

Autonomous Response: Threat Report 2019

Introduction

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 our digital businesses continue to grow in complexity, diversity, and scale.

In the past, when threat actors were less advanced and when networks were more predictable, a traditional approach to security was often adequate to keep cyber-threats at bay. By configuring security tools with some combination of rules or signatures, security teams 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 and subtlety of daily behaviors in a business, have gradually disarmed security teams who still rely on traditional controls. Traditional 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.

The fact is that novel threats will inevitably get inside, and so the industry’s attention has shifted to the question of how cyber 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.

While traditional approaches continue to pre-define the threat in advance, Darktrace’s unique application of AI focuses instead on learning the normal ‘pattern of life’ for individual businesses, and spotting subtle deviations indicative of a threat – whether known or unknown, external or internal, subtle or fast-moving. By learning ‘on the job’ and continuously adapting in light of new evidence, Darktrace’s artificial intelligence spots early indicators of cyber-threat that would otherwise go unnoticed, without relying on rules, signatures, or prior assumptions.

Summary

This report details seven case studies of attacks that were intercepted and neutralized by cyber defense AI, including insider threat, ransomware, and IoT attacks.

While all threat scenarios were distinct, some fast-moving and others slow and stealthy, in all cases the subtle indicators of suspicious activity were only detectable using Darktrace AI, which learns what is normal for the business environment and autonomously responds to attacks – before damage is done.

Fight Back With Darktrace Antigena

Yet as the volume and speed of attacks continues to rise, Darktrace’s cyber AI has evolved to not only detect but also intelligently fight back against in-progress attacks in real time. While traditional approaches have historically included some measure of automation in incident response as well – from the old-school IPS of the 1990s to the ‘next-gen’ antivirus tools and email gateways of the present day – these pre-programmed response tools have invariably been blunt and disruptive. What they lack, in a word, is context, and an evolving understanding of the normal ‘pattern of life’ for every user, device, and associated peer group in a business.

Armed with this rich and evolving understanding of ‘normal’ for the first time, Darktrace’s AI can not only respond to early indicators of cyber-threat before they do damage, but also do so in a highly targeted fashion. Rather than generating broad-brushed quarantines that would only serve to cause more disruption, Darktrace Antigena – the system’s autonomous response solution – works by surgically enforcing the normal ‘pattern of life’ for an infected device or disaffected employee, neutralizing the threat within seconds and sustaining normal operations by design.

In the fight against advanced cyber-criminals, Darktrace’s cyber AI is finally giving control back to the defenders, transforming even the most complex and vulnerable organization into a resilient, self-defending digital business.

Insider Threat

An Insider Scanning the Network for Vulnerabilities

Malicious and Persistent

Insider threat represents one of the more dangerous and common attack vectors in the enterprise. These threats originate from disgruntled, careless, or compromised employees who abuse their access to internal systems in varying degrees of severity and malice. At the more insidious end of the spectrum, malicious insiders pose an especially significant threat to the business, as their privileged access and knowledge of the network allow them to undertake extended attack missions and quietly exfiltrate or manipulate critical data without triggering suspicion.

Darktrace's AI identified and neutralized one such malicious insider at a major investment firm in South Africa. The self-learning AI was able to contain a persistent threat as it moved through multiple stages of the attack chain, from reconnaissance to script writes and script execution. By learning 'on the job', Antigena adapted to the threat as it evolved and effectively contained it at each stage.

Suspicious Behavior

The reconnaissance stage began with a laptop 'pinging' hundreds of internal IP addresses to identify those which were active. It then swept the network for the names of responsive machines, and scanned them for open channels of communication. Darktrace's AI flagged the suspicious behavior as unusual network-scanning activity, and instantly prompted Antigena to take action. Based on its dynamic evaluation of the threat, Antigena decided to enforce the device's group 'pattern of life' for one hour, preventing the laptop from deviating from its prior behavior or that of its peers.

Yet a few hours later, the threat returned. The laptop started running commands on hundreds of other internal computers in the IP range it had initially identified. This involved moving multi-purpose script files, and using a remote-administration tool. These programs could be exploited to locate sensitive information and documents, or to open a backdoor for an external attacker to hijack.

