

Catching the Silent Attacker, and the Next Phase of Cyber AI Threat Report 2018

Contents

Executive Summary	1
Internet of Things	2
Insider Threat	3
Cryptocurrency Mining	4
Ransomware & Worms	5
Quiet and Stealthy	6
AI & Immune Systems: The Next Phase	7

Executive Summary

Lurking beneath the surface of enterprises today are in-progress cyber-threats, ranging from new vulnerabilities to advanced attackers that have taken hold of critical information. And yet the leaders of those enterprises often have no clue, until it is too late. The company board may discuss the incident recovery plan but have little oversight of the way the company protects its data systems in the first place.

The cyber security breaches of the last year point to an inadequacy in the ability to see and detect emerging problems within our networks. As businesses explode in digital complexity and new types of threat emerge, organizations are increasingly outpaced and outmaneuvered by cyber-attackers.

The old approach to information security, concerned with keeping threats out by strengthening the network perimeter, doesn't work. Not against ransomware, which spreads too quickly to react in time, or slow and stealthy threats that fly under the radar, nor against insiders gone rogue, or hacked connected devices. The organizations that avoid the cyber headlines are embracing artificial intelligence to tackle these sophisticated and changeable adversaries.

With over 5,000 deployments, Darktrace is the world's leading cyber AI company, having identified over 60,000 threats that would otherwise have gone unnoticed. These are the attacks that didn't make front-page news.

This report highlights some of the most significant trends in the cyber-threat landscape, and feature case studies of real-world threats that Darktrace's Enterprise Immune System has uncovered. They range from criminals hijacking servers to mine for bitcoin to company employees stealing data – all had occurred despite traditional cyber security controls like firewalls and anti-virus, and were stopped before damage had occurred.

Internet of Things

The corporate network used to be a collection of desktop computers and servers. Today, it includes anything with an internet connection, from coffee machines to Fitbits, air-conditioning units to thermostats, video conferencing cameras to wind turbines. These devices are proliferating in the enterprise and have become a favored way for criminals to infiltrate a network. In 2017, Darktrace saw a 400% increase in the number of IoT (Internet of Things) security incidents in its customers' networks.

Connected objects are among the most vulnerable devices in the IT world. Half a million IoT devices were vulnerable to a botnet named Mirai, which launched mass-scale denial of service attacks in 2016. Darktrace has intercepted attackers that have targeted the biometric scanners used to keep unauthorized personnel out of manufacturing plants, internet-enabled drawing pads used by architects, and even a large fish tank that automatically dispenses fish food, but whose internet connection was exploited to exfiltrate data out of a large entertainment corporation.



increase in the number
of IoT security incidents

By taking control of such equipment, a criminal can not only steal information, but bring business to a grinding halt, either demanding payment to regain access or sabotaging the equipment beyond repair. They can also subtly alter data over long periods of time: for instance, by changing results obtained from a drilling company's sensors, a criminal could trick them into mining a depleted area.

Gartner forecasts that the number of connected things in use worldwide will reach 20.4 billion by 2020. And while attackers need to get lucky only once, businesses need to get it right every time. This imbalance is catching business leaders unawares and storing trouble for the future. The vulnerability of IoT devices cannot be eradicated but it can be very effectively mitigated, allowing a business to reap the rewards of connectivity without jeopardizing its data, systems and reputation.

Hacked Fingerprint Scanner at Manufacturing Plant

A global luxury goods manufacturer used biometric fingerprint scanners to restrict access to its warehouses. Unbeknownst to the security team, an attacker exploited vulnerabilities in one of these connected devices and started surreptitiously changing the biometric data in a suspected attempt to gain access to the highly secure facilities.

The compromise was not detected by the company's traditional security tools, because the targeted device was not monitored by the IT security team and, notwithstanding this, the activity was too novel to trigger alerts that flagged only 'known' malicious behaviors.

Darktrace AI instantly detected the foreign presence, as the infiltrated device started to exhibit highly anomalous patterns of behavior. The compromise was swiftly addressed, and the company unharmed.

Infiltrated Refrigeration System at Global Food Chain

Our future kitchens will feature connected fridges that report back to their owners on how full, cold, or clean they are – something that the food manufacturing industry already relies on.

