

WHITE PAPER

Cloud Application Security

Protect your SaaS with Self-Learning AI



CONTENTS

SaaS: A Brave New World	1	Threat Finds	6
The State of SaaS Security	2	/ Phishing Emails Lead to Multiple Microsoft 365 Account Compromises	6
Self-Learning AI in the Cloud	3	/ Data Exfiltration from Salesforce	8
Autonomous Response to Machine-Speed Attacks	4	/ Account Takeover across Office 365 and Share-Point	9
Augmented Intelligence: Uplifting Human Teams	5	/ Malicious File Download in Box	10
		/ Multiple Indicators of Attack Unveiled	11
		Cyber AI Loop™	12
		/ Multi-platform coverage	12

SaaS: A Brave New World

Following the mass adoption of remote and hybrid working patterns, more critical data than ever resides in cloud applications – from Salesforce and Google Workspace, to Box, Dropbox, and Microsoft 365.

While these applications fuel efficiency and innovation at scale, security teams are now tasked with defending a complex patch-work of services with native security controls that are generally reliant on historical attack data, and incompatible across platforms.

Meanwhile, the efficiencies and scalability that these applications offer also act as a pull for cyber-criminals, who can launch their attacks at blistering speed and scale up and across the organization.

As cloud applications look set to remain an integral part of the digital estate, organizations are being forced to rethink how they protect their users and data in this area.

Security leaders are turning away from tools focussed on the attack, reliant on Threat Intelligence and confined to a single area of the digital estate, and towards Self-Learning AI which understands the digital business, revealing subtle deviations that indicate a cyber-threat, and then actioning an autonomous, targeted response.

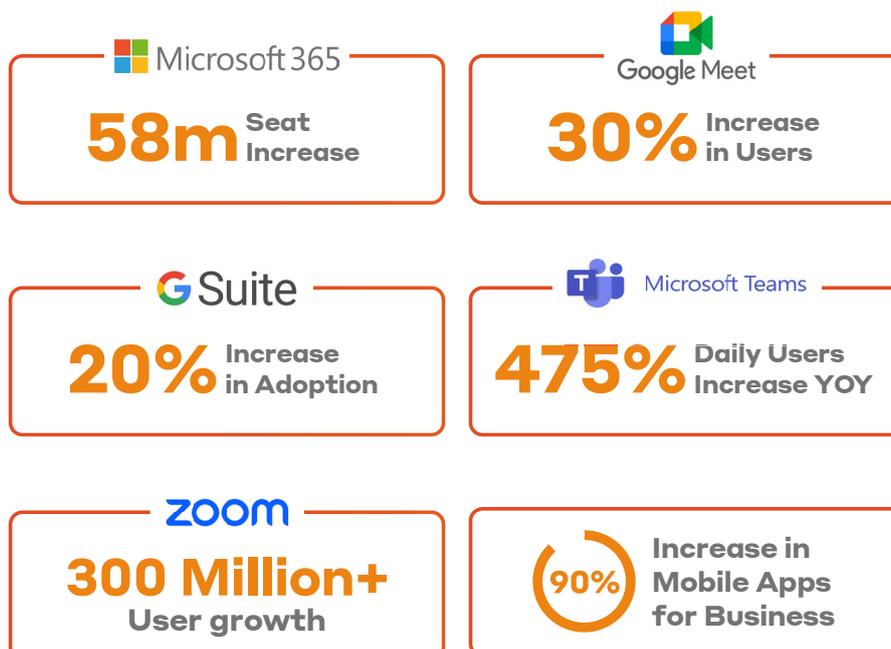


Figure 1: The explosion in SaaS usage since the global outbreak of COVID-19

The State of SaaS Security

Many organizations still rely on SaaS security tools that assume you can predict what the next attack will look like by referring to historical attack data. But with threat actors constantly innovating and updating their tools, techniques and procedures, novel and more sophisticated attacks bypass these static tools with ease.

Moreover, cloud security solutions have developed in siloes, and the complexity of managing this patchwork of disjointed technologies creates new vulnerabilities for stretched security teams. Whenever an innovative new SaaS platform begins to get adopted by business around the world, a familiar pattern can be found: a surge in micro-vendors focusing on a small set of security use cases, with no vision of the future or ability to innovate beyond features that are costly both in terms of money and time. In reality, defenders are having to protect an array of different platforms in an evolving threat landscape.

Common threat vectors that evade existing solutions include:



