APT35 ‘Charming Kitten' discovered in a pre-infected environment

darktrace.com/en/blog/apt-35-charming-kitten-discovered-in-a-pre-infected-environment

Max Heinemeyer, Director of Threat Hunting | Friday April 23, 2021

APT35, sometimes referred to as Charming Kitten, Imperial Kitten, or Tortoiseshell, is a notorious cyber-espionage group which has been active for nearly 10 years. Famous for stealing scripts from HBO’s Game of Thrones in 2017 and suspected of interfering in the U.S. presidential election last year, it has launched extensive campaigns against organizations and officials across North America and the Middle East. Public attribution has associated APT35 with an Iran-based nation state threat actor.

Darktrace regularly detects attacks by many known threat actors including Evil Corp and APT41, alongside large amounts of malicious but uncategorized activity from sophisticated attack groups. As Cyber AI doesn’t rely on predefined rules, signatures, or threat intelligence to detect cyber-attacks, it often detects new and previously unknown threats.

This blog post examines a real-world instance of APT35 activity in an organization in the EMEA region. Darktrace observed this activity last June, but due to ongoing investigations, details are only now being released with the wider community. It represents an interesting case for the value of self-learning AI in two key ways:

- **Identifying ‘low and slow’ attacks:** How do you spot an attacker that is lying low and conducts very little detectable activity?
- **Detecting pre-existing infections without signatures:** What if a threat actor is already inside the system when Cyber AI is activated?

Advanced Persistent Threats (APTs) lying low

APT35 had already infected a single corporate device, likely via a spear phishing email, when Cyber AI was deployed in the company’s digital estate for the first time.

The infected device exhibited no other signs of malicious activity beyond continued command and control (C2) beaconing, awaiting instructions from the attackers for several days. This is what we call ‘lying low’ – where the hacker stays present within a system, but remains under the radar, avoiding detection either intentionally, or because they’re focusing on another victim while being content with backdoor access into the organization.

Either way, this is a nightmare scenario for a security team and any security vendor: an APT which has established a foothold and is lying in wait to continue their attack – undetected.

Finding the infected device

When Darktrace’s AI was first activated, it spent five business days learning the unique ‘patterns of life’ for the organization. After this initial, short learning period, Darktrace immediately flagged the infected device and the C2 activity.

Although the breach device had been beaconing since before Darktrace was implemented, Cyber AI automatically clusters devices into ‘peer groups’ based on similar behavioral patterns, enabling Darktrace to identify the continued C2 traffic coming from the device as highly unusual in comparison to the wider, automatically identified peer group. None of its behaviorally close neighbors were doing anything remotely similar, and Darktrace was therefore able to determine that the activity was malicious, and that it represented C2 beaconing.

Darktrace detected the APT35 C2 activity without the use of any signatures or threat intelligence on multiple levels. Responding to the alerts, the internal security team quickly isolated the device and verified with the Darktrace system that no further reconnaissance, lateral movement, or data exfiltration had taken place.

APT35 ‘Charming Kitten’ analysis

Once the C2 was detected, Cyber AI Analyst immediately began analyzing the infected device. The Cyber AI Analyst only highlights the most severe incidents in any given environment and automates many of the typical level one and level two SOC tasks. This includes reviewing all alerts, investigating the scope and nature of each event, and reducing time to triage by 92%.

