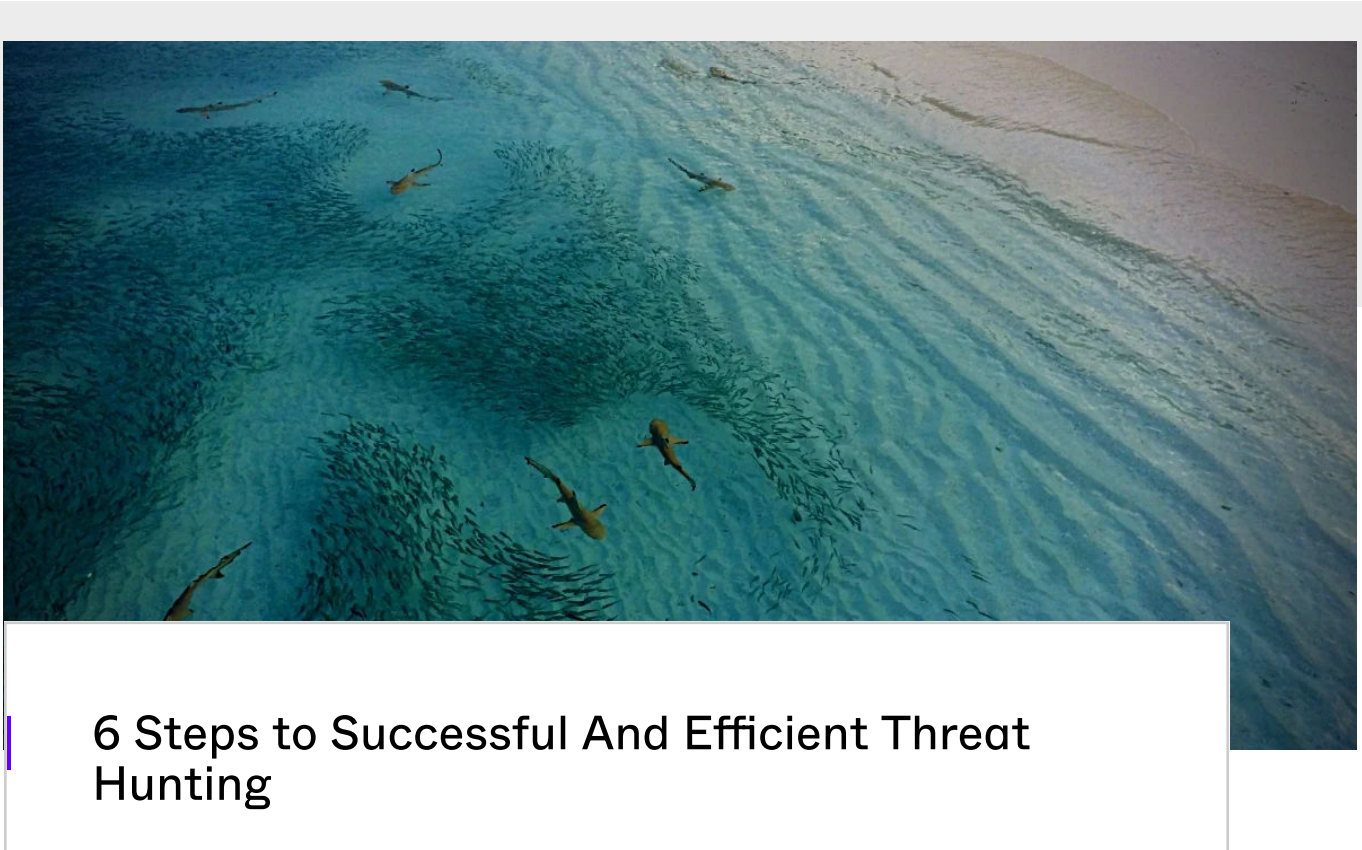

6 Steps to Successful And Efficient Threat Hunting

January 19, 2021
by Resha Chheda

Cybersecurity often feels like a game of cat and mouse. As our solutions get better at stopping an attack, adversaries have often already developed and started utilizing new tactics and techniques. According to Verizon DBIR, advanced threats lurk in our environment undetected, often for months, while they stealthily look to gather valuable information to steal or data to compromise. If you wait until these threats become visible or an alert is generated by traditional SOC monitoring tools, it can be too late. Threat hunting can help combat these challenges. Rather than waiting for an alert, threat hunters proactively assume that an advanced adversary operates inside the network and operates to find their existence.

In this post, we discuss threat hunting, why it's essential, and how you can enable your team to adopt efficient hunting strategies with the SentinelOne Platform.