Account Takeover

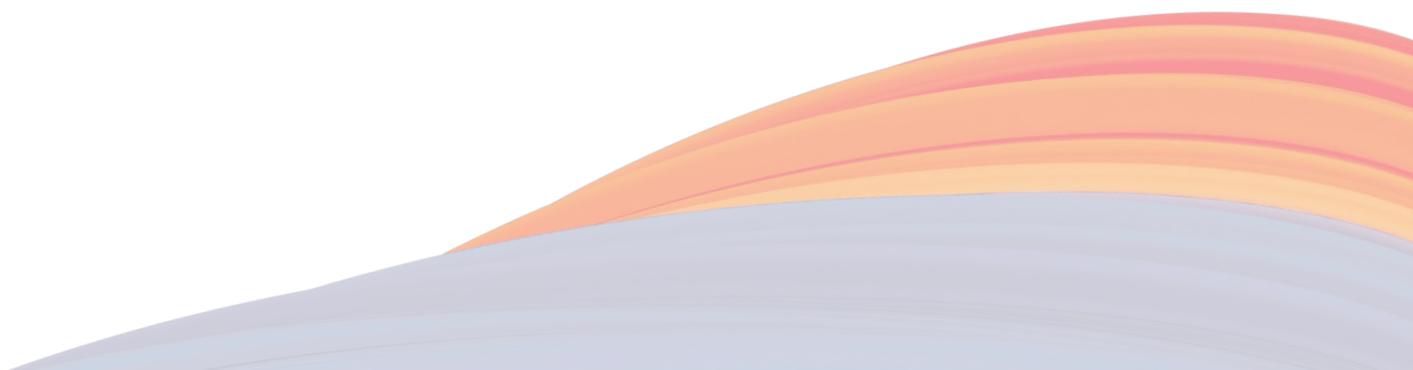
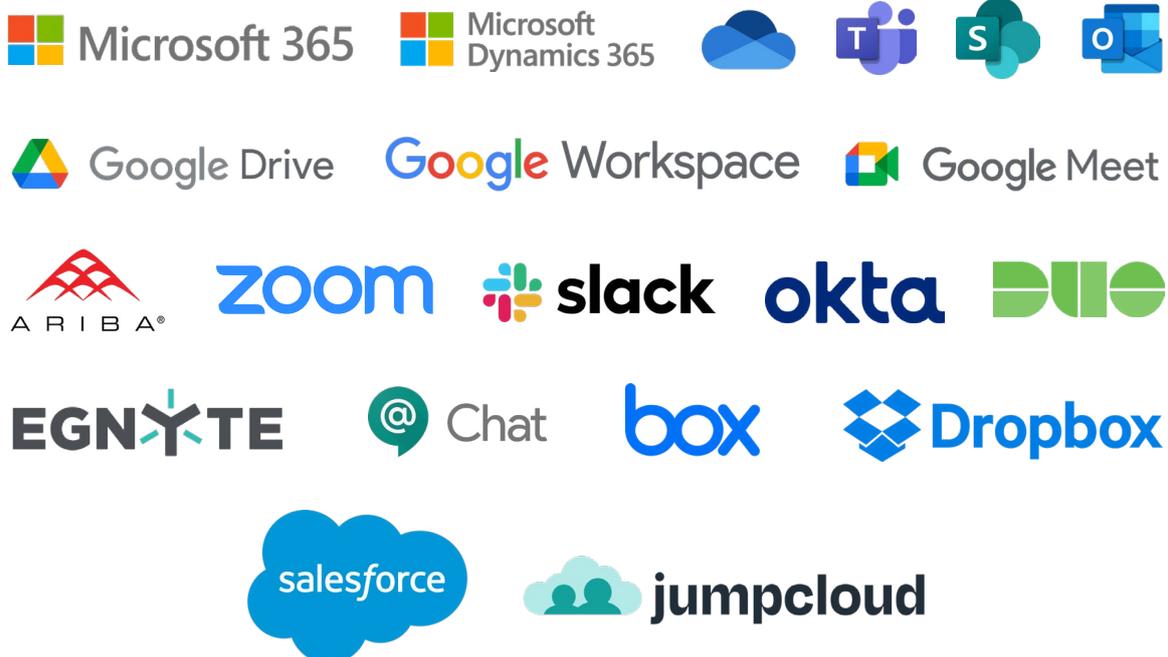


Insider Threat



Data Exfiltration

Organizations today are looking to simplify their security stack, and this requires a broader, more holistic approach to cloud security that protects multiple platforms with a single, proven approach to threat detection and response.



Self-Learning AI in the Cloud

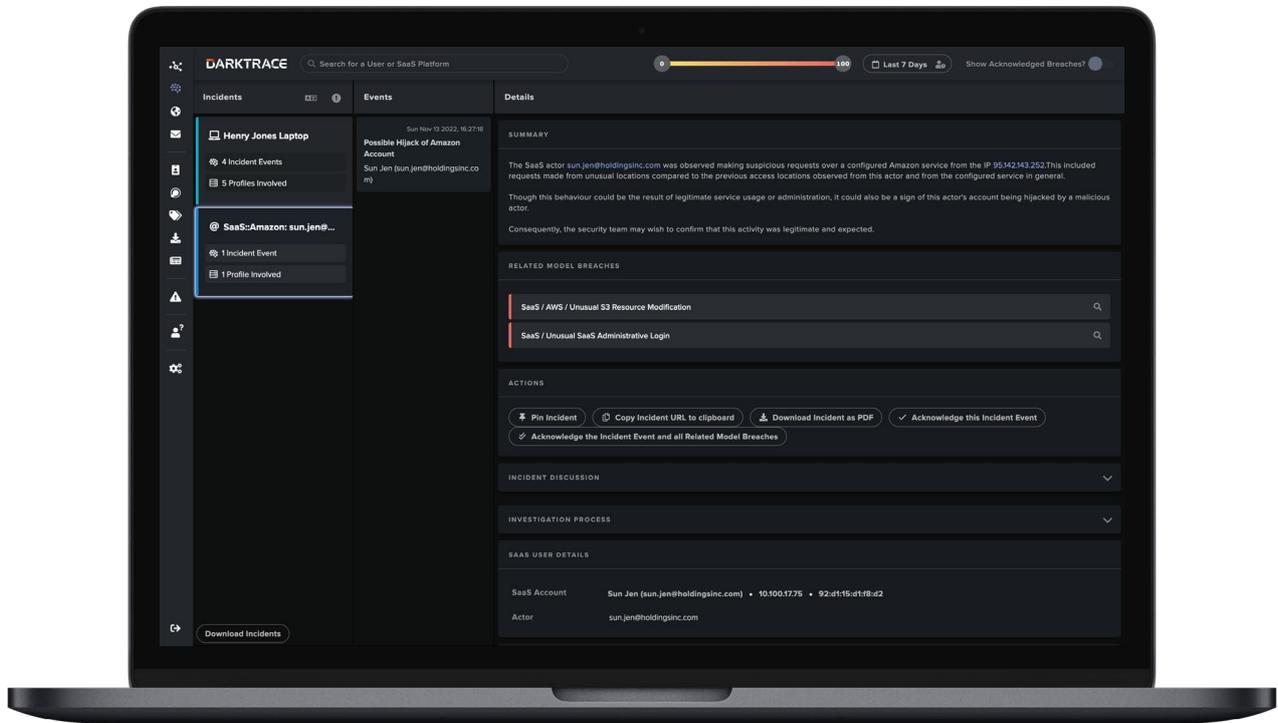


Figure 2: Darktrace's Threat Visualizer intuitively presents its findings

Darktrace is the first and only solution to deliver intelligent, self-learning protection that understands the full scope of your dynamic workforce in order to spot and stop threats across cloud applications.

Without relying on prior assumptions or pre-defined rules, its Self-Learning AI technology continuously learns the 'patterns of life' for every user and device across the business, analyzing workforce behaviors in context and spotting subtle deviations that indicate a genuine threat. By learning a sense of 'self' for your entire workforce, the technology detects threats such as compromised SaaS credentials and malicious insiders, no matter which application is targeted.