A fast food chain managed a significant issue where a flaw in the software running on their storage refrigerators could have allowed attackers to change the temperature of the units, which in turn could have caused widespread food spoilage. The reputational and financial costs of recovering from such a scenario could be crippling.

Darktrace AI spotted this latent vulnerability as soon as the technology was installed, as the refrigerators were sending mass-delivery spam emails. Before it had been exploited by a would-be saboteur, the company rectified the flaw.

Insider Threat

We might think of hackers as hooded teenagers or state-sponsored mafias, but 65% of early-stage threats that Darktrace detects entail insiders misusing legitimate access to damage their employer, either knowingly or unknowingly.

The difficulty of detecting insider threat means that it can go unseen for long periods of time, sometimes years. Whether copying and selling medical information, inadvertently sending sensitive files to personal accounts, or insider trading, these activities can do gradual and long-term damage to a business's competitiveness.



Amid the noise of everyday business activity, the potentially harmful behaviors of insiders are hard to see. Most corporate strategies focus instead on prevention. Training and awareness programs promote good cyber security practices, and employees are encouraged to report suspicious behavior. These programs reduce risk, but do not account for insiders with the intention to harm, who hide in plain sight. Critically, too, you cannot guarantee that your employees will take the right course of action 100% of the time.

Within this category of threat type, there is a small group of users that pose an even greater threat to corporations, while often being subject to the least amount of scrutiny – privileged access users. These tech-savvy employees operate under a cloak of legitimacy and know how to cover their traces to conceal their activity when acting deliberately.

So, who watches the watcher? And how do you know when a friend turns to a foe? AI is a critical ally in the fight against insider threat, because it is able to gather together many subtle traces of activity, over long and short periods of time – something that human teams cannot do, because there is so much data and complexity to handle. AI can spot the deviating behaviors that stand out, even only slightly, compared to what it knows is normal behavior for any device, user or network.

No organization can eliminate insider threat, but it must be managed if companies are to protect the crown jewels in the long run.

An Unsuspecting Insider

A major US pharmaceuticals company found itself infected with an aggressive malware, due to insider threat. An employee had accessed BitTorrent, a peer-to-peer network used for transferring large files, to download media content, including pirated films.

While the employee judged his activity to be harmless to the company, the malware that was downloaded onto his desktop as a result was highly dangerous. It was designed to profile the device's vulnerabilities and open a backdoor for the attacker behind the malware to exploit.

No damage was caused as Darktrace had observed the download and saw it trying to move laterally around the network. With autonomous response, Darktrace took action instantly, stopping the malware's connections back to its command and control center, and neutralizing the threat.

“

Insider threat is one of the most serious threats a company can face.

”

John Carlin, Former Assistant Attorney General for U.S. National Security

Cryptocurrency Mining

Most cyber-attackers aim to steal or jeopardize data, but there are a growing number of hackers that break into your systems to exploit the infrastructure in a different way. They use your computing power to make money by mining for cryptocurrencies – often without you ever finding out.

Mining for cryptocurrencies, such as bitcoin, is the process of authenticating and legitimizing the transactions of these decentralized currencies. The more power hackers can steal from PCs, servers, and other devices, the faster they can mine the coins. With many thousands of computers co-opted with mining malware, the rewards can be significant. It is estimated that a thousand-strong army of hijacked computers could make over \$200,000 per year.



Some hackers prefer the stealthy, low-profile nature of these campaigns. Unlike ransomware, which spreads quickly, announcing its criminality to its victims, cryptocurrency mining can run in the background for months and even years, without the owners of the infrastructures knowing.

In the past six months alone, Darktrace has detected and intercepted over 1,000 incidents of cryptocurrency mining, and uncovered signs of it in 25% of all customers' networks - some mining operations are even run by rogue company employees. While the attackers are stealing computing power on a daily basis, they also pose a risk to the wider infrastructure and critical data. An unknown presence on the network is as unpredictable as it is dangerous.

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 its network for a period of 5 months.

After installing AI defense technology, 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. Without the AI that caught the anomalous behaviors, the operation could have continued for many months – long after the internship had ended.

“

AI-powered cyber defense offers the best chance to detect and fight back against crypto-mining attacks.