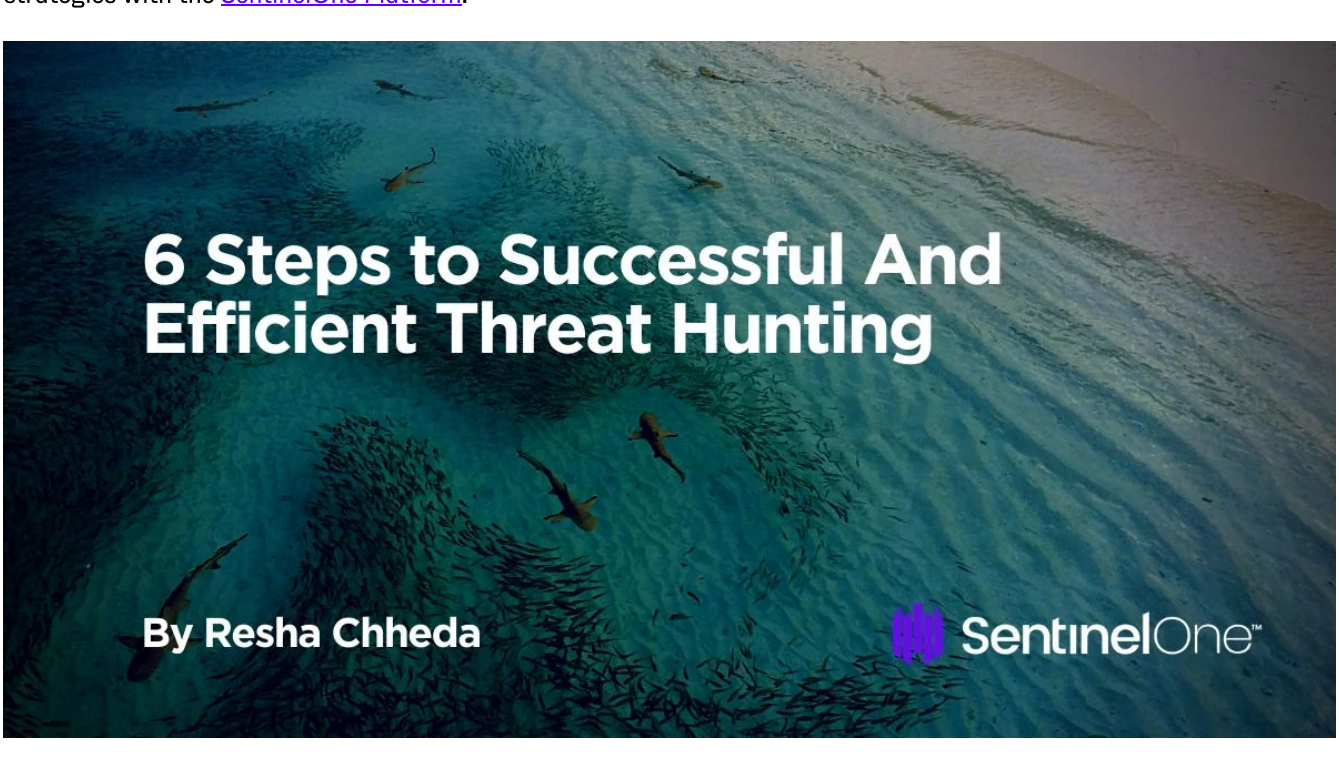

What is Threat Hunting?

Threat hunting has been defined by some as a *"computer security incident response before there is an incident declared"*. Others define it as *"threat detection using the tools from incident response"* or even *"security hypothesis testing on a live IT environment."*

> We define threat hunting as *the process of searching across networks and endpoints to identify threats that evade security controls before they can execute an attack or fulfill their goals.*

Rather than simply relying on security solutions to detect threats, threat hunting is a proactive approach to finding threats hidden in your network.

Unlike the Security Operations Center (SOC) and Incident Response (IR) teams, threat hunters not only respond to threats; they actively search for them. This process involves making hypotheses on the existence of potential threats, which are then either confirmed or disproven on the basis of collected data and analysis.

Threat hunting is also quite a different activity from either incident response or digital forensics. The purpose of DF/IR methodologies is to determine what happened after a breach was discovered. In contrast, when a team engages in threat hunting, the aim is to search for attacks that may have already slipped through your defensive layers.

Threat hunting differs from penetration testing and vulnerability assessment, too. These attempt to simulate an attack, ask questions such as what 'could' happen if someone compromised my security. Whereas threat hunters work from the premise that an attacker is already in the network and then look for indicators of compromise, lateral movement, and other tell-tale artifacts that may provide evidence of the attacker.

Why Do You Need To Incorporate Threat Hunting?

On average, cybercriminals spend 191 days inside a network before being discovered, and that's more than enough time to cause some damage.

Simply stated, if you aren't looking for threat actors inside your network, you may never know they are there. What if the attackers lock you out of the systems before you notice that you are under attack? With an efficient threat hunting program, you don't have to stress over such possibilities.

Threat hunting is human-driven, iterative, adaptive, and systematic. Hence, it effectively reduces damage and overall risk to an organization, as its proactive nature enables security professionals to respond to incidents more rapidly than would otherwise be possible. It reduces the probability of an attacker being able to cause damage to an organization, its systems, and its data.

Threat hunting also reduces your reliance on external vendors that may not know your network or normal employee behavior as well as your threat hunting team might.

Finally, threat hunting will force you to learn your networks, systems, applications, and users.

Understanding all of these components is a critical element of a robust security framework.

Six Steps To Creating An Efficient Hunting Program

So how do you create a perfect and efficient hunting program? Well! In reality, the perfect hunting program rarely exists! You need your hunting program to be an iterative combination of processes, tools, and techniques continually evolving and adaptive to suit your organization. Here are six steps that will help you create an efficient threat hunting program in your organization.

1. Ensure You Have The Right Data.

No data, no hunt! Period!

