add even further value as threat landscapes evolve and attackers continue developing automated, machine-speed campaigns.

Increasingly fluid workforces, partner networks, and supply chains, as well as new technologies and services, will make it imperative to leverage innovations like Cyber AI Analyst that bring clarity to the complexity of digital life and defend against all types of threats.

By applying advanced machine learning to the data-heavy investigative tasks that AI can do best, Cyber AI Analyst gives you time for the critical decision-making that only humans can do best.

Key Takeaways

- Minimizes time-to-meaning and time-to-response, augmenting overworked and understaffed teams
- Replicates meaningful human investigations at scale and machine-speed - only possible with AI
- Understands what's important to you amid the complexity of digital environments
- Automatically investigates, 24/7 – highlighting which incidents to prioritize and all the relevant information in a human-readable, actionable format
- Transfers years of insights from Darktrace's expert analysts to your unique environment

Discover how Darktrace can protect your dynamic workforce.

[Sign up for a free trial here](#)

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

[@darktrace](#)