VECTRA
SECURITY THAT THINKS

THREAT REPORT

Breaking down the SolarWinds breach: an inside look at the methods used

TABLE OF CONTENTS

Vectra® protects business by detecting and stopping cyberattacks.

As a leader in network detection and response (NDR), Vectra® AI protects your data, systems and infrastructure. Vectra AI enables your SOC team to quickly discover and respond to would-be attackers —before they act.

Vectra AI rapidly identifies suspicious behavior and activity on your extended network, whether on-premises or in the cloud. Vectra will find it, flag it, and alert security personnel so they can respond immediately.

Vectra AI is *Security that thinks*®. It uses artificial intelligence to improve detection and response over time, eliminating false positives so you can focus on real threats.

The Vectra AI detection models provide real-time early warning and continuous visibility across the attack progression from on-premise to cloud without any dependency on IoCs, signatures, or other model updates.

HIGHLIGHTS

- Multiple communication channels, phases, and tools were used to establish interactive, hands-on-keyboard control. Each phase was designed to minimize the chance of detection, with techniques that defeat IDS tool signatures, EDR, manual threat hunting, and even common approaches to ML-based detection.
- The DGA used in this attack was different: a single, unique subdomain was generated for each victim, compromised of a globally unique ID calculated from local attributes and an encoding of the victim hostname.
- Vectra's AI will see through the evasion tactics applied and detect the tunnels as soon as they go active.
- Vectra uniquely protects the entire network of hybrid, on-premise, and cloud connectivity with learning behavioral models that understand both hosts and identities—tracking and stopping attackers earlier in the kill chain.

Summary

The SolarWinds Orion hack, now known as Sunburst or Solorigate, clearly illustrates the need for AI-powered Network Detection and Response (NDR). Preventative security and endpoint controls, while raising the bar, are insufficient, and legacy, signature-based Intrusion Detection Systems (IDS) have again been proven ineffective when detecting new attacks where indicators of compromise (IoCs) do not yet exist.

The SolarWinds attackers have put in significant effort and skill in evading preventive controls that involved network sandboxes, endpoint, and multifactor authentication (MFA), including:

- extensive checks to ensure that it was not in a sandbox or other malware analysis environment
- use of code signing and legitimate processes to evade common endpoint controls
- novel in-memory dropper to evade file-based analysis in distributing the Command and Control (C2) beacon
- MFA bypass using stolen Security Assertion Markup Language (SAML) session signing keys

The level of skill and focus required to cleanly bypass endpoint controls is a tribute to recent advances in Endpoint Detection and Response (EDR). However, it is also a reminder that a determined and sophisticated adversary will always be able to bypass prevention and endpoint controls.

Leveraging network detection and response—where network is defined broadly as everything outside of the endpoint—is a better approach in defending against this class of attack. The Vectra AI detection models provide real-time early warning and continuous visibility across the attack progression from on-premise to cloud without any dependency on IoCs, signatures, or other model updates. All of this works to identify and stop attacks like Sunburst/Solorigate/SolarWinds before damage is done.

Mapping the attack progression from on-premise to cloud

The objective of the Orion code compromise was to establish a reliable and stealthy C2 channel from the attackers to a trusted and privileged infrastructure component inside the datacenter—SolarWinds —which would provide the attackers with both an initial privileged accounts, as well as a pivot point to carry the attack forward.

Multiple communication channels, phases, and tools were used to establish interactive, hands-on-keyboard control. Each phase was designed to minimize the chance of detection, with techniques that defeat IDS tool signatures, EDR, manual threat hunting, and even common approaches to ML-based detection.

Below, we outline the attack progression from the initial backdoor through the establishment of persistent access to cloud environments, with a specific focus on Microsoft Office 365/email, which appears to have been a key objective of the attacks.

The Vectra AI coverage—with no reliance on IoCs or signatures—begins as soon as the initial C2 channel is established. The combination of behaviors observed directly on the SolarWinds server caused it to be prioritized as “Critical” even before any lateral movement occurred, allowing for early containment. If the attack were to progress, additional detections would provide full visibility into each subsequent phase even as the attack pivots into the cloud and Office 365.