Rather than focussing on a narrow set of use cases, the technology protects the full range of cloud platforms – including the full Microsoft 365 product suite, collaboration platforms like Slack, and file sharing applications like Dropbox. It is also uniquely able to counter threats in Salesforce.

This self-learning approach has been proven to protect over 5,500 organizations in all industries around the world, containing the most sophisticated threats in cloud applications and beyond.

Huffmaster companies added Darktrace to our cyber security stack and it has already paid off – we are notified within minutes when a Microsoft 365 account is breached.

/ Director of IT, Huffmaster

Autonomous Response to Machine-Speed Attacks

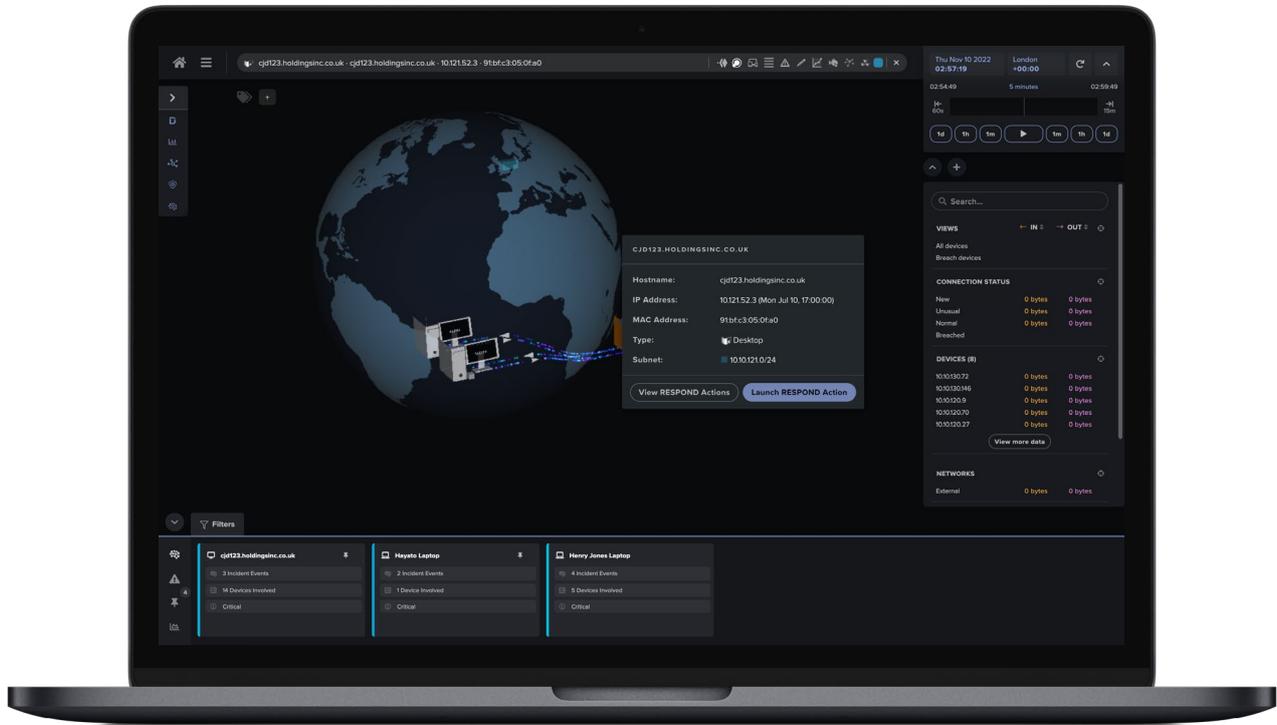


Figure 3: Darktrace RESPOND depicted within the Darktrace user interface

Cloud application security has historically put emphasis on detection over response, generally relying on human teams to initiate an action. Where response mechanisms are used, they usually come in the form of blunt, heavy-handed actions and pre-programmed responses, which tend to miss more sophisticated threats and can cause disruption to normal business.

Organizations are finding that cyber-criminals are increasingly striking out-of-hours – on weekends or holidays – when they know human response time will be slower. The agility and scalability of the cloud drastically quickens the dwell time of a typical cyber-attack, and so a fast response is paramount.

Autonomous Response has changed the paradigm of cloud application security. As Self-Learning AI understands the business, Autonomous Response actions a targeted, proportionate response that sustains normal operations while interrupting the threat.

Its responses are universal, with the technology taking co-ordinated action across a user's mobile device, tablet, home laptop, and office desktop. And it reacts at machine-speed, stopping account takeovers, insider threat, and data exfiltration in seconds.

For us, Autonomous Response technology combats the most sophisticated ransomware attacks out there and it does that within seconds of the threat emerging. This is the future of security.

