SentinelOne

The Art of Threat Hunting

Putting Theory into Practice

Threat hunting is a proactive cybersecurity practice where security analysts use experience and ingenuity to search for, identify, and eradicate advanced threats that evade existing security measures. In many organizations it represents the best opportunity to uncover the most sophisticated threat actors.

Threat hunters are detectives, methodically searching for clues using a combination of tools, telemetry, threat intelligence, and knowledge of adversaries' tactics, techniques, and procedures (TTPs). To many, threat hunting may seem like a dark art, to be practiced only by the most skilled security practitioners. The reality is much less intimidating: successful threat hunting requires just a few tools, a little know-how, and the willingness to roll up your sleeves.

Step 1: Survey your hunting grounds

Do you have the basic tools you need for effective threat hunting?

Do you have the right data at your disposal? An effective hunt requires data - lots of it.

Foundational Data

- Endpoint
- Network
- Authentication
- Threat intelligence feeds

Rich Data

- Command line logs
- Access logs
- Email
- Cloud infrastructure
- API and Application logs

High-volume Data

- NS
- DHCP
- Netflow

- **How far back does it go?** Some threats lurk for months, even years, before being discovered. Keeping more data, for longer periods, unlocks more productive retrospective hunts.
- **What facilities do you have to query the data?** Do analysts need special skills to build complex structured queries, or do your repositories support more intuitive interfaces, such as natural language queries?
- **Can you join and query multiple different types of data at once?** Connecting logs from different sources, such as network logs with endpoint processes, can make suspicious behavior stand out more clearly.
- **Can you generate statistics and trends over time?** When a rare behavior becomes frequent, it may be time to investigate why.
- **How is performance?** The more queries you can run, the more hunting questions you can answer.

Step 2: Choose your target

What are you hunting for? Threat hunting involves developing and then testing a series of hypotheses around how an attack might slip through your defenses.

If your hypothesis was true, what traces would the attacker leave? What can you query within your data sets in order to surface these signals? Examples might include:

Hypothesis: Ransomware

A known ransomware actor is lurking in my environment.

Potential signs of a ransomware actor:

- Execution of scripts from temp directories
- Unusual use of built-in Windows admin tools such as "wmic", "net", and "query"
- Usage of encoding/decoding tools such as "certutil"
- Execution of Sysinternal tools
- Suspicious manipulation of scheduled tasks
- Signs of credential dumping using Mimikatz or "procdump"

Hypothesis: Insider Threat

An insider is collecting large volumes of data to prepare it for exfiltration.

Potential signs of insider data staging:

- Abnormal data access patterns
- Increased usage of compression or encryption tools
- Communication with unsanctioned cloud services
- Usage of alternate data transfer methods such as file sharing sites, peer-to-peer networks, or physical media
- Installation of unauthorized software

Hypothesis: Webshell

A threat actor has deployed a webshell and is using it as a network backdoor.

Potential signs of webshell activity:

- Modifications to config files such as "index.php", ".htaccess" and "web.config"
- Executable content such as .php .asp, and .jsp, in unexpected locations in web directories
- Unexpected outbound network traffic
- Excessive use of "POST" requests in application logs
- Error logs indicating failed access attempts, including syntax errors or access denied errors

Step 3: Unleash the hounds!

Now it's time to develop and execute our hunting queries.

Engineering threat hunts requires a different mindset from engineering automated detections. Threat hunts should pull in as many potential results as feasible, to maximize the likelihood of surfacing a stealthy threat.

Step 4: Sift through the results

Are there any true positive findings in your query outputs? Time to roll up our sleeves and investigate.

Separating benign behaviors from suspicious ones can often be accomplished by applying progressively more strict filters to the query output, eliminating layers of benign results, until you're left with nothing but potential threats that merit deeper investigations

Example: Our threat hunting efforts have uncovered hundreds of examples of the use of the Sysinternals tool "psexec".

- **Summarizing these results by time**, we see that the majority of these happen on a regular cadence, associated with known scheduled tasks. Filtering these out leaves 50 that remain under investigation.
- **Summarizing these results by user** shows that many are triggered by known administrative accounts for remote software installation and system maintenance. Filter these out leave 5 that remain under investigation.
- **Examining the command lines** reveals that 3 of the remaining instances were executed by a developer who is performing troubleshooting in a staging environment.
- **The remaining two instances** show lateral movement from a user's laptop to a file server during off hours. This highly suspicious result merits some immediate follow-up by security analysts.

Step 5: Sharpen your tools for the next time

Every threat hunt is an opportunity to learn and improve. Key questions to ask include:

- Is there additional data that would introduce more context that would make the hunt more efficient in the future?
- Can the hunting queries be refined to produce more signal and less noise next time?
- Was query performance satisfactory? Are there opportunities to tune performance to reduce the time to generate results?
- Can this hunt be automated for the future? Did the hunt uncover any TTPs that can be built into autonomous detections?

Effective threat hunting requires arming your defenders with the right data and a platform that makes it accessible with minimal friction. SentinelOne's Singularity™ Platform and Singularity™ AI SIEM provide a platform that operates with unprecedented speed and nearly infinite scale to make threat hunting a practical reality.

Ready for a Demo?

Visit the SentinelOne website for more details,
or give us a call at +1-855-868-3733

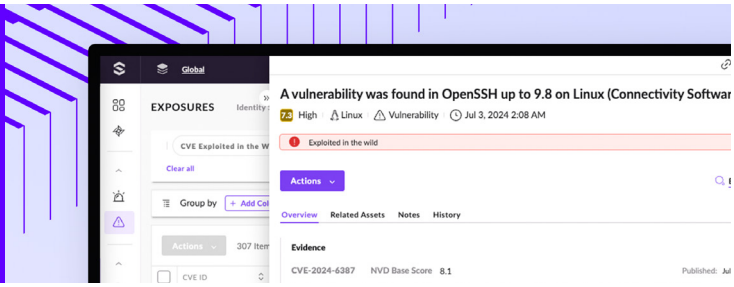

Innovative. Trusted. Recognized.

Gartner

A Leader in the 2023 Magic Quadrant for Endpoint Protection Platforms

MITRE ENGENUITY

Record Breaking ATT&CK Evaluation

+ 100% Protection. 100% Detection
+ Outstanding Analytic Coverage, 4 Years Running
+ 100% Real-time with Zero Delays

Gartner Peer Insights

96% of Gartner Peer Insights™

EDR Reviewers Recommend SentinelOne Singularity

FedRAMP · TEVORA PCI DSS Attestation HIPAA Attestation · AICPA SOC · STAR Level One · VB100 VIRUS · SE Labs Best Innovator Winner 2021 · AAA · Trusted Cloud Provider CSA

About SentinelOne

SentinelOne is the world's most advanced cybersecurity platform. The SentinelOne Singularity™ Platform detects, prevents, and responds to cyber-attacks at machine speed, empowering organizations to secure endpoints, cloud workloads, containers, identities, and mobile and network-connected devices with intelligence, speed, accuracy, and simplicity. Over 11,500 customers—including Fortune 10, Fortune 500, and Global 2000 companies, as well as prominent governments—all trust SentinelOne to Secure Tomorrow.

sentinelone.com

sales@sentinelone.com
+1 855 868 3733