Antigena decided to enforce the device's group 'pattern of life' for one hour

Antigena Steps In

No other similar file-writes were seen across the network during this period of time, which showed up as highly unusual to Darktrace's AI. Given its evolving understanding of the threat in the context of the network and its previous autonomous response, Antigena decided to block all outgoing connections using the SMB file-transfer channel, instantly containing any lateral movement across the network.

Once the threat had been neutralized, the security team was able to investigate and confirm that the laptop belonged to a member of the IT team who had been using an illegitimate scanning tool to look for weaknesses in the network. This was an especially revealing example of the power of Darktrace's AI, and how Antigena can step in at different phases of an attack chain and neutralize persistent threats at an early stage.

Zero-Day Trojan

Suspicious Download and Connections

Novel Strain of Malware

While legacy security tools can often identify known threats that have already been discovered 'in the wild', artificial intelligence can uniquely spot the weak and subtle signals of a never-before-seen cyber-threat. This capability has become necessary in recent years, as advanced cyber-criminals continue to develop novel tactics, techniques, and procedures specifically designed to evade controls that have been pre-programmed with signatures of past attacks.

Darktrace's ability to react to these subtle indicators was critical for an American manufacturer of industrial IoT controls, when it was hit by a zero-day trojan.

At 1:30pm on a Thursday, the AI alerted the company's IT Manager to a suspicious download of a file named 'OfficeActive.bin'. While the file looked like a Microsoft product, Darktrace indicated that the file was being downloaded from an unidentified source that was 100% rare for the network.

While the file looked like a Microsoft product, Darktrace indicated that the file was being downloaded from an unidentified source that was 100% rare for the network

Building Trust in AI Response

Antigena was configured in 'Passive Mode' at the time – a starter mode that restricts the AI to communicating what it would have done in response to the threat, without actually taking action – enabling the team to build trust in the system's decision-making. The IT team was able to see how Antigena would have stopped the attack at an early stage, and also how it adapted to a novel threat as it escalated.

In response to the highly unusual pattern of activity, Antigena first recommended enforcing the device's group 'pattern of life' for two hours, which would have stopped the threat in its tracks while sustaining normal operations.

As it observed more suspicious downloads, Antigena escalated its response, enforcing the device's individual 'pattern of life' for five minutes. And when the device attempted to make a new external connection, Antigena responded again, suggesting that the AI surgically block all outgoing connections from the device for one hour.

Remediating the Threat

Within minutes of identifying the alert, the IT Manager had contacted the end user and performed an emergency recompose to remediate the threat on the machine. The entire process was completed within 20 minutes. Once the threat had been neutralized, the IT Manager copied the trojan's URL and file name into Virus Total to check whether the threat had been observed and recorded elsewhere. The search came up with nothing, confirming that this was indeed a zero-day trojan uniquely discovered by Darktrace's AI.

IoT Hack: CCTV

Corporate Espionage?

Compromised Security Camera

The increasing connectivity of everyday devices has introduced a significant blind spot in the enterprise. IoT devices, often designed with unintegrated, basic security controls, are routinely targeted by threat actors and used as stepping stones into the network.

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 instead became a security risk itself.

The AI fought back at machine-speed, preventing a serious breach

Quick Reaction

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. 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.

IoT Hack: Smart Locker

Sensitive Customer Data Targeted

IoT Vulnerability

At an amusement park in North America, a threat actor attempted to steal sensitive customer data via a vulnerable IoT device: a 'smart' locker used by visitors to store personal belongings.

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.

Low and Slow

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.

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.

For this amusement park and others, Darktrace's cyber AI has neutralized countless 'low and slow' attacks at an early stage. By learning 'on the job', the system spots subtle threats that other tools miss. It continuously revises its understanding in light of new evidence, and generates autonomous actions that adapt to the threat as it unfolds.

