

A large, light blue wireframe globe in the top left corner, composed of many small triangles and lines, representing a global network or data distribution.

Cyber AI for SaaS Security

Protecting Your Dynamic Workforce

A smaller, light blue wireframe cloud shape in the top right, composed of triangles and lines, representing cloud computing or data storage.A larger, light blue wireframe cloud shape in the middle right, composed of triangles and lines, representing cloud computing or data storage.A smaller, light blue wireframe cloud shape in the bottom right, composed of triangles and lines, representing cloud computing or data storage.A smaller, light blue wireframe cloud shape in the bottom middle, composed of triangles and lines, representing cloud computing or data storage.

Introduction

Contents

Introduction	1
An Immune System Approach	2
Microsoft 365 Account Compromised and Sabotaged	3
Malicious File Download in Box	3
Microsoft 365 and SharePoint Infiltration	4
Microsoft 365 Account Compromise Exposes Financial Documents	4
Account Takeover at Panamanian Bank	5
Automated Brute-Force Attack	5
Attempted Access from Rural Japan	6
Phishing Link Leads to Account Compromise	6

Today's dynamic workforce is dispersed, agile, and unpredictable. From a security perspective, user behaviors are more disjointed than ever – cutting across a wide range of email, cloud, and SaaS services, and often operating well beyond the corporate network.

This fractured arrangement not only renders the traditional paradigm of the network perimeter obsolete, but also eludes the static and siloed nature of most security tools – from traditional on-prem defenses and email gateways, to cloud-specific controls developed for narrow use cases, such as access policy, governance, and workload compromise.

Perhaps the most critical locus of workforce activity today resides in SaaS applications, with users increasingly leveraging cloud services from Salesforce and G Suite, to Box, Dropbox, and Microsoft 365. These platforms fuel efficiency and innovation at an unprecedented scale, and organizations of all shapes and sizes have adopted them to manage highly sensitive data and mission-critical operations.

Yet with the increased adoption of SaaS, security teams must now defend a complex patchwork of services with native security controls that are not only rules based and static, but also incompatible across platforms. While a variety of third-party tools have been developed to unify this patchwork, their approach to threat detection is equally rigid, and their limited scope fails to track dynamic workforce behaviors that inevitably span the rest of the digital business.

Existing controls in this area represent the 'protective skin' of SaaS security – critical for preventing simple threats, but ultimately insufficient in the face of more advanced attacks. More often than not, these risks arise from compromised SaaS credentials and insider threats, especially from malicious administrators with privileged access. Since even legitimate workforce behaviors are often elusive and nearly impossible to predict, a static and pre-programmed approach to SaaS security is simply not enough.

To protect today's dynamic workforce, security teams must be equipped to discern when and how a trusted account has been leveraged for nefarious purposes. This requires more than just 'protective skin' – it requires an 'immune system' approach to security that is not only adaptive, but also grounded in a unified and behavioral understanding of your entire workforce.



An Immune System Approach

Darktrace's Enterprise Immune System is the first and only solution to deliver intelligent, self-learning protection that understands the full scope of your dynamic workforce – no matter where they work, or the nature of their applications.

As a complement to your protective skin, Darktrace's 'immune system' works by continuously learning the normal 'patterns of life' for every user and device in the business, correlating complex behaviors and relationships across email, cloud, and the corporate network. Without relying on prior assumptions or pre-defined rules, Darktrace's Cyber AI can thus analyze workforce behaviors in context, spotting subtle deviations that indicate a genuine threat.

By learning a sense of 'self' for your entire workforce, the Enterprise Immune System detects abnormal use of SaaS credentials and malicious insiders in seconds. This unique self-learning approach enables security teams to identify and contain even the most sophisticated threats in SaaS environments and beyond.

Investigating SaaS Threats with Cyber AI Analyst

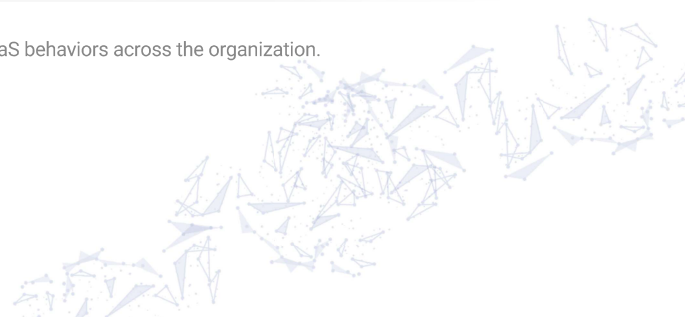
Darktrace not only identifies hijacked accounts and insider threats, but automatically investigates the full scope of SaaS-based security incidents. Every time the Enterprise Immune System detects a pattern of suspicious behavior, Darktrace's 'Cyber AI Analyst' launches into an enterprise-wide investigation, stitching together disparate anomalies before settling on a high-level conclusion about the nature and root cause of the wider security incident.

