Darktrace Cyber AI: Compromised Email Credentials

Business leaders rarely consider how valuable a corporate inbox can be until it falls into the wrong hands. Yet once inside, threat actors enjoy a wide range of attack options and pivot points from which to choose. The ease with which attackers can gain access – whether through phishing campaigns, brute-force attempts, or exchanges on the Dark Web – should be cause for alarm.

In many cases, attackers will pillage your inbox for the valuable data it contains. Personal information from private chats or sensitive billing details can be leveraged for fraud or blackmail, while old email threads may contain highly confidential company information. Customer lists, pricing documents, and even roadmap and IP details are often just a few search terms away from being discovered.

In other cases, criminals will use the account as a launching point for the next stages of an attack. They may sit quietly in the background to gather intelligence about high-value executives or partners, reviewing documents, reading conversations, and learning how to blend in for when they inevitably strike. As with supply chain account takeovers, the ability to read an ongoing email thread and follow up with a plausible reply is often the most effective way to achieve an attack mission without triggering suspicion.

Figure 1: Darktrace's Threat Visualizer displaying the geographical login locations

Enterprise-Wide Context

While the possibilities for attackers are nearly endless, the options for defenders are limited. Corporate account takeovers are typically monitored for by simple and static defenses, including 'impossible travel' rules that rarely catch attackers who know how to hide. Thanks to its enterprise-wide view, however, Darktrace's Enterprise Immune System offsets the limitations of rules-based approaches and catches cyber-threats that inevitably get through.

Just like the human immune system, Darktrace's Cyber AI uses a deep knowledge of 'self' to spot malicious behavior that would otherwise go unnoticed. By learning the normal 'pattern of life' of every user and device in the business, Darktrace detects subtle deviations that reveal even the most careful criminals – whether those deviations are made manifest in suspicious login behaviors, inbox rule creations, or edits to user permissions. As cyber-threats become more advanced, leveraging self-learning AI across the entire digital business will be the only viable way to keep criminals out of your inbox.

Case Study 1: Automated Brute-force Attack

At another company, Darktrace detected several failed login events on a Microsoft 365 account using the same credential, 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 each day 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 alerted the security team. Were it not for Darktrace correlating multiple weak indicators and illuminating the subtle signs of an emerging threat, this automated attack could have continued for weeks or months, making educated guesses at the user's password based on other information it had already gathered.

Figure 2: A graph illustrating the repeated login attempts

Case Study 2: 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.