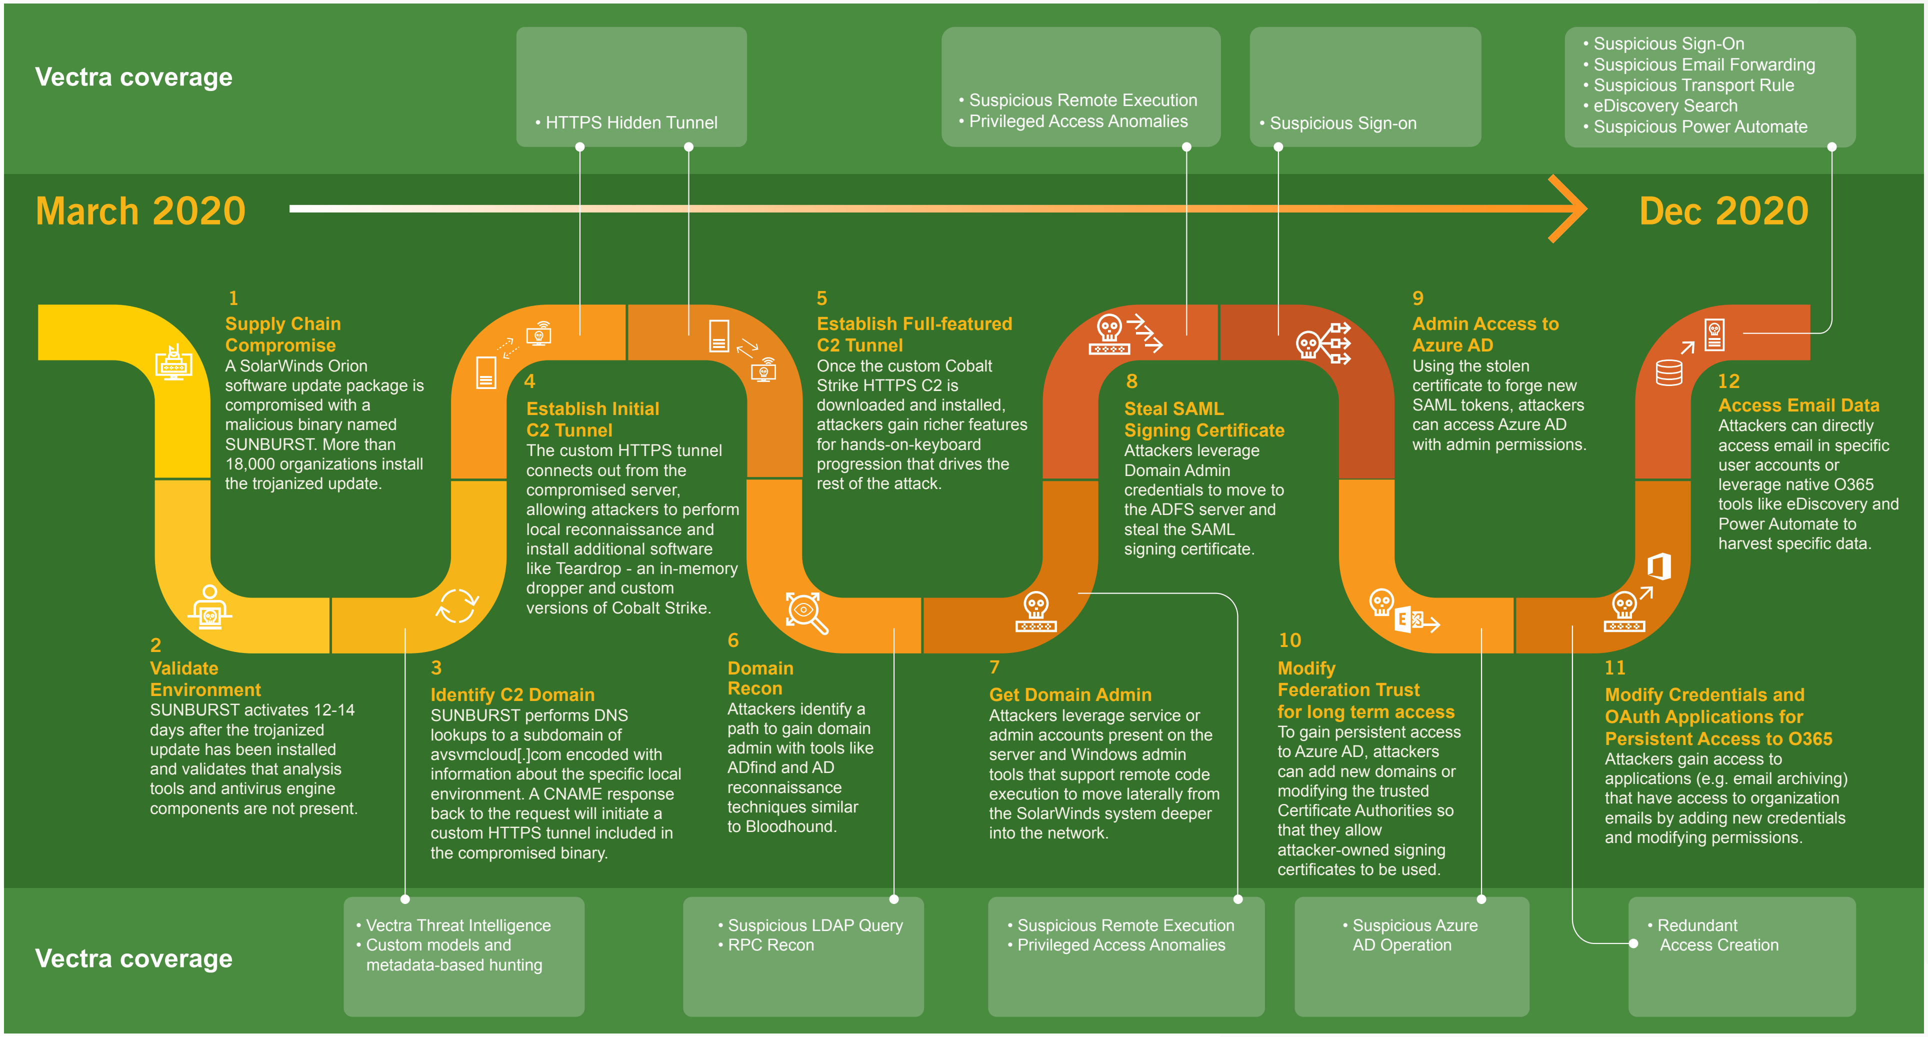

Vectra coverage

- HTTPS Hidden Tunnel

- Suspicious Remote Execution
- Privileged Access Anomalies

- Suspicious Sign-on

- Suspicious Sign-On
- Suspicious Email Forwarding
- Suspicious Transport Rule
- eDiscovery Search
- Suspicious Power Automate

March 2020 → **Dec 2020**

1
Supply Chain Compromise
A SolarWinds Orion software update package is compromised with a malicious binary named SUNBURST. More than 18,000 organizations install the trojanized update.

2
Validate Environment
SUNBURST activates 12-14 days after the trojanized update has been installed and validates that analysis tools and antivirus engine components are not present.

3
Identify C2 Domain
SUNBURST performs DNS lookups to a subdomain of avsvmcloud[.]com encoded with information about the specific local environment. A CNAME response back to the request will initiate a custom HTTPS tunnel included in the compromised binary.

4
Establish Initial C2 Tunnel
The custom HTTPS tunnel connects out from the compromised server, allowing attackers to perform local reconnaissance and install additional software like Teardrop - an in-memory dropper and custom versions of Cobalt Strike.

5
Establish Full-featured C2 Tunnel
Once the custom Cobalt Strike HTTPS C2 is downloaded and installed, attackers gain richer features for hands-on-keyboard progression that drives the rest of the attack.

6
Domain Recon
Attackers identify a path to gain domain admin with tools like ADfind and AD reconnaissance techniques similar to Bloodhound.

7
Get Domain Admin
Attackers leverage service or admin accounts present on the server and Windows admin tools that support remote code execution to move laterally from the SolarWinds system deeper into the network.

8
Steal SAML Signing Certificate
Attackers leverage Domain Admin credentials to move to the ADFS server and steal the SAML signing certificate.

9
Admin Access to Azure AD
Using the stolen certificate to forge new SAML tokens, attackers can access Azure AD with admin permissions.

10
Modify Federation Trust for long term access
To gain persistent access to Azure AD, attackers can add new domains or modifying the trusted Certificate Authorities so that they allow attacker-owned signing certificates to be used.

11
Modify Credentials and OAuth Applications for Persistent Access to O365
Attackers gain access to applications (e.g. email archiving) that have access to organization emails by adding new credentials and modifying permissions.

12
Access Email Data
Attackers can directly access email in specific user accounts or leverage native O365 tools like eDiscovery and Power Automate to harvest specific data.