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 resource-strained security teams in a position to take action. In the threat finds that follow, Darktrace's Cyber AI identified threats in SaaS environments that not only evaded static defenses, but were also contained and actioned well before they could escalate into a crisis.



Figure 1: Darktrace's Threat Visualizer showing a summary of anomalous SaaS behaviors across the organization.



Microsoft 365 Account Compromised and Sabotaged

In one international non-profit, Darktrace detected an account takeover in Microsoft 365 that bypassed Azure’s AD 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-criminals 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.

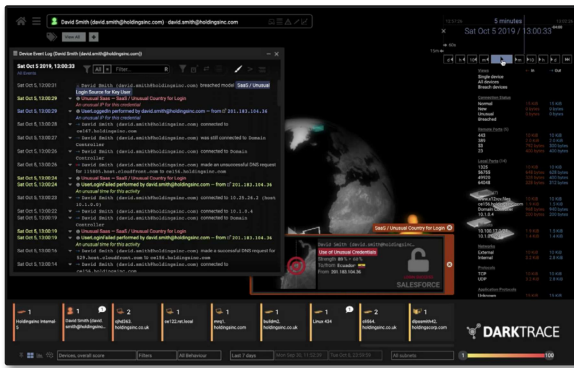


Figure 2: Darktrace’s AI detects the unusual SaaS login location

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 Cyber AI’s bespoke knowledge of ‘self’ for every member of the organization’s workforce, the technology was able to identify the threat immediately. The Enterprise Immune System 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 Cyber AI’s deep understanding of user behavior and granular visibility within Box allowed it to spot the subtle signs of account compromise. Moreover, when Darktrace 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.

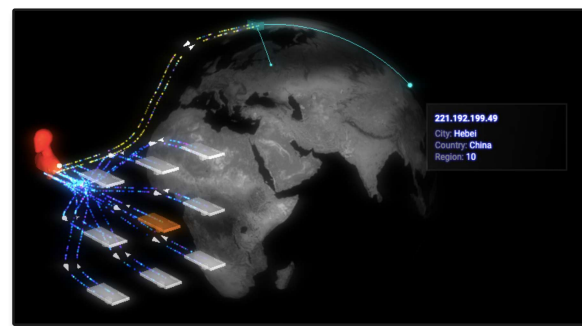


Figure 3: Darktrace showing the location of the unusual IP address

Microsoft 365 and SharePoint Infiltration

At a US-based insurance company, Darktrace Cyber 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, Cyber AI identified the behavior as anomalous, 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 Cyber 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, Cyber AI picked up on 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 assuring the threat was neutralized before serious damage could occur.

Microsoft 365 Account Compromise Exposes Financial Documents

At a global company, a threat actor infiltrated an employee's Microsoft 365 account to access sensitive financial documents hosted in SharePoint, including pay slip and banking details. The attacker went on to make configuration changes to the hacked inbox, deleting items and making updates that may have allowed them to cover their tracks.

When Darktrace AI Analyst investigated the incident, it immediately connected the dots between the anomalous activities to paint a comprehensive picture of the attack.

Darktrace Cyber AI detected that the threat actor logged in from a range of previously unseen external IP addresses and networks, located in Nigeria and Bulgaria. AI Analyst understood that the legitimate user had never logged in to this particular SaaS account from either of these addresses, indicating that logins from both locations were the actions of an imposter.

Darktrace further saw that the activity following these two anomalous logins occurred at an unusual time for the employee. With its ability to automatically analyze events to piece together attack narratives, AI Analyst was able to put together these weak signals of a threat and illuminate the likely account compromise.



Account Takeover at Panamanian Bank

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.

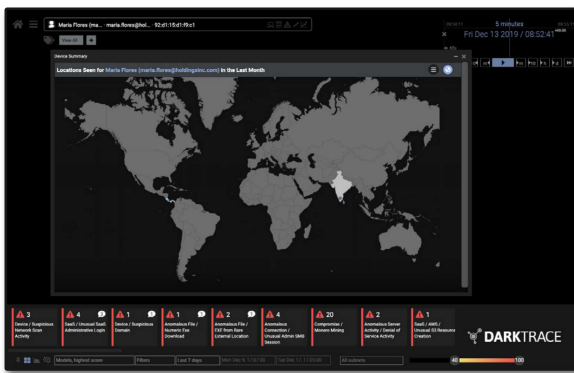


Figure 4: The user interface showing login locations

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.