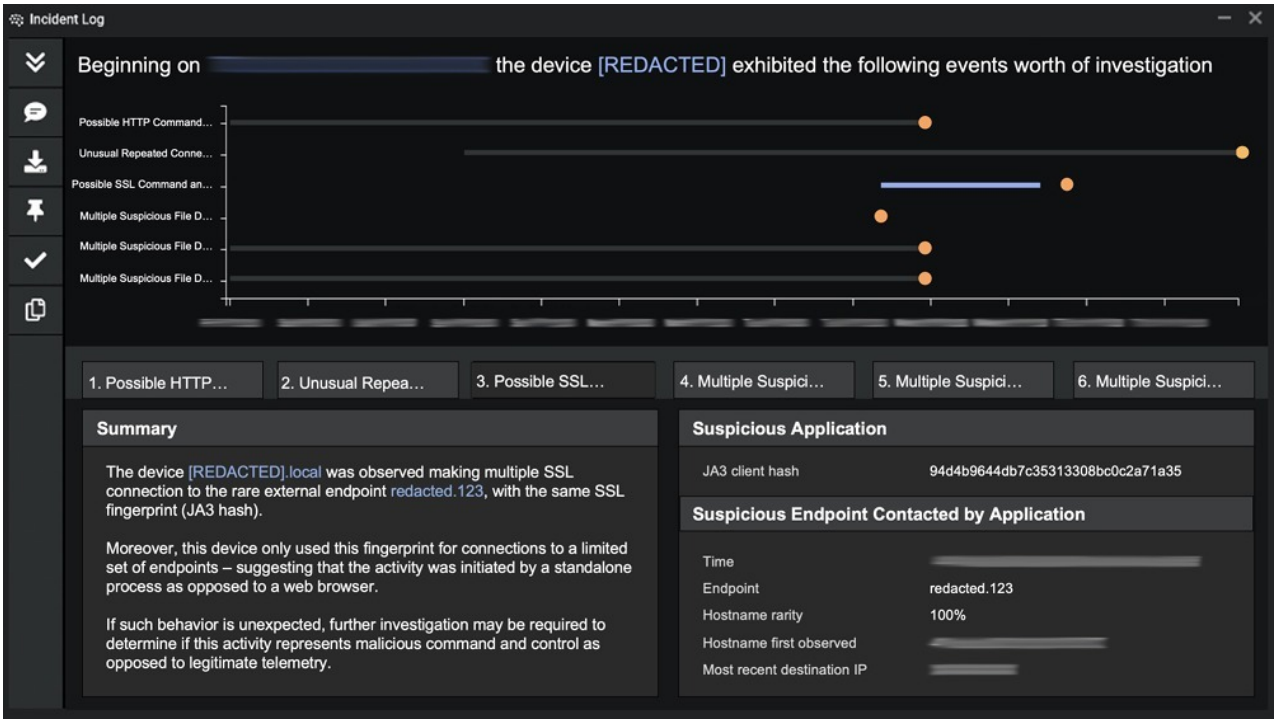

Figure 1: Similar Cyber AI Analyst report observing C2 communications

Numerous factors made the C2 activity stand out strongly to Darktrace. Combining all those small anomalies, Darktrace was able to autonomously prioritize this behavior and classify it as the most significant security incident in the week.

Figure 2: Example list of C2 detections for an APT35 attack

Some of the command and control destinations were known to threat intelligence and open-source intelligence (OSINT) – for instance, the domain cortanaservice[.]com is a known C2 domain for APT35.

However, the presence of a known malicious domain does not guarantee detection. In fact, the organization had a very mature security stack, yet they failed to discover the existing APT35 infection until Darktrace was activated in their environment.

Assessing the impact of the intrusion

Once an intrusion has been identified, it is important to understand the extent of it – such as whether lateral movement is occurring and what connectivity the infected device has in general. Asset management is never perfect, so it can be very hard for organizations to determine what damage a compromised device is capable of inflicting.

Darktrace presents this information in real time, and from a bird's-eye perspective, making the assessment very simple. It immediately highlights which subnet the device is located in and any further context.