All successful threat hunting begins with having the right data to answer the right questions. Without the right data, you will not be able to conduct a successful and meaningful hunt. You need to ensure you have telemetry that captures a wide range of activity and behaviors across multiple operating systems and which can serve as a base for all your threat hunting efforts. Device telemetry should include data like network traffic patterns, file hashes, processes, user activity, network activity, file operations, persistence activity, system and event logs, denied connections, and peripheral device activity.

Just having the raw data is not enough; you also need to ensure that **you have context surrounding the data**. Knowing which data to combine, correlate, or extend is critical. Ideally, you want tools that allow a clear overview of all the above data with powerful capabilities to automatically contextualize and correlate different events into unified detections that minimize the amount of manual sifting through raw logs.

> SentinelOne patented Storyline^{TM} technology provides analysts with real-time actionable correlation and context and lets security analysts understand the full story of what happened in your environment.

Each autonomous SentinelOne Agent builds a model of its endpoint infrastructure and real-time running behavior. Every element of a story has the same Storyline. This gives you the full picture of what happened on a device and what caused it to happen. SentinelOne automatically correlates related activity into unified alerts that provide Campaign Level Insight. This reduces the amount of manual effort needed, helps with alert fatigue, and significantly lowers the skillset barrier of responding to alerts.

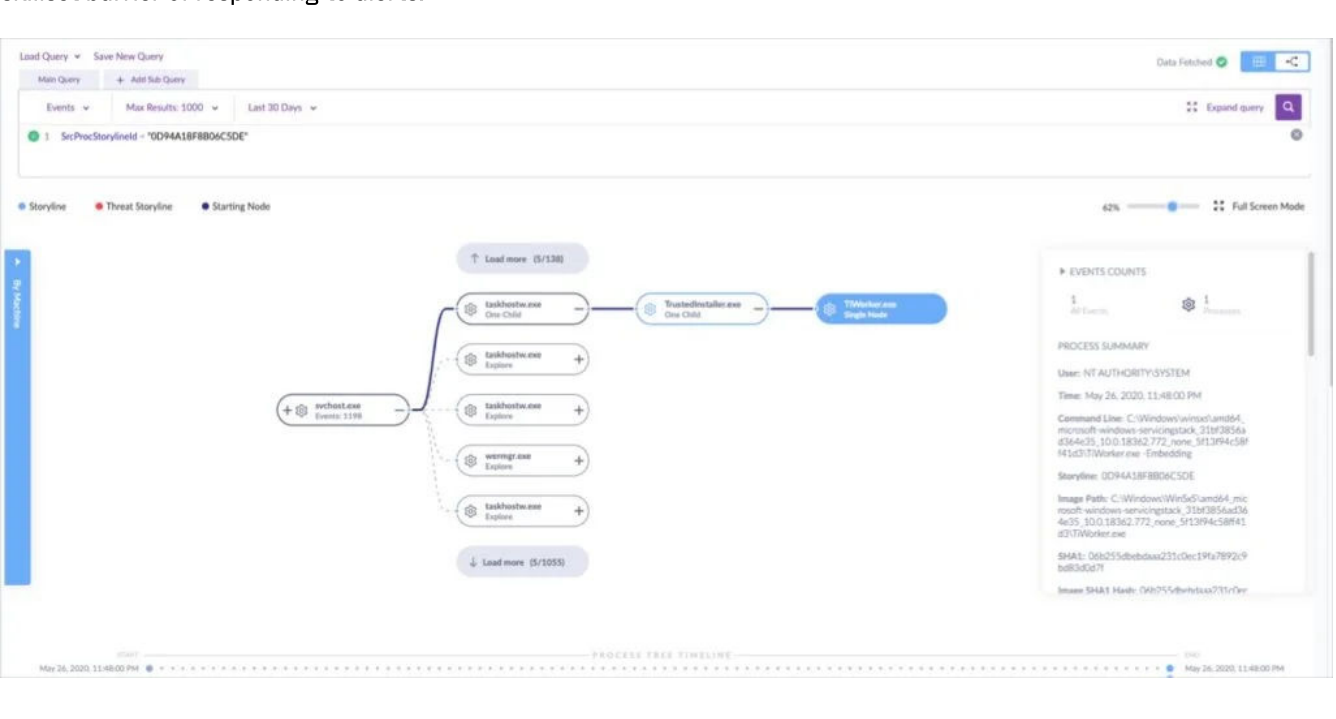

2. Baseline To Understand What's Normal In Your Environment

Threat hunters need a solid understanding of the organization's profile, business activities that could attract threat actors, such as hiring new staff or acquiring new assets, and companies.

A critical component of threat hunting is having the data to baseline 'normal' and find outliers (outlier analysis). Attackers will often want to blend in with ordinary users to acquire user credentials from phishing campaigns, so understanding a user's typical behavior is a useful baseline for investigating anomalous file access or login events.

Combining that with understanding what company data is of value to attackers and where it is located can lead to creating hypotheses such as *"Is an attacker trying to steal data located at a specific location?"* This, in turn, could prompt data collection that answers questions like: *"Which users have accessed that location for the first time in the last n days?"*

> SentinelOne's behavioral AI engine leverages advanced data science methods to teach systems the difference between regular day-to-day operations and actual threat behavior.

