VECTRA
SECURITY THAT THINKS.®

SOLUTION BRIEF

Where Signatures and Simple ML Fall Short: Detecting a Novel New Attacker

At the request of a federal agency, Vectra was asked to prove the capability and promise of AI to stop a real world cyberattack scenario. The Vectra platform harnessing Security AI-driven Attack Signal Intelligence™ uses AI-driven detection, triage and prioritization so security teams can take an automated risk-based approach to cyberattacks. Vectra's AI-driven detections go beyond signatures and anomalies to understand attacker behavior and zero in on attacker TTPs across the cyber kill chain. For this exercise, once the rules of engagement were agreed upon, Vectra's offensive experts went to task and built a new 'sensor stimulation' event that followed a few unique asks of the agency:

1. Unknown to any prior signatures-based capabilities, CVEs or vulnerabilities.
2. Simulated how a sophisticated nation-state insider threat would behave.
3. Behave naturally: operate low and slow, taking four (4) weeks to worm around the environment and ultimately evade the capabilities in place for detection.

After following this process, Vectra progressively detected the escalating behaviors after additional integrations were enabled. During this pivot toward casting suspicion on the specific "insider" account regular scans, and brute force methods were emulated.

Think like an attacker

In support of customer requirements to validate their Vectra implementation, an exercise was designed that would simulate an advanced insider threat scenario. This approach was selected to delineate differences via Vectra AI/ML detection and signature, while also providing a realistic method to avoid signature-based detection in order to plan meaningful exploitation scenarios that test the limits of detection.

To achieve this, the team planned a simulated supply chain compromise along with a malicious insider that would simulate lateral movements and remote desktop access to specific nodes. The approach taken was differentiated in several ways:

1. As an insider threat, initial compromise and upstream navigation are not needed which reduces external reconnaissance.
2. The "insider" character was a developer who was specifically trained for this role, which permitted an additional insider to impersonate behavior by duplicating these trained behaviors.
3. All code that ran on endpoints was specifically written for this engagement to reduce signature detection and leverage native tooling — "living off the land".
4. After one week, the operational model was switched to generating detection, tilting the evidence to be linked to the "insider" as needed to protect the other "insider". This represented the compromised "supply chain" which provided a rationale for an insider compromising recently gained access.
5. All behavior models in the exercise were kept within a sophistication level compatible with standard technical workers as no esoteric hacking tools were used.

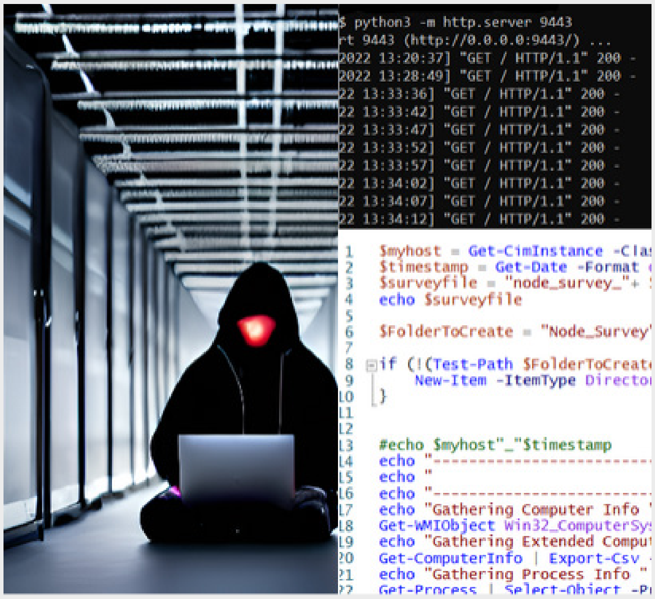

```
$ python3 -m http.server 9443
rt 9443 (http://0.0.0.0:9443/) ...
2022 13:20:37] "GET / HTTP/1.1" 200 -
2022 13:28:49] "GET / HTTP/1.1" 200 -
22 13:33:36] "GET / HTTP/1.1" 200 -
22 13:33:42] "GET / HTTP/1.1" 200 -
22 13:33:47] "GET / HTTP/1.1" 200 -
22 13:33:52] "GET / HTTP/1.1" 200 -
22 13:33:57] "GET / HTTP/1.1" 200 -
22 13:34:02] "GET / HTTP/1.1" 200 -
22 13:34:07] "GET / HTTP/1.1" 200 -
22 13:34:12] "GET / HTTP/1.1" 200 -
```

```
$myhost = Get-CimInstance -Clas
$timestamp = Get-Date -Format
$surveyfile = "node_survey_"+
echo $surveyfile

$FolderToCreate = "Node_Survey"

if (!(Test-Path $FolderToCreat
    New-Item -ItemType Directo
}


#echo $myhost"_"$timestamp
echo "----------------------------
echo "
echo "----------------------------
echo "Gathering Computer Info "
Get-WMIObject Win32_ComputerSys
echo "Gathering Extended Comput
Get-ComputerInfo | Export-Csv
echo "Gathering Process Info "
Get-Process | Select-Object -P
```

"For the first time in my career a defensive technique not only identified a presence and malicious activity, but characterized my motivation and nuanced differences in behavior"

–The Bad Actor

The bad actor in the exercise adhered to strict rules of engagement to ensure that no data out of scope could be manipulated, exfiltrated or otherwise removed. The offensive expert used what little information they could ascertain to design their attack. In addition to the four-week actual exercise, the bad actor leveraged the prior six weeks to design their plan, and shared as little information as possible with the target to obtain concurrence in the approach while not giving away their methods.