/ CSO, Sunlife

Augmented Intelligence: Uplifting Human Teams

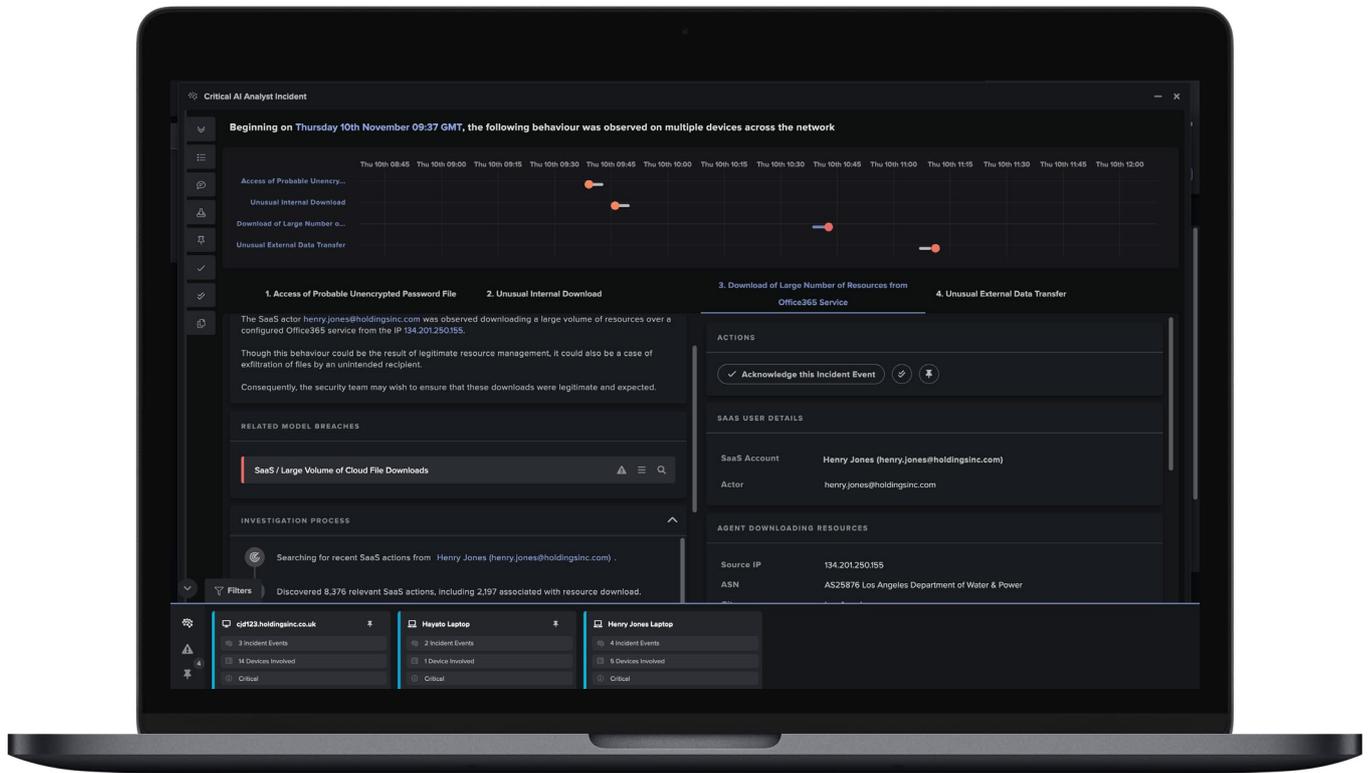


Figure 4: Cyber AI Analyst presents key information around an incident

Self-Learning AI not only reveals the subtle signs of advanced threats, but augments human teams by automating the process of threat triage, investigation, and reporting. Darktrace's Cyber AI Analyst launches an enterprise-wide investigation into every security event, connecting the dots between suspicious events in different areas of the business before settling on a high-level conclusion about the nature and root cause of the wider security incident.

Trained on an ever-growing data set of expert analyst behavior, Cyber AI Analyst automates analyst workflows at speed and scale, reporting on security incidents characterized by innovative attack techniques that would be impossible to capture with pre-defined playbooks.

Cyber AI Analyst produces a dynamic situational dashboard as well as written reports that immediately put strained security teams in a position to take action.

The new capabilities in Cyber AI Analyst have added real value to my team, especially the ability to launch on-demand investigations and query SaaS data or suspicious devices at any time.

/ CISO, Calligo

Threat Finds

/Phishing Emails Lead to Multiple Microsoft 365 Account Compromises

During a trial at a logistics company, Darktrace's AI detected a mass compromise when a phishing email led to multiple account takeovers.

The email came from a trusted source with a well-established history of correspondence. It therefore slipped under the radar of traditional tools, and the recipient was coaxed into clicking on a malicious link leading to a fake login page, where they were tricked into divulging their credentials.

Four days later, the attacker logged into the account from an unusual location, and proceeded to read files with sensitive information.

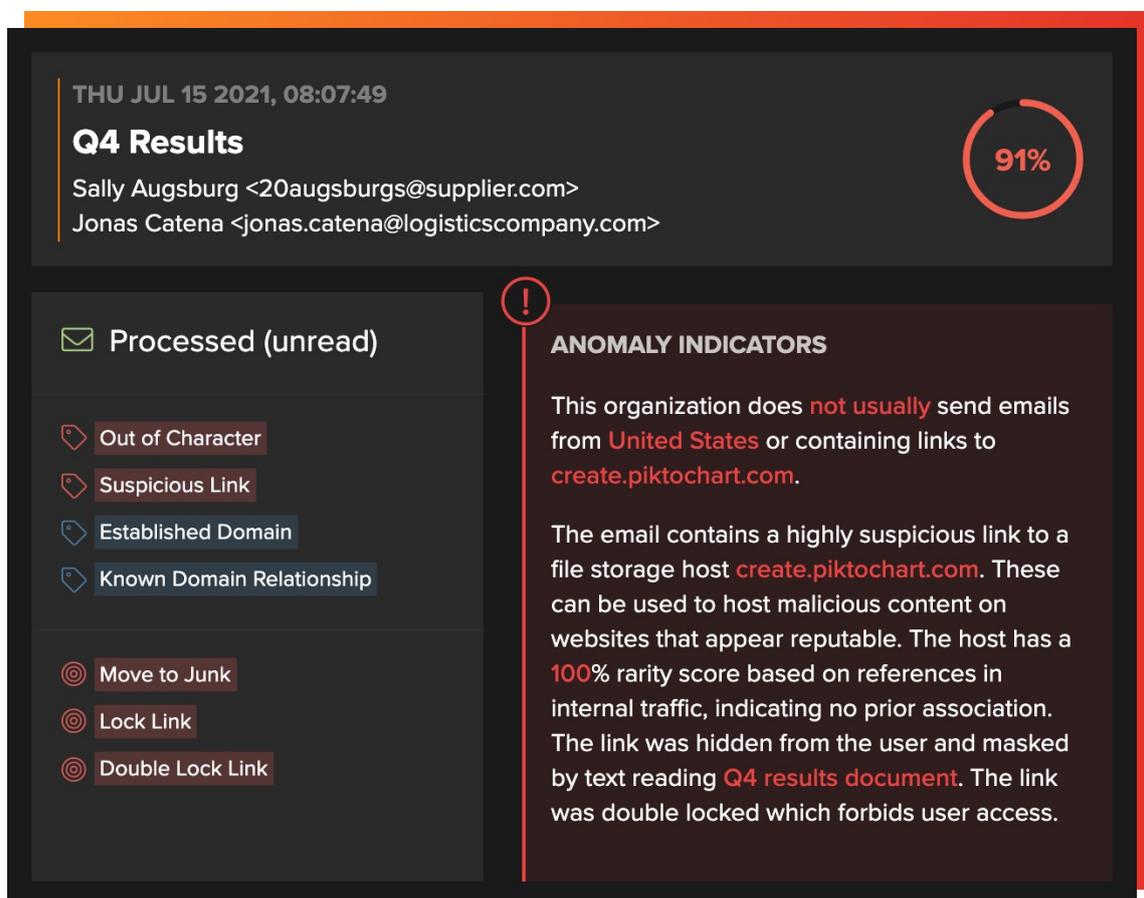


Figure 5: Darktrace reveals the indicators of threat in plain English (data has been anonymized)