This provides the analyst with the complete picture and any additional context needed to help them understand what normal looks like and enable them to spot any outliers. An alert is triggered if a pattern emerges, such as repeated login attempts from a country that is not the usual norm in your environment, which may indicate a potential brute force attack. This helps make threat detection and hunting faster and more accurate. SentinelOne also retains historical data from 14 days to 365+ days, available to query in near real-time, so that the hunting team can understand and analyze data over large periods of time.

3. Develop A Hypothesis

Many hunts start from an intel source that uses Indicators of Compromise (IoCs), hash values, IP addresses, domain names, network or host artifacts provided by third-party data sources such as Information Sharing and Analysis Center (ISAC) or the FBI. Hunts can also be incident driven; given any incident, you need to answer how and when it happened. However, not all threats are known. In fact, a large number of threats are unknown, so hunting cannot solely rely on utilizing known methodologies.

In a hypothesis-driven workflow, a hunt starts with creating a hypothesis, or an educated guess, about some type of activity that might be going on in your environment. Using Open-source intelligence (OSINT) tools and frameworks like MITRE ATT&CK works effectively if you know what you are looking for.

That brings us to one of the essential components of threat hunting: hypothesis formation and testing. Hypotheses are typically formulated by hunters based on tools and frameworks, social intelligence, threat intelligence, and past experiences. Generalized questions could include, *"If I were to attack this environment, how would I do it? What would I attempt to gain access to? What would be my targets?"*. Other examples could include questions like *"Why do I see encrypted HTTPS, FTP traffic to countries in the East, in my environment?"* or *"Why do I see an abnormal volume of DNS queries from a single machine?"*

Ideas can be derived from the following sources:

- **MITRE ATT&CK framework:** a vast knowledge base of attack tactics, techniques, and procedures. Studying the MITRE techniques and their simulation in test environments can serve as a foundation for developing hypotheses.
- **Threat Intelligence reports**: contain useful information about attack techniques and procedures based on real incidents. Systematic analysis of such reports should spark some thought and give rise to many threat hunting ideas.
- **Blogs, Twitter, and conference talks**: information about new attack techniques appears for the first time via research blogs, and conferences, even before the attackers start actively using it. The timely study of such information will allow threat hunters to be proactive and prepare before the new attack technique becomes widespread.
- **Penetration testing**: attackers tend to use tools similar to those applied by experienced pen testers. Therefore, studying pen-testing practices creates a treasure trove of knowledge for generating threat hunting hypotheses.

> SentinelOne's patented Deep Visibility lets you quickly and iteratively query and pivot across endpoint telemetry captured from endpoint devices to validate hypotheses.

SentinelOne automatically correlates all related objects (processes, files, threads, events, and more) of a threat. For example, suppose a process modifies a different process by injecting code. When you run a query, all interaction between the source process, target process, and parent process shows clearly in the cross-process details. This lets you quickly understand the data relationships: the root cause behind a threat with all of its context, relationships, and activities. Analysts can also leverage historical data to map advanced threat campaigns across time to enable efficient hypothesis generation.

You can create powerful hunting queries with easy-to-use shortcuts. As a threat hunter, the MITRE ATT&CK framework has likely become one of your go-to tools. SentinelOne makes hunting for MITRE ATT&CK Tactics, Techniques, and Procedures (TTPs) fast and painless. It's as easy as entering the MITRE technique ID and using this to perform a hunt.

> SentinelOne provides a query library of hunts using data from various open, commercial, and bespoke sources curated by SentinelOne research.

These hunts are the output of hypotheses that are proven across research data and are generic. For example, the use of unmanaged, unsigned PowerShell is likely abnormal in most environments and would commonly require additional investigation. Both of the above examples are not malicious in and of themselves but fit in a hunting workflow as they are descriptive of anomalies.

4. Investigate & Analyze Potential Threats

After generating the hypothesis, the next step is to follow up on it by investigating various tools and techniques to discover new malicious patterns in the data and uncover the attacker's TTPs. If the hypothesis is correct and evidence of malicious activity is found, then the threat hunter should immediately validate the nature, extent, impact, and scope of the finding.

Although threat hunting starts with a human-generated hypothesis, threat protection tools, like SentinelOne, make the investigation more efficient. SentinelOne's Deep Visibility empowers rapid threat hunting capabilities thanks to Storyline. Each autonomous SentinelOne Agent monitors endpoint activity and real-time running behavior. A Storyline ID is an ID given to a group of related events in this model. When you find an abnormal event that seems relevant, use the Storyline ID to quickly find all related processes, files, threads, events, and other data with a single query.

With Storyline, Deep Visibility returns full, contextualized data that lets you swiftly understand the root cause behind a threat with all of its context, relationships, and activities revealed from one search.

> Storyline allows threat hunters to understand the full story of what happened on an endpoint and enable them to see the complete chain of events, saving time for your security teams.

5. Rapidly Respond To Remediate Threats

Once you uncover a new TTP, you need to make sure you can effectively respond and remediate the threat.

The response should distinctively define both short term and long term response measures that will be used to neutralize the attack. The main goal of the response is to immediately put an end to the ongoing attack to prevent the system from damage by a perceived threat. But it is also essential to understand the cause of the threat to improve security and prevent attacks of a similar manner in the future. All necessary steps must be taken to ensure that similar attacks are not likely to happen again.