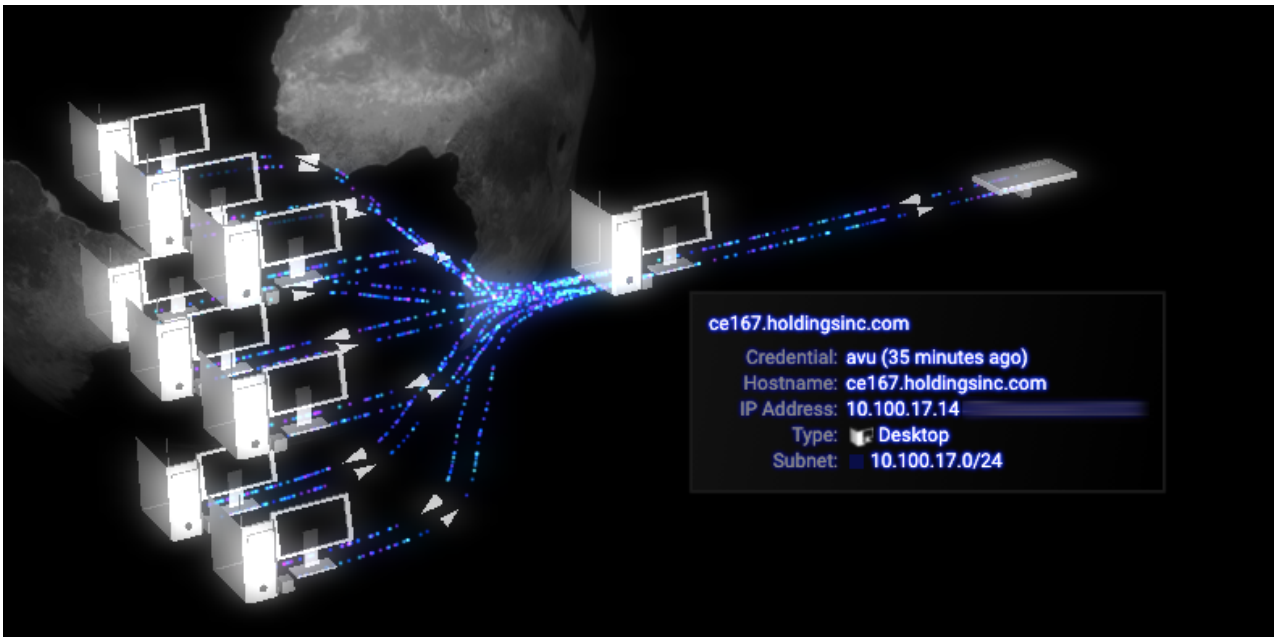

Figure 3: Darktrace's Threat Visualizer displaying the connectivity of a device

Based on this information, the organization confirmed that it was a corporate device that had been infected by APT35. As Darktrace shows any credentials associated with the device, a quick assessment could be made of potentially compromised accounts.

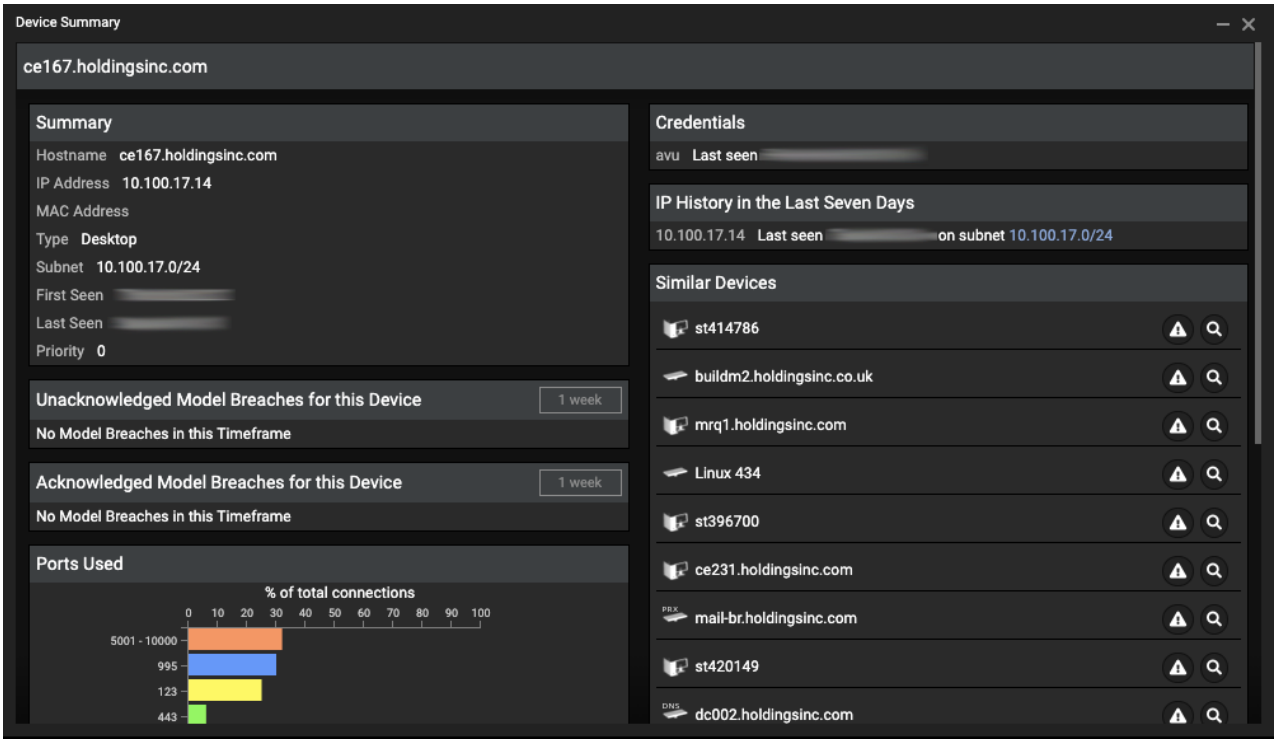

Figure 4: Similar and associated credentials of a device

Luckily, only a single local user account was associated with the device.

The exact level of privileges and connectivity which the infected device had, as well as the extent to which the intrusion might have spread from the initially infected device, was still uncertain. By looking at the device's event log, this became rapidly clear within minutes.

Filtering first for internal connections only (excluding any connections going to the Internet) gave a good idea of the level of connectivity of the device. A cursory glance showed that the device did indeed have some level of internal connectivity. It made DNS requests to the internal domain controller and was making successful NetBIOS connections over ports 135 and 139 internally.

By filtering further in the event log, it quickly became clear that in this time the device had not used any administrative channels, such as RDP, SSH, Telnet, or SMB. This is a strong indicator that no lateral movement over common channels had taken place.

It is more difficult to assess whether the device was performing any other suspicious activity, like stealthy reconnaissance or staging data from other internal devices. Darktrace provided another capability to assess this quickly – filtering the device's network connections to show only unusual or new connections.