The following day, Darktrace detected a new email rule from another unusual location. Almost immediately, a large volume of outbound emails was sent from the account, all containing the same suspicious link.

After this set of outbound emails, Darktrace alerted to unusual activity from anomalous locations on other company accounts – these users had been tricked into giving away their details from the emails supposedly sent by their colleague.

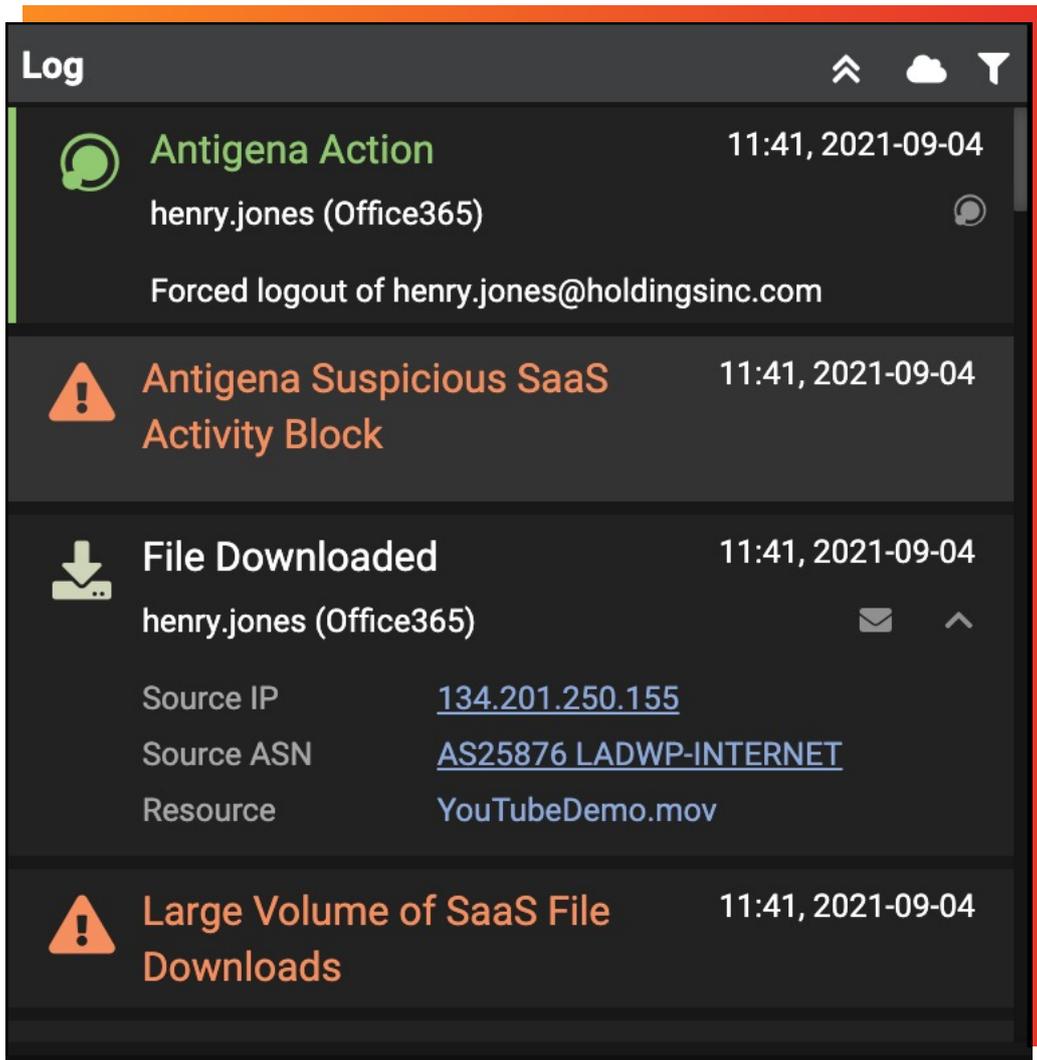


Figure 6: An example of RESPOND autonomously locking a SaaS account

More sensitive customer files were read, followed by a second spike in outbound emails from these hijacked accounts. This time, the emails were sent to external contacts. In total, over 450 phishing emails were sent to a wide range of third parties. Many of these third parties in turn had their credentials compromised – repeating the cycle once again.

Darktrace surfaced multiple subtle anomalies which revealed the initial email to be suspicious, and Darktrace RESPOND – Darktrace’s Autonomous Response technology – would have held the email back.

However, as this was a trial, it was set in passive mode, so the attack was allowed to unfold.

We can also see how Darktrace RESPOND would have recognized the unusual SaaS activity following the account hijacks, and intervened to lock the compromised accounts.