Ransomware

Fast and Deadly

Automated Extortion

At 7:05pm on a Friday, an employee at a large telecommunications firm accessed his personal email from a corporate smartphone and was tricked into downloading a malicious file containing ransomware. Seconds later, the device began connecting to an external server on the Tor network.

Darktrace AI responded in moments. Just nine seconds after the start of the SMB encryption activities, Darktrace raised a prioritized alert signifying that the anomaly required immediate investigation. As the behavior persisted over the next few seconds, Darktrace revised its judgment and activated Antigena.

While the security team had left the office for the weekend, Darktrace Antigena responded autonomously, interrupting all attempts to write encrypted files to network shares. This instantly neutralized the threat before it could spread across the telco's sprawling infrastructure, giving the security team time to catch up.

As automated strains of ransomware continue to emerge on the Dark Web and in corporate networks around the world, organizations will need to fight back with AI to keep pace. Here as elsewhere, Darktrace's cyber AI response has become a critical component in the fight – containing fast-acting attacks before they have time to encrypt critical data and bring the business to a halt.

While the security team had left the office for the weekend, Darktrace Antigena responded autonomously, interrupting all attempts to write encrypted files to network shares.

Spear Phishing

A Targeted Email Attack

Email Attack

A well-known municipality in the United States recently fell victim to a targeted email-borne attack. While most phishing emails are launched as part of indiscriminate 'drive-by' campaigns, this campaign bore the markings of a coordinated and sophisticated cyber-crime. Each email was well-crafted and tailored to the intended recipient. The threat actor had also gotten hold of the city's address book, as the attack was delivered to recipients alphabetically, from A to Z.

Yet while each email appeared harmless and was customized to the recipient, the messages all contained a malicious payload hiding behind a button that was variously disguised as a link to Netflix, Amazon, and other trusted services.

Antigena spotted the campaign at the letter 'A', legacy tools woke up to the threat at 'R'

Hidden Links

Darktrace's AI was able to analyze these hidden links in connection with the normal 'patterns of life' of the intended recipients in the network. When the first email came through, Antigena immediately recognized that neither the recipient nor anyone in his peer group or the rest of the city's staff had visited that domain before. Antigena instantly raised a high-confidence alert, and suggested autonomously locking each link as it entered the network.

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

Supply Chain Attack

An Imposter Who Exploited A Trusted Relationship

Hijacked Email Account

Some of today's more resourceful cyber-criminals have learned that the easiest way into the enterprise is often through the front door, provided they can gain the trust of a legitimate user. By hijacking the account details of a trusted colleague, business associate, or vendor along the supply chain, threat actors can trick recipients into clicking a malicious link or transferring millions out of the business.

Darktrace's AI caught one such attack targeting a film production studio in LA, after the account details of a contact at a trusted supplier had been compromised.

Account details can be leveraged for many nefarious purposes, but in this case, the criminal seems to have used them to read through the contact's historical correspondence with an employee at the studio. After reviewing previous threads and learning how the contact and employee typically communicate, he sent a plausible reply to the employee's latest email.

The email was convincing – it mirrored the contact's writing style and tone

Social Engineering

The email was convincing – it mirrored the contact's writing style and tone, and made sense in the context of the relationship and previous discussions. It also included a malicious link that would have seemed harmless to any sensible employee receiving a link from a familiar contact at a familiar firm. These types of attacks are increasingly common, and very difficult to detect.

Darktrace's cyber AI discerned the weak indicators that revealed this 'trusted contact' to be a hijacked account controlled by an attacker. The AI response primed the network with the knowledge that the email and its content were outside the 'pattern of life' of the supposed sender. The employee was alerted and the malicious payload was neutralized.

Crucially, Antigena's decision was informed by the fact that this particular link would have been rare for both the sender and recipient given their prior communications, and the employee's normal 'patterns of life' in the network. The security team felt confident in its security posture knowing that Darktrace's AI didn't treat the recipient in the network as a mere email address. Rather, Antigena recognizes that the full scope of an employee's 'pattern of life' is often made manifest in disparate corners of the network, and in a way that can be correlated and analyzed intelligently by cyber AI.

About Darktrace

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

The company has over 900 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