SentinelOne enables analysts to take all the required actions needed to respond and remediate the threat with a single click.

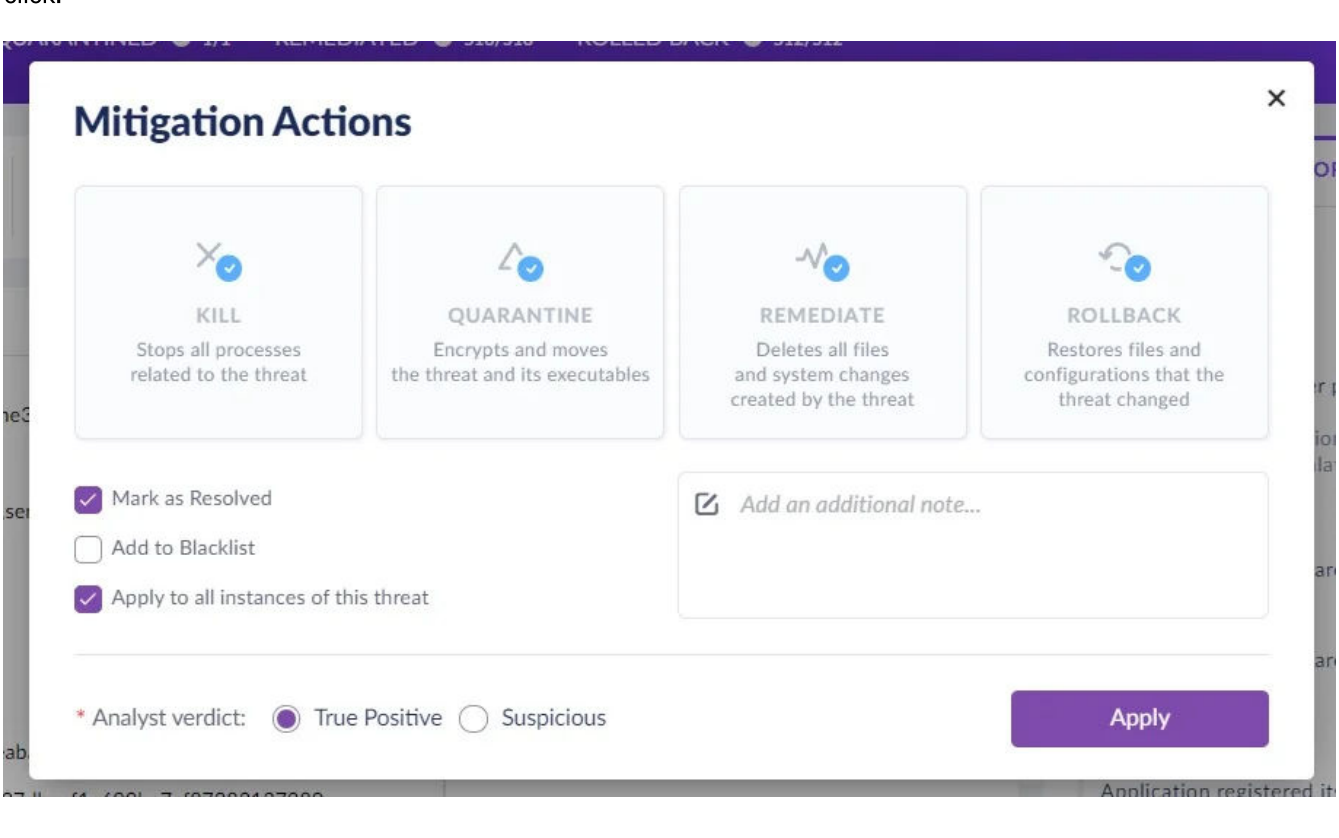

With one click, the analyst can rollback the threat or perform any other available mitigation actions. Rollback functionality automatically restores deleted or corrupted files caused by ransomware activity to their pre-infected state without needing to reimage the machine.

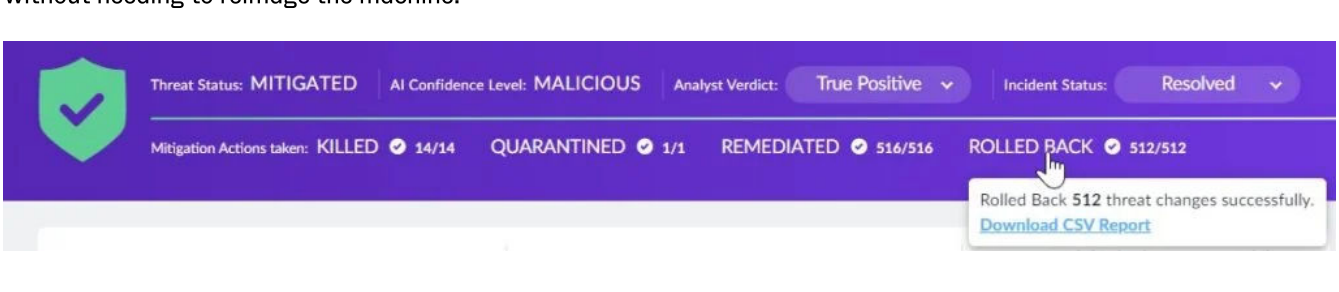

The threat can be added to Exclusions, marked as resolved, and notes can be added to explain the rationale behind the decisions taken. SentinelOne also offers full Remote Shell capabilities to give your security team a quick way to investigate attacks, collect forensic data, and remediate breaches no matter where the compromised endpoints are located, eliminating uncertainty and significantly reducing any downtime that results from an attack.