/ Data Exfiltration from Salesforce

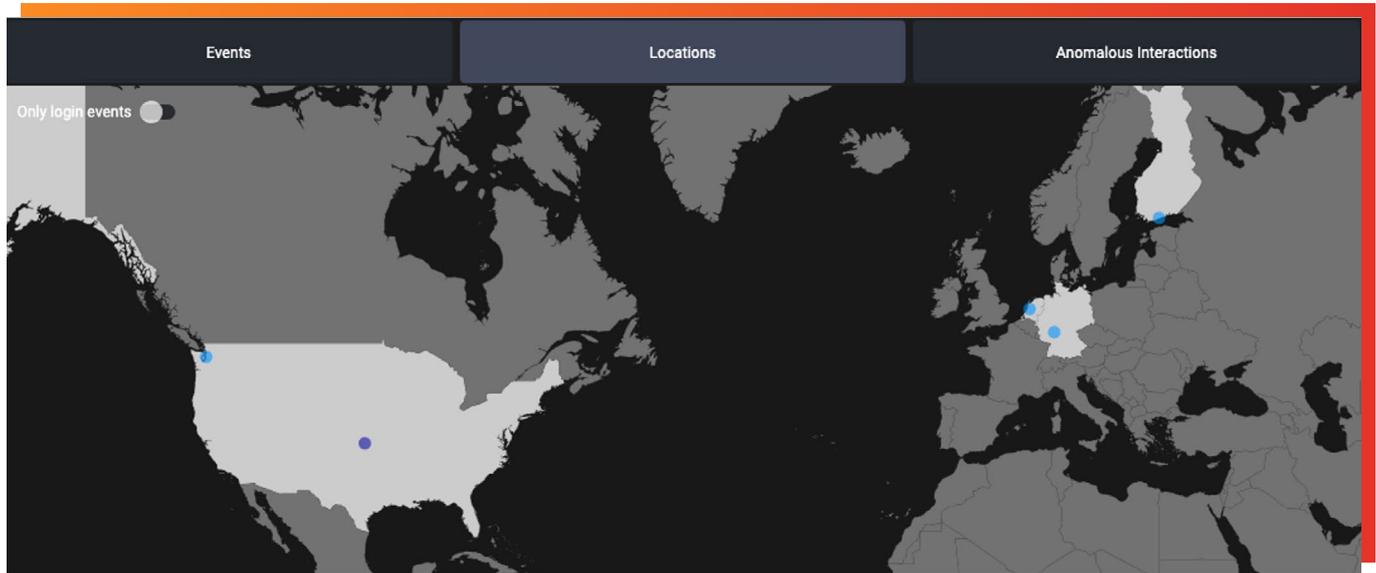


Figure 7: The geographical locations associated with the account, including the anomalous location in Finland

At a financial services organization, an employee was supposedly logged in and active on both their Salesforce and their Google Workspace account, from two separate locations, one of which was in a country not associated with the company in any way.

Darktrace picked up on further suspicious activity when a large amount of data was exported from the Salesforce account to an unusual destination, correlating multiple indicators of threat, and locking the attacker out of Salesforce, preventing the data exfiltration.

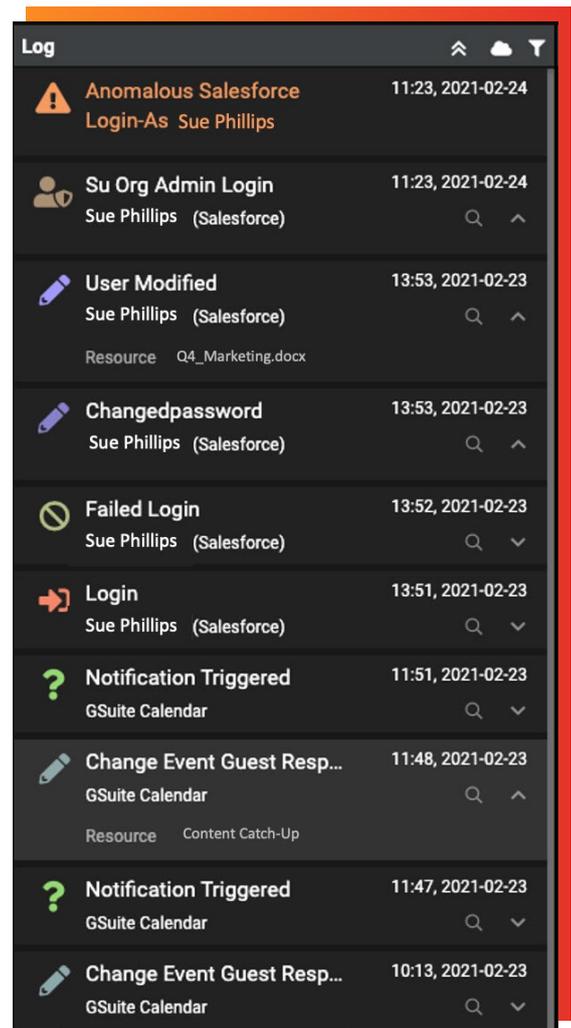


Figure 8: Darktrace correlated activity across multiple SaaS platforms to conclude that this was highly suspicious behavior



/ Account Takeover across Office 365 and SharePoint

At a US-based insurance company, Darktrace's bespoke understanding of the business across multiple applications 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 suspicious, as no other Microsoft 365 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 further used their illegitimate access to read and write to files on the user's personal SharePoint account.

Darktrace 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. Self-Learning AI stitched together these signs of anomalous behavior and locked the account, preventing the intruder from gaining further access.