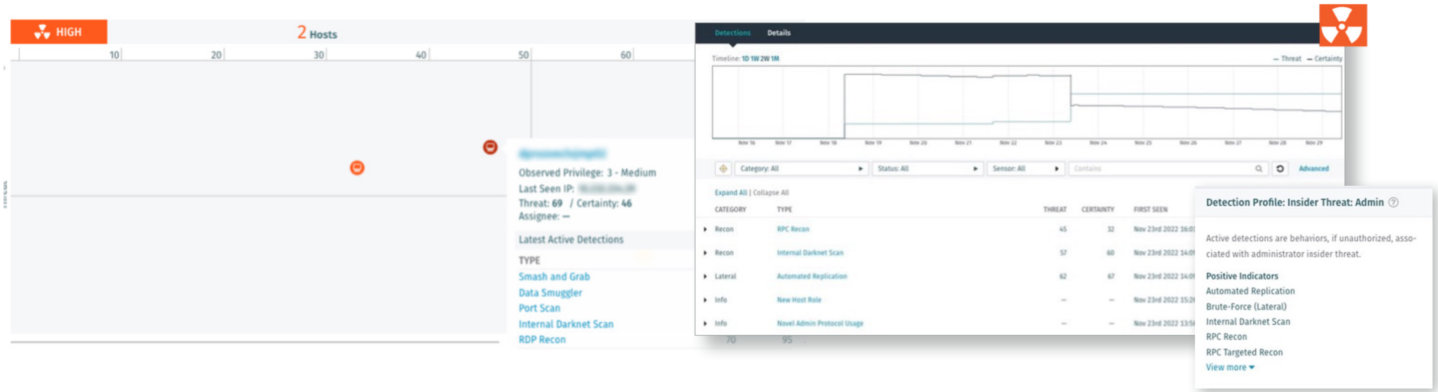

[redacted]
Observed Privilege: 3 - Medium
Last Seen IP: [redacted]
Threat: 69 / Certainty: 46
Assignee: —

Latest Active Detections

TYPE
Smash and Grab
Data Smuggler
Port Scan
Internal Darknet Scan
RDP Recon

CATEGORY	TYPE	THREAT	CERTAINTY	FIRST SEEN
Recon	RPC Recon	45	32	Nov 23rd 2022 16:0
Recon	Internal Darknet Scan	57	60	Nov 23rd 2022 14:0
Lateral	Automated Replication	62	67	Nov 23rd 2022 14:0
Info	New Host Role	—	—	Nov 23rd 2022 15:2
Info	Novel Admin Protocol Usage	—	—	Nov 23rd 2022 13:5

Detection Profile: Insider Threat: Admin

Active detections are behaviors, if unauthorized, associated with administrator insider threat.

Positive Indicators
Automated Replication
Brute-Force (Lateral)
Internal Darknet Scan
RPC Recon
RPC Targeted Recon

The Power of Security-driven AI: Knowing what is malicious

Detecting a *known* bad actor is quite novel with signature-based capabilities. With the exercise in this scenario taking on the form of a sophisticated adversary, traditional mechanisms fall short. Detecting the *unknown* threat has always been the challenge in cybersecurity. How do we protect against what we don't know exists? Patching, vulnerability management and Zero Trust principals only get us so far.

Following the low-and-slow approach, the bad actor was able to live in the environment for a week before the host and privilege account activity moved the host into the medium quadrant. At that time, the actor was not doing anything of significance, just learning the environment, looking for their next move, and maintaining persistence.

However, once our bad actor made a move, Vectra detections were triggered (much to the surprise of our attacker who thought they would completely evade detection). We can see within the outputs that while the user did nothing that was in an existing signature, their ultimate combination of successive behaviors including the Novel SMB and Protocol usage, Automated Replication, scans, ultimate RPC Recon and Bruce Force movement — together triggered the alerts. In the end, the exfiltration of 4Gb of data triggered a 'smash and grab' activity.

Behind the action

Within the attack simulation exercise, the bad actor leveraged multiple Linux and Windows hosts to perform their tasks. The exercise ultimately landed three of the hosts in the high quadrant and one of the hosts in the medium quadrant. This is enough to where any analyst or SOC team would be taking immediate investigation or possible automatic remediation within their SOAR platform.

The Vectra platform was operationally monitoring the entire infrastructure fabric of physical switching via a passive tap along with vSensors on the virtualized environment. No threat intelligence was leveraged for the exercise, and only Kerberos, LDAP and Active Directory integrations were provided to enrich the algorithm learning.

As the bad actor moved their attack simulation forward, the accounts associated with the users also began to be flagged as 8/10 within the Privilege Use Anomaly algorithms. In a real-world scenario like this one, a system-wide view must be had to correlate users, hosts, and potential GovCloud and SaaS workloads.

Ultimately, running an actual use case like this is necessary to prove our innovation. There are many promises of single solutions to solve all the challenges facing our cyber operators and defenders that are just that, promises. A coordinated and orchestrated set of capabilities clearly provides specific missions with the right set of capabilities. Underlying all the mission partner's needs, is a message that we must leverage the power of AI for the heavy lifting of a tier-1 and tier-2 analyst. Without it, we cannot scale to properly protect our national security.

About Vectra

Vectra® is the leader in cyber threat detection and response for hybrid cloud. Vectra's patented Attack Signal Intelligence™ detects and prioritizes threats across public cloud, SaaS, identity, and networks in a single platform. Vectra's Attack Signal Intelligence goes beyond simple anomaly detection to analyze and understand attacker behavior. The resulting high-fidelity signal and deep context enable security operations teams to prioritize, investigate and respond to cyber attacks in progress sooner and faster. Organizations worldwide rely on the Vectra platform and MDR services to stay ahead of modern cyberattacks.