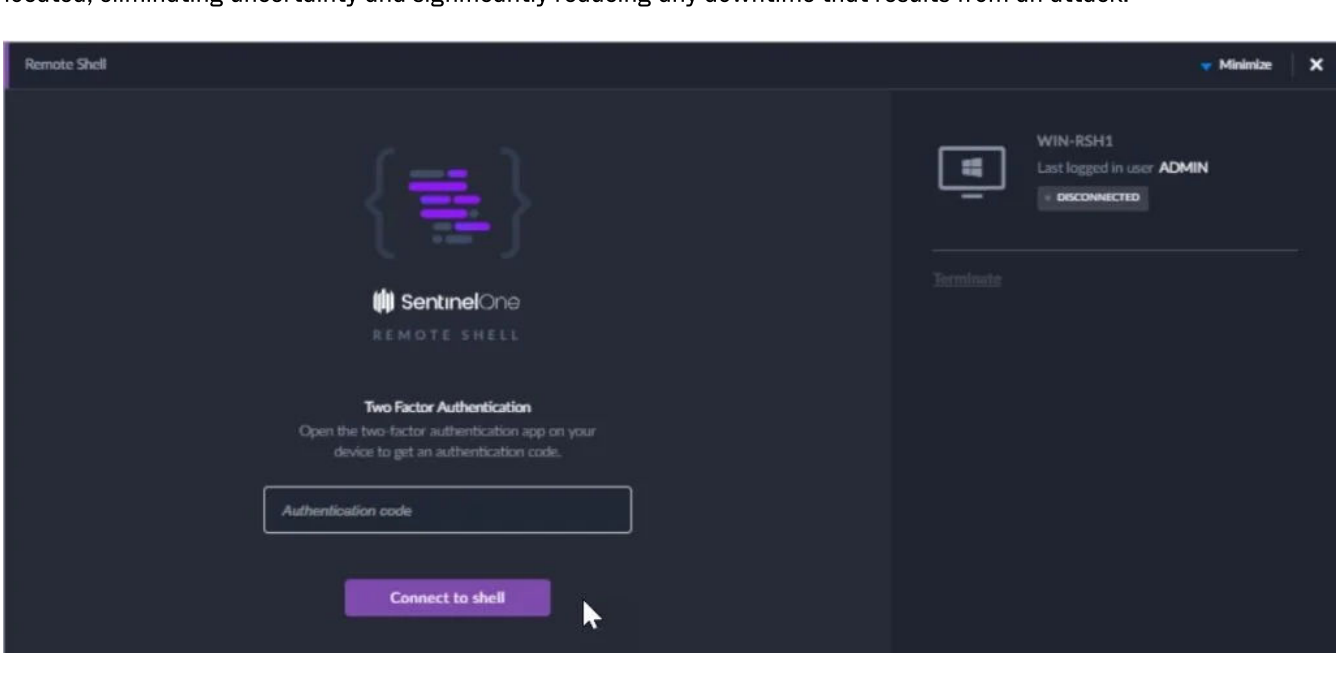

SentinelOne also can detect threats in advance through the aid of its machine learning and intelligent automation. It can anticipate threats and attacks by deeply inspecting files, documents, emails, credentials, browsers, payloads, and memory storage. It can automatically disconnect a device from a network when it identifies a possible security threat or attack.

6. Enrich And Automate For Future Events

Finally, successful hunts form the basis for informing and enriching automated analytics. The final step in the threat hunting practice is to use the knowledge generated during the threat hunting process to enrich and improve EDR systems. This way, the organization's global security is enhanced thanks to the discoveries made during the investigation.

Advanced threat hunting techniques will try to automate as many tasks as possible. Monitoring user behavior and comparing that behavior against itself to search for anomalies, for example, is far more effective than running individual queries. However, both techniques are likely to be required in practice. Both are made easier if you have tools like SentinelOne with a rich set of native APIs enabling full integration across your security software stack.

> SentinelOne is designed to lighten the load on your team in every way, and that includes giving you the tools to set up and run custom threat hunting searches.

With Storyline Auto-Response (STAR) custom detection rules, you can turn Deep Visibility queries into automated hunting rules that trigger alerts and responses when rules detect matches. STAR gives you the flexibility to create custom alerts specific to your environment that can enhance alerting and triaging of events