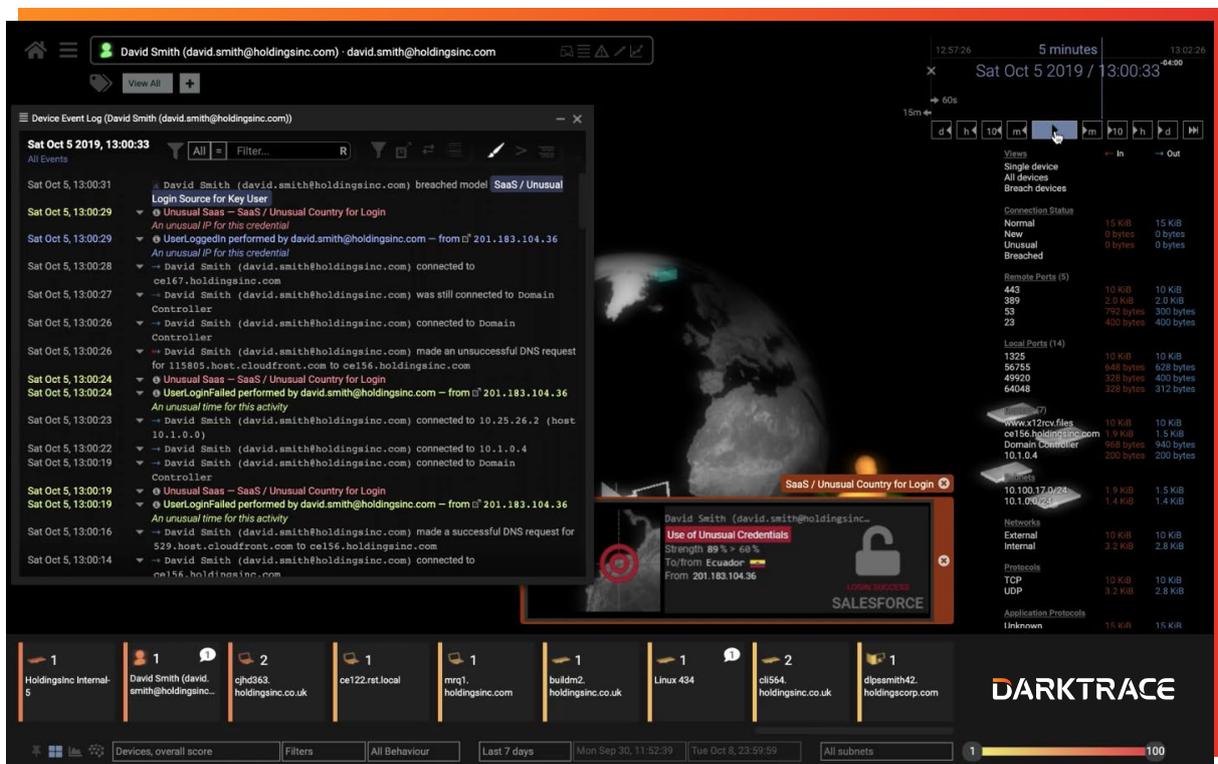


Figure 9: Darktrace's AI detects the unusual SaaS login location

/ Malicious File Download in Box

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

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

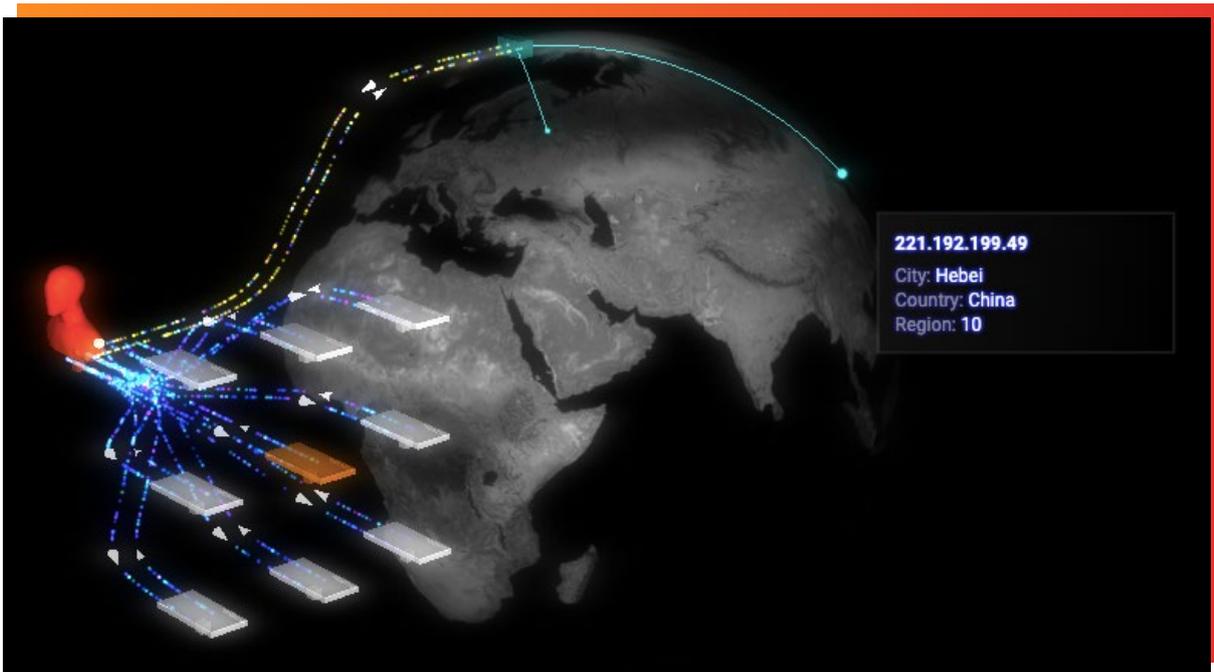


Figure 10: Darktrace showing the location of the unusual IP address

/ Multiple Indicators of Attack Unveiled

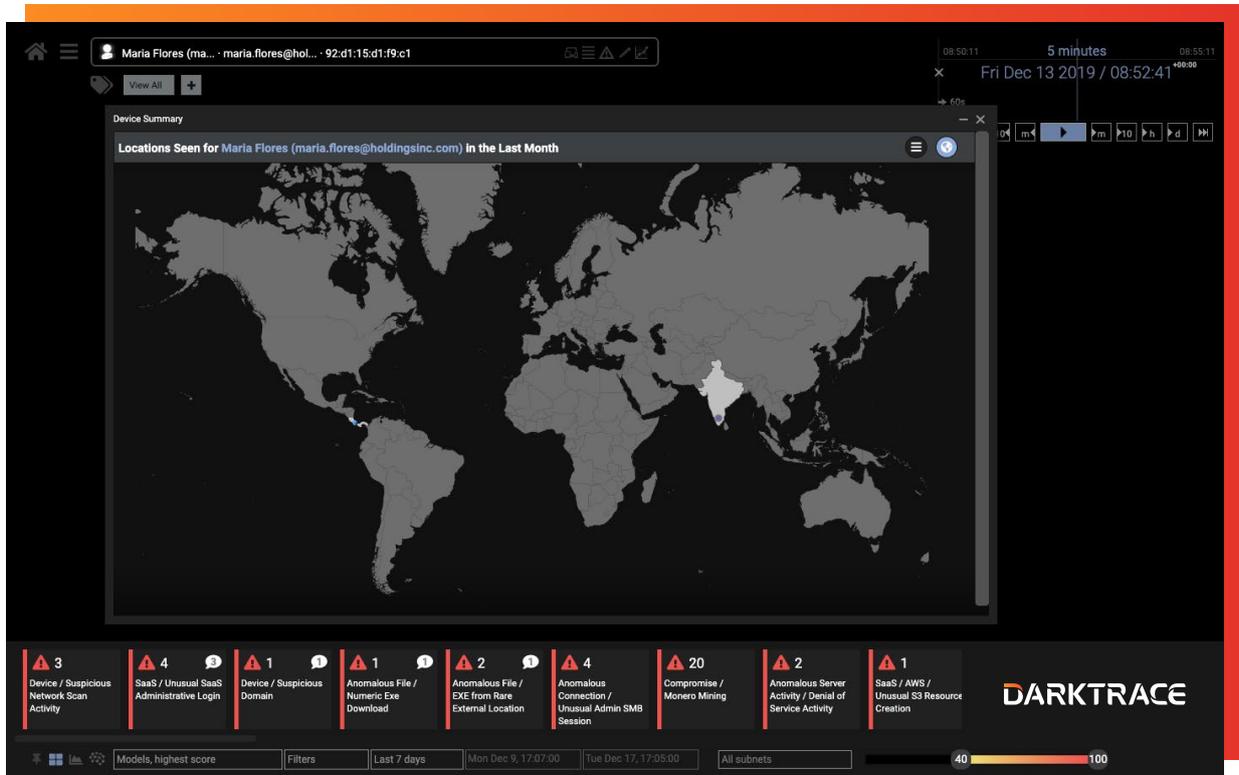


Figure 11: The user interface showing login locations

One Microsoft 365 account was used in a brute force attack against a well-known bank in Panama, with logins originating from a country that deviated from the normal 'patterns of life' of the company's operations. Darktrace identified 885 logins over a period of 7 days.

While the majority of authentications originated from IP addresses in Panama, 15% of the authentications originated from an IP address that was 100% rare and located in India. A further analysis revealed that this external endpoint was included in multiple spam blacklists, and that it had recently been associated with abusive behavior online – possibly unauthorized Internet scanning or hacking.

Darktrace then witnessed what appeared to be an abuse of the password reset function, as the user in India was observed changing account privileges in a highly unusual manner.

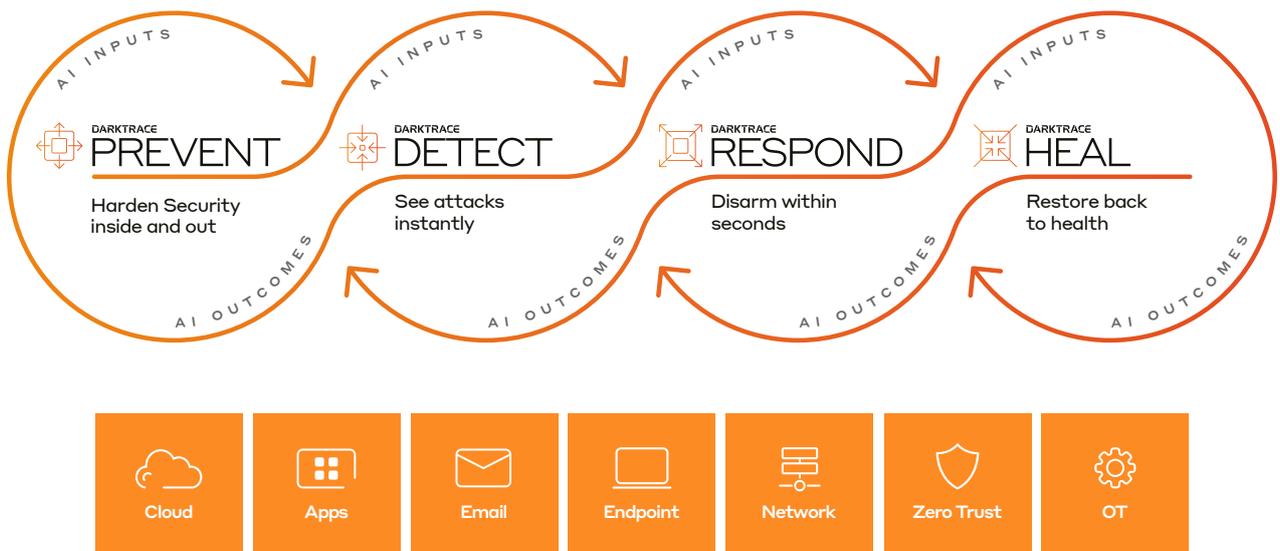
What marked the activity as particularly suspicious was that after the password reset, failed log-in attempts from an IP normally associated with the organization were observed, suggesting the legitimate user was locked out.