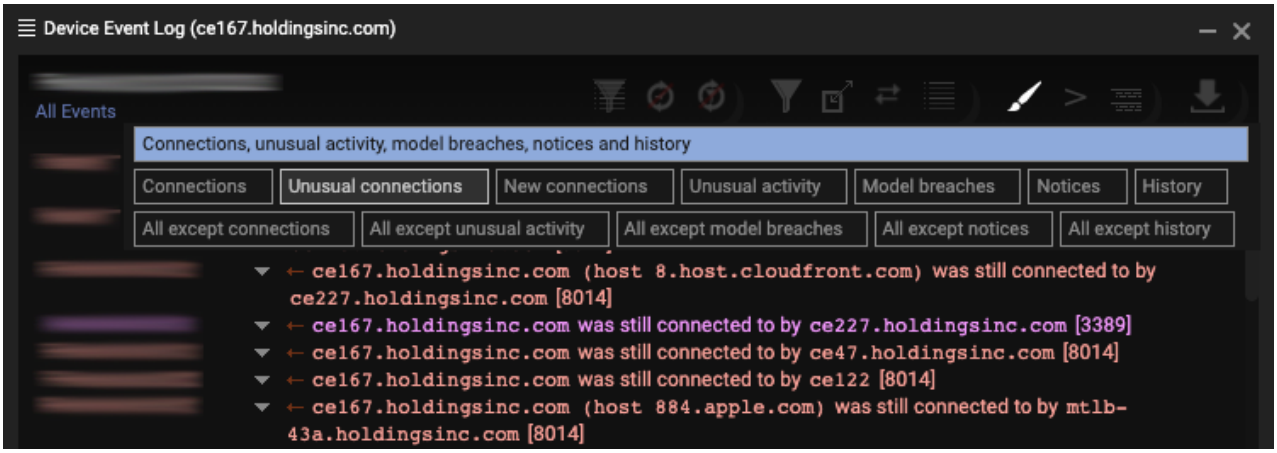

Figure 5: Event device log filtered to show unusual connections only

Darktrace assesses each individual connection for every entity observed in context, using its unsupervised machine learning to evaluate how unusual a given connection is. This could be a single new failed internal connection attempt, indicating stealthy reconnaissance, or a connection over SMB at an unusual time to a new internal destination, implying lateral movement or data staging.

By filtering for only unusual or new connections, Darktrace's AI produces further leads that can be pursued extremely quickly, thanks to the context and added visibility.

No further suspicious internal connections were observed, strengthening the hypothesis that ATP35 was lying low at that time.

Unprecedented but not unpreventable

Darktrace's 24/7 monitoring service, Proactive Threat Notifications, would have alerted on and escalated the incident. Darktrace Antigena would have responded autonomously and enforced normal activity for the device, preventing the C2 traffic without interrupting regular business workflows.

It is impossible to predefine where the next attack will come from. APT35 is just one of the many sophisticated threat actors on the scene, and with such a diverse and volatile threat landscape, unsupervised machine learning is crucial in spotting and defending against anomalies, no matter what form they take.

This case study helps illustrate how Darktrace detects pre-existing infections and 'low and slow' attacks, and further shows how Darktrace can be used to quickly understand the scope and extent of an intrusion.

Learn how Cyber AI Analyst detected APT41 two weeks before public attribution

Shortened list of C2 detections over four days on the infected device:

- Compromise / Sustained TCP Beaconing Activity To Rare Endpoint
- Compromise / Beaconing Meta Model

- Compromise / Beaconing Activity To External Rare
- Compromise / SSL Beaconing To Rare Destination
- Compromise / Slow Beaconing To External Rare
- Compromise / High Volume of Connections with Beacon Score
- Compromise / Unusual Connections to Rare Lets Encrypt
- Compromise / Beacon for 4 Days
- Compromise / Agent Beacon

Observed C2 destinations:

Cortanaservice[.]com, port 443, beacon intervals at 100 seconds

Max Heinemeyer

Max is a cyber security expert with over a decade of experience in the field, specializing in a wide range of areas such as Penetration Testing, Red-Teaming, SIEM and SOC consulting and hunting Advanced Persistent Threat (APT) groups. At Darktrace, Max oversees global threat hunting efforts, working with strategic customers to investigate and respond to cyber-threats. He works closely with the R&D team at Darktrace’s Cambridge UK headquarters, leading research into new AI innovations and their various defensive and offensive applications. Max’s insights are regularly featured in international media outlets such as the BBC, Forbes and WIRED. When living in Germany, he was an active member of the Chaos Computer Club. Max holds an MSc from the University of Duisburg-Essen and a BSc from the Cooperative State University Stuttgart in International Business Information Systems.