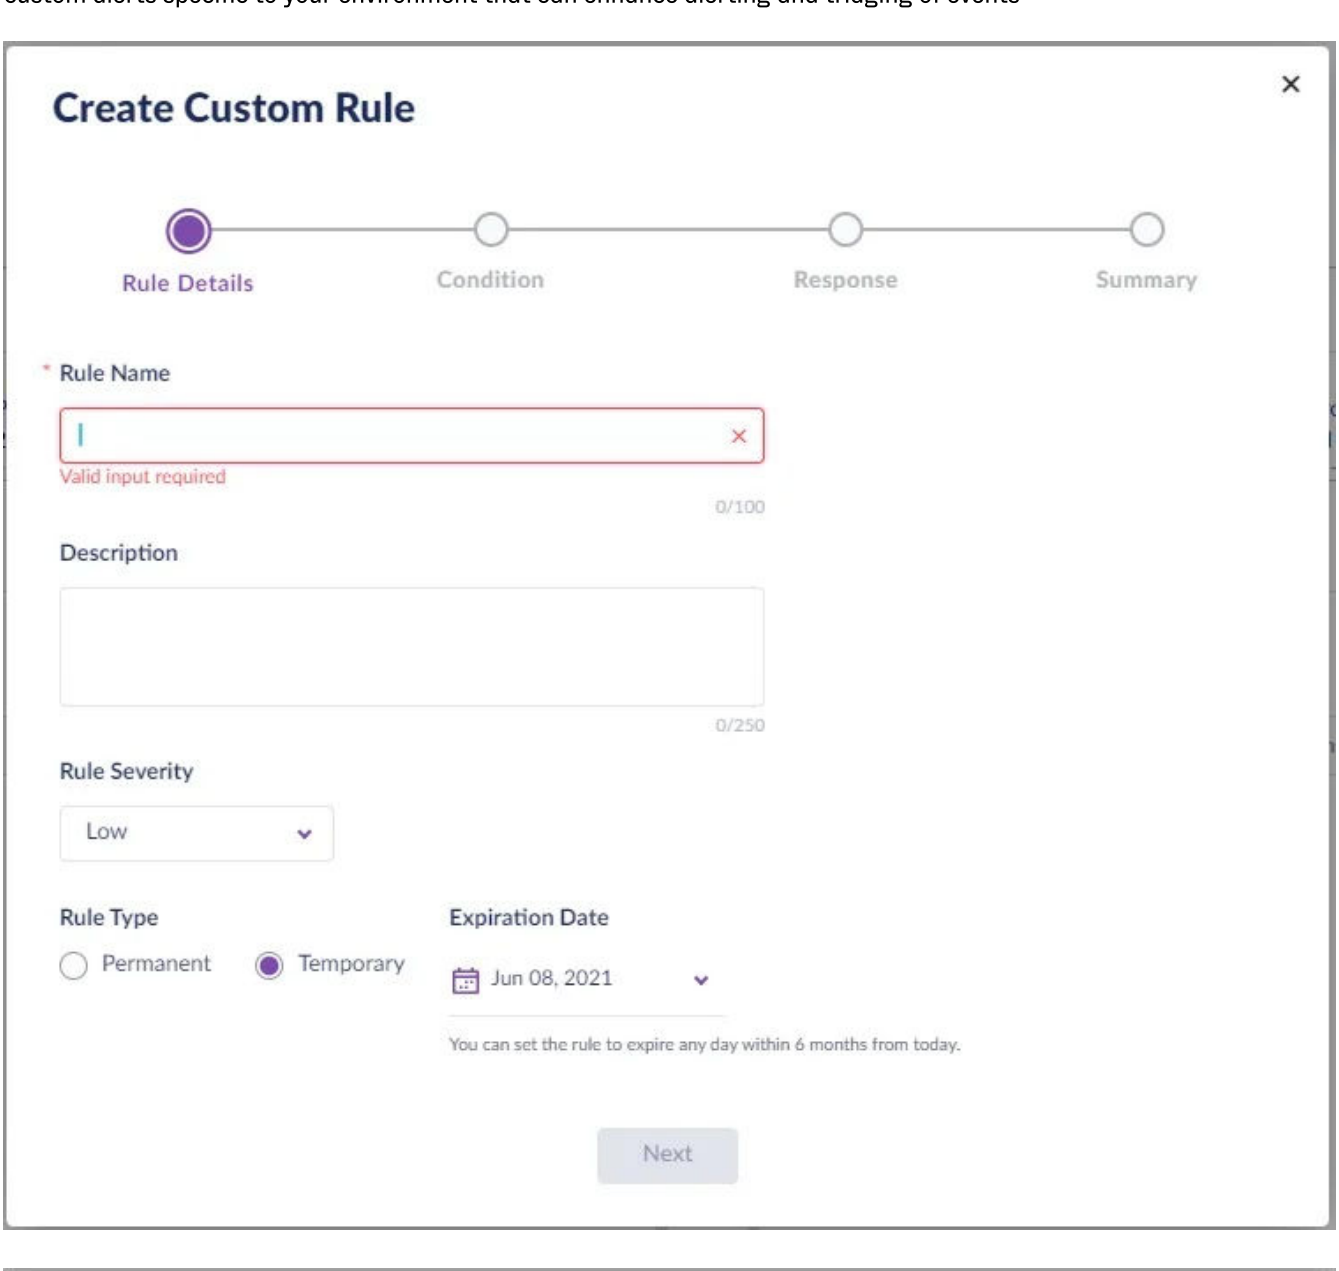