Vectra coverage

- Vectra Threat Intelligence
- Custom models and metadata-based hunting

- Suspicious LDAP Query
- RPC Recon

- Suspicious Remote Execution
- Privileged Access Anomalies

- Suspicious Azure AD Operation

- Redundant Access Creation

Establishing a stealthy C2 channel

The C2 channel is not only carefully designed to evade signatures, but also simple-yet-common ML approaches. Vectra's C2 models use a more sophisticated AI approach, seeing through the evasion tactics to detect the underlying C2.

The Domain Generation Algorithm that doesn't behave like a DGA

As soon as the backdoor activates, a DGA is used to redirect the attackers' main C2 infrastructure. While the domain contacted by the backdoor is dynamically generated, this DGA is different than most. DGAs are typically used by botnets to stay ahead of threat intel feeds and blocklists. The victim system cycles through many random-looking, algorithmically-generated domains—most of which return NXDOMAIN responses—until they eventually resolve the currently-active C2 domain.

However, the DGA used in this attack was different: a single, unique subdomain was generated for each victim, compromised of a globally unique ID calculated from local attributes and an encoding of the victim hostname. The backdoor beaconed at random intervals to this single domain. The domain was successfully resolved each time. The victim continued the DNS beacon until it received a CNAME response, which was the signal to move to the next phase of the attack and provided the primary C2 domain connection.

The DNS beacon was designed in such a way that it would be nearly impossible to detect by security solutions:

- The domain in use was years old and was moved into Microsoft IP space approximately 2 months before initial use, defeating reputation or domain-aging approaches.
- No NXDOMAINs were returned, taking away a primary signal for common DGA detection.
- The algorithmically generated domain was a deep subdomain (e.g. 3mu76044hgf7shjf.appsync-api.us-est-2.avsvmcloud.com). Many legitimate services use randomization at this level and thus detecting on randomization would be extremely noisy and impractical.
- Beacons to a single sub-domain with a single piece of information, which evades DNS tunnel detectors that are looking for varying subdomains that would be used in order to encode more information.

Despite the methods used by the attackers, this phase of the attack is accurately detected by Vectra Threat Intelligence and custom models based on known IoCs.

Switching over to HTTPS

Once the C2 channel is activated, the binary establishes an HTTPS tunnel to the domain indicated by the CNAME response to the DNS beacon. HTTPS tunnels provide higher interactivity, more efficient data transfer, and are harder to detect than DNS tunnels.

Vectra's AI detection models kick into high gear as soon as the HTTPS tunnel is established.

The initial HTTPS phase uses a fully custom tunnel with a slow maximum beacon rate (1 minute, meaning low interactivity) and limited feature set. It appears to have been used for two main purposes:

(1) Collect additional detail about the compromised SolarWinds server, including an administrator account.

(2) Download the TEARDROP dropper and install a Cobalt Strike HTTPS tunnel customized to evade tool signatures.

After the Cobalt Strike HTTPS tunnel was established, the active, hand-on-keyboard phases of the attack progression leveraged this instead.

Both HTTPS tunnels use relatively standard methods to communicate, transferring commands and results over a periodic, encrypted beacon. There is little-to-no innovation in this tooling.

However, as with the DNS beacon, the domains were chosen to minimize the probability of detection: the domains were all well-seasoned/aged and hosted in Amazon AWS IP spaces. This simple tactic will, surprisingly, evade not just passive DNS and reputation systems but most ML-based detection systems as well. Why? Most ML detection starts by finding beacons in an environment, but due to the sheer number of legitimate beacons in networks, those models need a way to choose which beacons to alert on and which to ignore.

A simplistic way to address this problem is to only alert on beacons using domains under a few weeks old, effectively giving a pass to any domain with even modest age. It's also common to whitelist some CSP IP ranges, especially from mobile devices. This approach works fine for most red teams but, as this case clearly shows, does not defend against a sophisticated adversary that has planned the attack in advance.

Vectra uses a beacon detector as a foundation for tunnel detection. However, rather than taking the approach outlined above, Vectra's Hidden HTTP/S tunnel models use an advanced LSTM Recurrent Neural Network applied to time-series traffic metadata to identify C2 interaction over the beacon while ignoring things like stock tickers. As a result, Vectra's AI will see through the evasion tactics applied here and detect the tunnels as soon as they go active.

> Vectra's AI will see through the evasion tactics applied here and detect the tunnels as soon as they go active.

Progressing the attack: Identity and the cloud pivot

Once the C2 channel has been successfully established, it is used to provide hands-on-keyboard control to move the attack forward. Given the human element, details will vary from environment to environment. However, there is a common progression outlined in the research on both Cozy Bear TTPs in general and the SolarWinds attack in particular.