03/12 20:45:39	SaaS:Admin	Regular	UpdateUser
03/12 20:45:39	SaaS:Admin	Regular	ChangeUserLicense
03/12 20:26:43	SaaS:Login	Regular	UserLoggedIn
03/12 20:26:43	SaaS:FailedLogin	Regular	UserLoginFailed
03/12 20:26:36	SaaS:FailedLogin	Regular	UserLoginFailed
03/12 18:31:31	SaaS:Login	Regular	UserLoggedIn
03/12 17:57:46	SaaS:Admin	Regular	ChangeUserLicense
03/12 17:57:46	SaaS:Admin	Regular	UpdateUser
03/12 17:06:57	SaaS:Admin	Regular	UpdateUser

Figure 5: The activity associated with the SaaS account, highlighting the changed credentials

Automated Brute-Force Attack

Darktrace detected several failed login events on a Microsoft 365 SaaS account every day over the course of a week. Each batch of login attempts was performed at precisely 6.04pm on six days. The consistency in both the time of day and the number of login attempts was indicative of an automated brute force attack, which is programmed to discontinue after a certain number of failed attempts in order to avoid lockouts.

Darktrace considered this pattern of failed attempts highly anomalous and so alerted the security team. Were it not for Darktrace correlating multiple weak indicators and fleshing out the subtle signs of emerging threat, this automated attack could have continued for weeks or months, making educated guesses at the users' password based on other information it had already gathered.

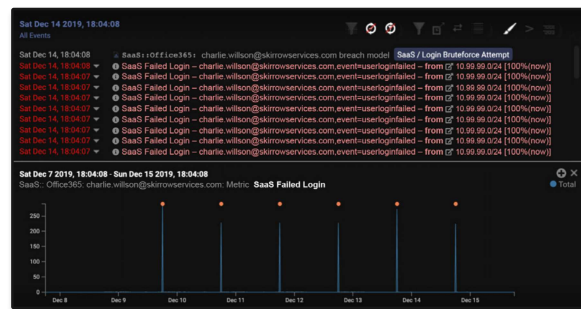


Figure 6: A graph illustrating the repeated login attempts

Attempted Access from Rural Japan

At a financial services corporation based in Europe, an Microsoft 365 credential was observed logging in from an unusual IP address linked to a location in rural Japan.

Although access from remote locations is possible if a user travels or uses a proxy service, this could also be a strong indicator of compromised credentials and malicious access by an unauthorized user. Given that the access point was substantially different from the usual accessing IPs, Darktrace flagged this as anomalous and immediately suggested further investigation.

The security team was able to remotely lock the Microsoft 365 account and reset the credentials, preventing the malicious actor from further activity. Had this activity gone unnoticed, the threat actor could have used their access privileges to deploy malware in the organization or solicit a fraudulent payment.

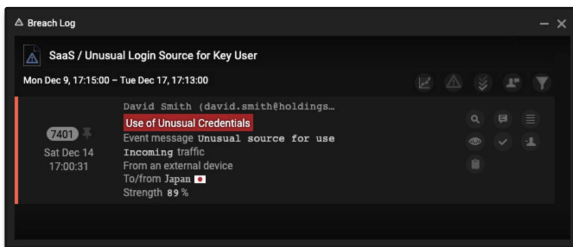


Figure 7: The login from Japan breached several models

Phishing Link Leads to Account Compromise

An employee clicked on a malicious link in an email that led to a spoofed login page, which captured her email address and password. Now armed with her credentials, the attacker pivoted back to Microsoft 365 and used them to login remotely.

Once inside the account, the attacker's next step was to propagate the attack to more victims, sending 99 emails with the subject line 'remittance advice' to a wide range of destination companies. This behavior might be expected in some instances, but not for this employee.

Darktrace also saw the creation of a new inbox rule. Forwarding rules are often created by attackers to spread spam or hide their activities: by automatically deleting emails after sending, the evidence trail is destroyed within the email system.

However, by independently monitoring the emails, and SaaS account activities, Darktrace was able to see the full picture of the attacker's activities, correlating the suspicious behaviors to identify the account compromise.

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](https://twitter.com/darktrace)