”

Justin Fier, Director of Cyber Intelligence and Analytics, Darktrace

Ransomware & Worms

Ransomware attacks encrypt computers' files en masse, rendering them inaccessible, and then demand a ransom payment in cryptocurrency in return for the decryption key. The operational, financial and reputational damage caused to victim organizations can be devastating.

While ransomware has been a part of the threat landscape for many years, in 2017, we saw the return of the network worm: malware that replicates itself to spread to other computers. The WannaCry ransomware attacks spread so quickly because they were self-propagating, using a worm. After the initial infection had taken place, the code spread from computer to computer, and across organizations globally, with no need for any action on the part of the humans behind the attack.

Ransomware is a 'noisy' attack, in that thousands of connections are made as the malware sets about encrypting the files on the hijacked computers. This blatantly abnormal activity typically triggers multiple alarm bells. The problem is the speed. Alarm bells are meaningless unless they have the capability to respond, swiftly and precisely, and shut the attack down before it has done damage. 'Autonomous response' AI that can intelligently fight back in real time is critical in these scenarios.

AI and Autonomous Response

The value of artificial intelligence lies in its ability to gather together lots of subtle pieces of information and draw conclusions that humans cannot easily arrive at.

For cyber security, advances in AI have enabled the creation of an immune system at the heart of the organization: a self-learning system that automatically forms an understanding about its digital environment and can see and respond to the early signs of a compromise or threat.

The next stage of cyber security is embracing AI to not only detect the previously undetectable, but also to fight back against the threats – like antibodies in our bodies.

In a world where cyber-attackers' movements can be measured in milliseconds, autonomous response is indispensable to keep up with the adversary. It buys your security team the time that they desperately need to catch up.

AI Stops Ransomware Before it Spreads

At 7.05pm on a Friday, an employee at a large telecommunications firm accessed his personal email from his corporate smartphone and was tricked into downloading a malicious file. Seconds later, the device began connecting to an external server on the Tor network.

Darktrace AI discovered that a new and advanced strain of ransomware had been deployed. The attack was automated and spread faster than ordinary ransomware.

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 24 seconds, Darktrace revised its understanding of the deviation, judging it to be an active, fast-moving threat.

While the security team had left the office for the weekend, Darktrace AI responded autonomously, and in a surgical manner. It interrupted all encryption attempts, while allowing normal activity to continue uninterrupted.

“

AI fights the most important battles for us.

”

Michael Sherwood,
CIO, City of Las Vegas

Quiet and Stealthy

Most of the cyber-attacks that make the news are not in fact single events, but a string of multiple data breaches and compromises that date back many months or even years. The complexity of digital infrastructures has hit the point where security teams struggle to keep abreast of new changes, emerging anomalies, and even current security procedures that should, in theory, be followed.

Sophisticated attackers take advantage of this complexity, taking a number of steps to avoid detection. Lying low in the noise of the network, they may only take action for a few milliseconds every day, trickling stolen data out very slowly. They are well under the radar and bypass traditional security controls.

And while data theft remains a significant objective, the integrity of data itself is a growing target for hackers and saboteurs. By manipulating data slowly, attackers can gradually raise doubts and erode confidence in data systems in the minds of key stakeholders. They can also do systemic damage, particularly to industries such as healthcare and financial services, where the reliability of data is a prerequisite.

Today's cyber defense strategies and technologies must tackle both fast and automatic offensives, as well as the slow and stealthy ones lurking beneath the surface.

Password-Hacking Tool Attempts to Evade Detection

A university in Italy was hit with an advanced attack containing 'active defense mechanisms' that helped it avoid detection by standard security monitoring.

Malware that logs keystrokes and sends user passwords to attacker-controlled destinations was employed. The program, called Smoke Malware Loader, uses numerous tricks to hide from prying eyes, modifying itself to avoid detection, and creating a smokescreen of redundant traffic to hide behind.

Despite these advanced techniques, an attacker still leaves breadcrumbs. In this case, Darktrace AI had observed many signals over a period of days; the download of a file, rare destinations outside of the network, and the transmission of small packets to advertise the attacker's presence to its command center (known as beaconing). The subtle signals are easily missed in everyday network traffic, but AI excels in making sense of such information. The incident was remediated within the week.