Moving laterally to get to Domain Admin

An attacker in controlling trusted infrastructure (SolarWinds) with access to various privileged accounts, including SolarWinds service accounts and in some cases even admin accounts, provides a short path to full domain admin permissions.

An obfuscated version of ADfind has been commonly reported as part of the attack tooling, used for domain enumeration including identification of domain admin accounts. Vectra covers domain enumeration with the **Suspicious LDAP Query** detection model.

Given an understanding of the group relationships and domain admin accounts, the next logical step is to use RPC to map the path to domain admin. Vectra detects this with an **RPC Recon** model.

Once the path is mapped, lateral movement will begin. Research suggests that task scheduler invoked via Windows Management Instrumentation (WMI) with accounts discovered on the SolarWinds server is the most common tactic used. Vectra provides detection coverage for this lateral movement in two ways:

- **Privileged Access Anomalies**: Vectra identifies and maps the relationships of all privileged accounts, services, and hosts. Attempts to use privileged accounts and services in unusual ways, including to execute commands, will trigger detection. For a more detailed description of Vectra's Privileged Access Analytics, see our solution brief on zero trust and privileged access.

- **Suspicious Remote Execution**: This model zeroes in specifically on RPC UUIDs associated with remote code execution via DCE/RPC, WMI, and DCOM, with learnings on a [src, dst, account, UUID] tuple. Remote code execution outside of the model learnings will trigger detection.

SAML Golden Ticket and Modifying Federation Trusts

With a Domain Admin account in hand, the attackers looked for opportunities to extend their presence into cloud environments. Given the prevalence of MFA on Azure AD and other federated identity providers, the attackers instead targeted and stole SAML signing certificates to forge new SAML token—used to bypass MFA. The SAML signing certificate is available in memory on the Active Directory Federation Services (ADFS) server or similar. Techniques to move laterally to the ADFS server would be the same as those to get to domain admin, but now with a more powerful account.

Vectra coverage for the move to ADFS/SAML server is the same as in earlier phases, specifically **Privileged Access Anomalies** and **Suspicious Remote Execution**.

The SAML signing certificate is then used to forge a SAML token, enabling global admin access to Azure AD. Vectra's **Azure AD Suspicious Sign-On** detects this access, analyzing multiple dimensions of the sign-on—including IP, geolocation, sign-on methods, and host—to identify unusual access.

Once global admin access is achieved, the adversary gains persistent access by modifying the domains Federated Trust configuration. This can be achieved by either adding new trusted domains the attacker controls, or new trusted Certificate Authorities (CAs) for creating and signing new certificates. These changes enable the attacker to issue new SAML tokens on an ongoing basis without needing access to the ADFS above. Vectra detects unusual changes to the Federated Trust configuration, as well as other Azure AD configuration, with the **Suspicious Azure AD Operation** detection model.

Persisting Access to Office 365 Email and Data

Office 365 holds a wealth of data, driven even more in 2020 by the increase of collaboration and document sharing to support remote work during the COVID-19 pandemic. Gaining ongoing email access appears to have been a key objective of the attackers, at least for specific targets in the U.S. Federal government. One technique discovered employed modification of the credentials and/or X509 keys for OAuth applications with email read/write access (e.g. Email Archiving applications). These changes ensured easy access to all email communications on an ongoing basis via API calls, with very low chance of detection.

Vectra's **Azure AD Redundant Access Creation** will trigger on these types of changes to OAuth app permissions.

With Office 365 access at admin level, a variety of other techniques could also be used to maintain access to both email and other Office 365 data. These include:

- Setting up eDiscovery and Compliance searches to look for specific information across all channels in Office 365 – Email, Teams, OneDrive, SharePoint. These applications are covered by Vectra's **Suspicious eDiscovery Search** and **Suspicious Compliance Search** detection models.

- Creating Power Automate flows to automatically exfil the data, either over HTTP or to external destinations such as attacker-controlled Google Drive or Box accounts. Vectra's **Suspicious Power Automate** detection addresses this vector.

- Other less stealthy techniques such as email forwarding or transport rules are also options. These are also covered by Vectra detection models – **Suspicious Email Forwarding** and **Suspicious Transport Rule**.

The Value of NDR

The SolarWinds hack demonstrates the utility—and necessity—of NDR solutions when taking steps to detect breaches that have bypassed preventative security, and to protect data. Network-based technologies are critical when countering the increasing sophistication of threats.