Using AI, Darktrace can detect and respond to email-borne threats and cloud-based attacks that other tools miss.

/ CIO, City of Las Vegas

Cyber AI Loop™

Self-Learning AI empowers a complete, always-on solution with autonomous feedback continuously improving the state of security.



Darktrace is delivering the first ever always-on feedback system that creates a virtuous cycle in which each capability strengthens and hardens the entire security ecosystem. Darktrace/Apps fits into the DETECT and RESPOND components of this ecosystem, sharing its insights with the rest of the security stack – and human teams through Explainable AI.

/ Multi-platform coverage

A cross-platform understanding of user behavior across email, network and cloud applications allows Darktrace to reveal the full scope of an incident – with each stage of the attack presented to security teams.

Darktrace/Apps + Darktrace/Email™

Darktrace/Apps brings additional insight from inbox behavior and works with Darktrace/Email to highlight the full picture of account takeovers and other cloud-based attacks (for example how an application login from an unusual destination relates to mass outbound emails from the same user).

Darktrace/Apps + Darktrace/Network™

Working with Darktrace/Apps, Darktrace/Network also provides an important piece of the puzzle, highlighting any lateral movement between the two areas of your digital estate. Key use cases include data exfiltration and insider threat.

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 2,100 people around the world and protects over 7,700 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 97242 2011

info@darktrace.com

[in](#) [twitter](#) [youtube](#)
darktrace.com