Audio Files Leaked from a Video Conferencing Unit

Dripping data out slowly over an extended period of time is more likely to go unnoticed, and this was the intention of attackers that broke into the systems of an international sports manufacturer.

New video conferencing equipment in the company's boardroom was exploited via an unauthenticated remote access tool, and small audio files were leaked to an unknown external server, bit by bit. Playing a high-risk game, with highly confidential boardroom conversations being targeted, the attackers were careful. The individual leaks were never over 10 KB and were performed within office hours, so as not to trigger suspicion.

The hack could have been highly successful if it had continued for long enough. But according to Darktrace AI, the behaviors that this device, the video conferencing system, was manifesting were highly anomalous. That's because it had learnt how that device normally behaved, and recognized the difference, even though the changes were slow and slight. It saved the company a major data breach.

Attackers Gathering Reconnaissance

The sophisticated planning stages of a cyber-attack were observed over a period of 4 weeks at a healthcare company, where sensitive data was surreptitiously being diverted through two unauthorized devices.

Two devices began to show signs of anomalous activity within the network. Shortly after joining the network, the devices started to act like web gateways, funneling internet traffic. The MAC addresses of these devices identified them as Raspberry Pis, low-cost single board computers, which served a suspicious external website to unwitting network users, and presented a fake login page.

While the objective of the attack was to harvest user credentials, and then use them to attack internal systems, the attackers didn't get this far. The set-up and early communications exhibited by the devices were subtle, but AI revealed that together they made up a concerning picture of abnormal activity.

AI & Immune Systems: The Next Phase

The cyber-threat case studies in this report are from companies that have identified cyber security issues *before* they became major incidents or disasters.

Darktrace invented the AI that does this, known as 'immune system' technology: it understands what's normal inside the digital environment and deals with it, just as your human immune system protects you from quiet viruses and harmful bacteria. And it is adaptive – it's always learning, always getting better.

Its widespread adoption by enterprises and governments is shifting security investments from the reactive mode that too many organizations rely on, when their perimeter defenses inevitably fail. While traditional security tools such as firewalls, antivirus and their newer iterations can fend off known threats that have been identified elsewhere, they are powerless to find the subtler and better-disguised ones.

Conversely, AI can deal with both the complexity of the network and the evolving sophistication of new threats – and, most importantly for a field of technology that is at its cutting edge, it is not a concept but an application of AI that is proven every second, every day, in networks around the world. It installs in just one hour, and starts to learn then and there.

The next phase of cyber AI and 'immune system' defense takes the technology one step further. Proven to distinguish the strange but benign, from the unremarkable yet dangerous, AI is now capable of taking action against threats before they have a chance to start their cycles of compromise and destruction.

Autonomous response is the new frontier in cyber security, helping businesses act earlier and faster against the most challenging cyber-threats, including automatically-propagating viruses that take just seconds to spread their infection.

Cyber AI tackles threats at both ends of the spectrum then. It has the skill and subtlety to uncover what the human eye struggles to make out, including the wolf in sheep's clothing slowly waiting to pounce. But it also has the speed and precision to fight back against the fastest of malwares, without disrupting critical operations.

Business leaders have a responsibility to fortify their organizations in the long run – not just against today's attackers, but tomorrow's too. AI is their essential ally to fulfil this duty, and level the playing field.

“

With Darktrace, talk about AI in cyber security has turned into action.

Ovum

”

About Darktrace

Darktrace is the world's leading AI company for cyber security. Created by mathematicians, the Enterprise Immune System uses machine learning and AI algorithms to detect and respond to cyber-threats across diverse digital environments, including cloud and virtualized networks, IoT and industrial control systems. The technology is self-learning and requires no set-up, identifying threats in real time, including zero-days, insiders and stealthy, silent attackers. Darktrace is headquartered in San Francisco and Cambridge, UK, and has over 30 offices worldwide.

Contact Us

North America: +1 (415) 229 9100

Europe: +44 (0) 1223 394 100

Asia-Pacific: +65 6804 5010

Latin America: +57 322 942 6401

info@darktrace.com

darktrace.com

🐦 @darktrace