> Vectra uniquely protects the entire network of hybrid, on-premise, and cloud connectivity with learning behavioral models that understand both hosts and identities—tracking and stopping attackers earlier in the kill chain.

If you're ready to change your approach to detecting and responding to cyberattacks like these, and to get a closer look at how the Vectra Cognito Platform can find attacker tools and exploits, schedule a demo with Vectra today.

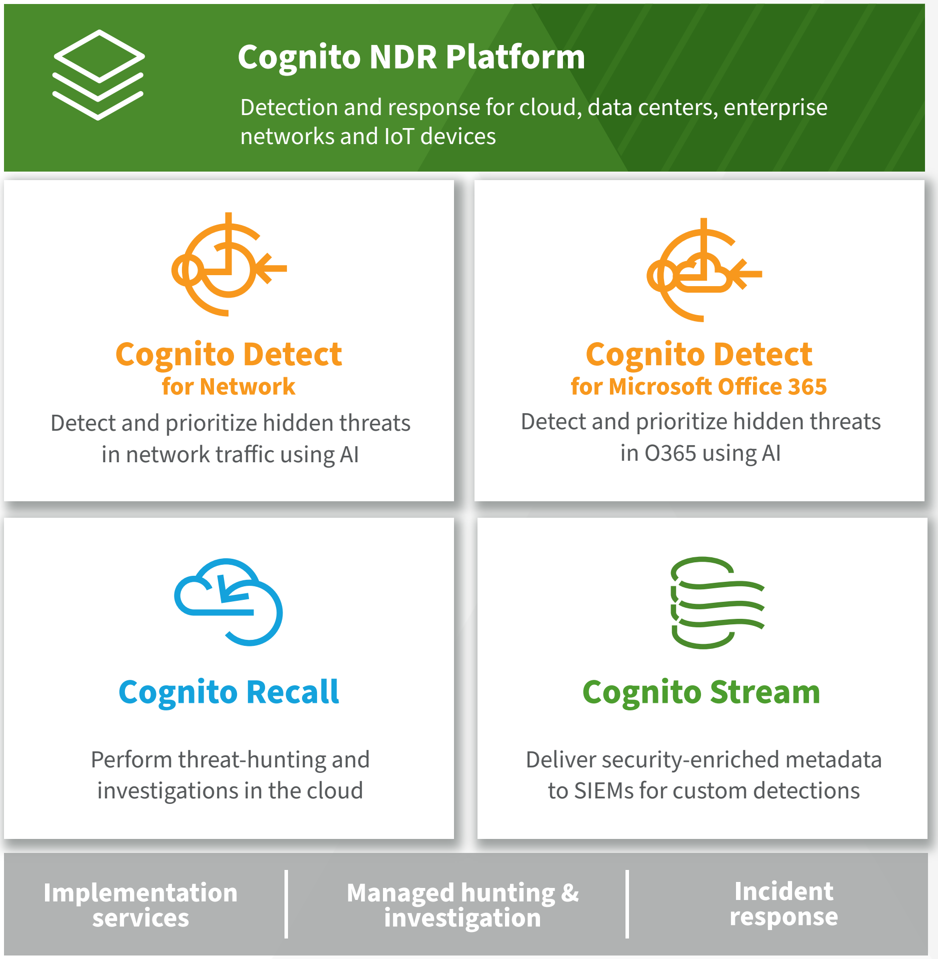

Cognito NDR Platform

Detection and response for cloud, data centers, enterprise networks and IoT devices

Cognito Detect
for Network
Detect and prioritize hidden threats in network traffic using AI

Cognito Detect
for Microsoft Office 365
Detect and prioritize hidden threats in O365 using AI

Cognito Recall
Perform threat-hunting and investigations in the cloud

Cognito Stream
Deliver security-enriched metadata to SIEMs for custom detections

Implementation services | **Managed hunting & investigation** | **Incident response**

For more information please contact a service representative at info@vectra.ai.

Email info@vectra.ai vectra.ai

SolarWinds Hack: Retrospective. Part 2: What caused the breach and what… | by Teri Radichel | Cloud Security | Dec, 2020 | Medium
Unraveling Network Infrastructure Linked to the SolarWinds Hack (domaintools.com)
Continuous Eruption: Further Analysis of the SolarWinds Supply Chain Incident (domaintools.com)

Version: **122320**