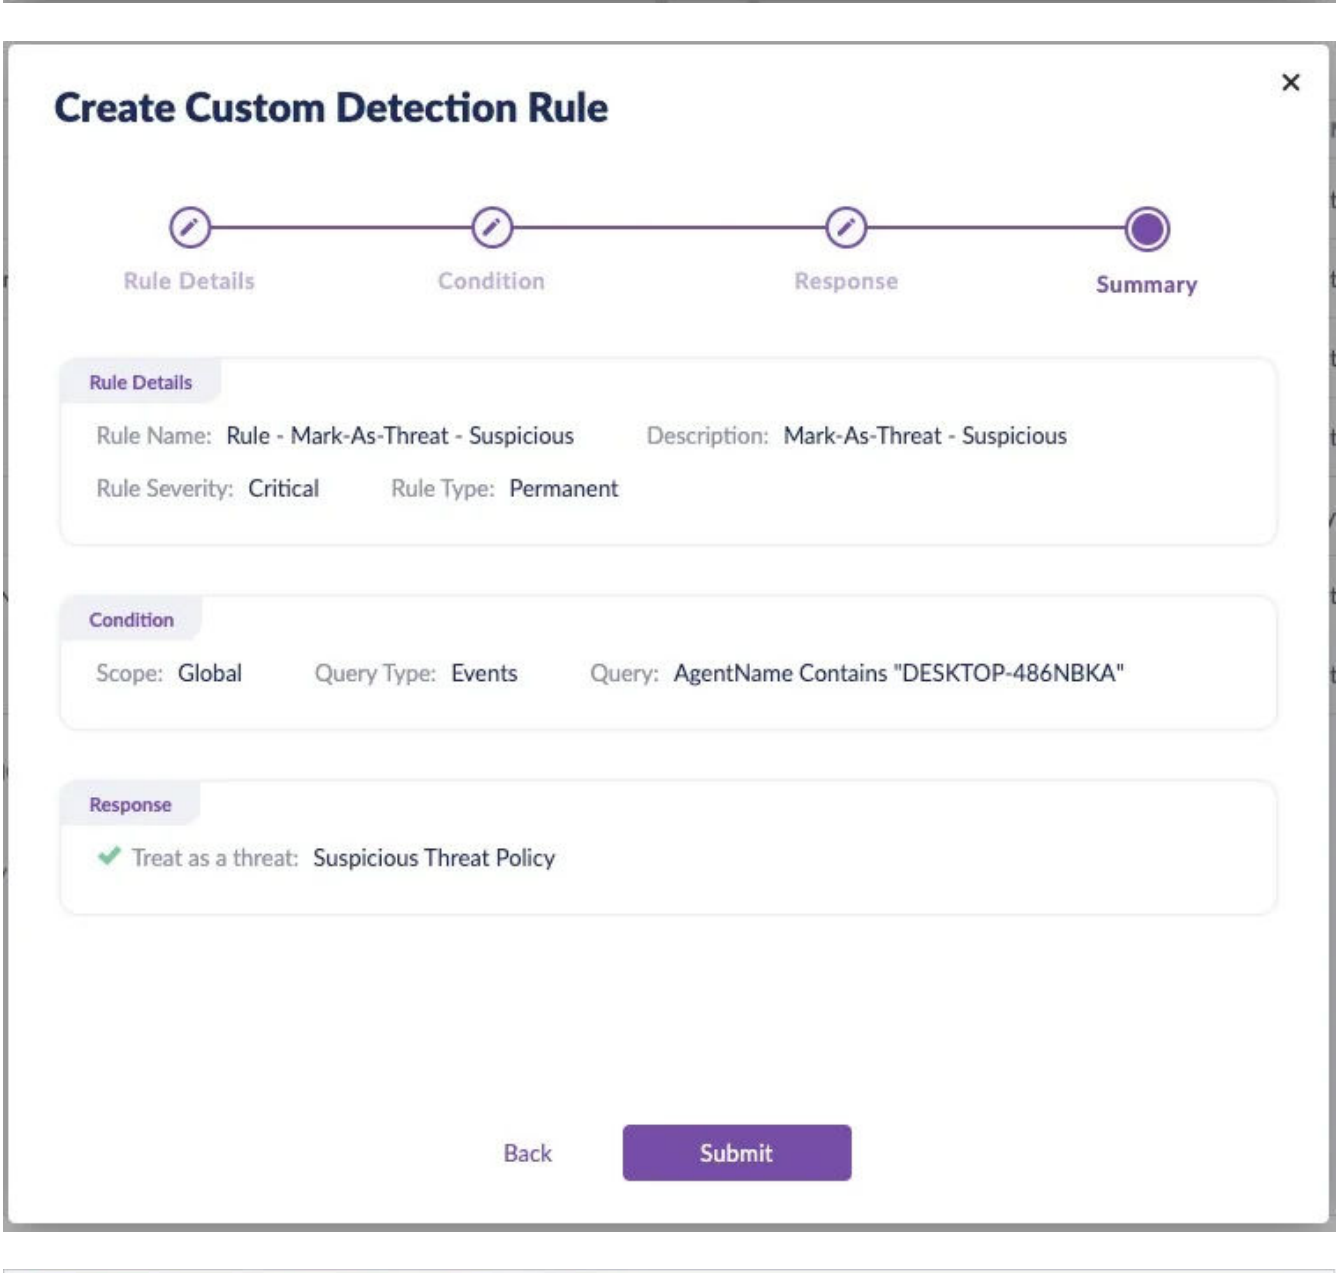

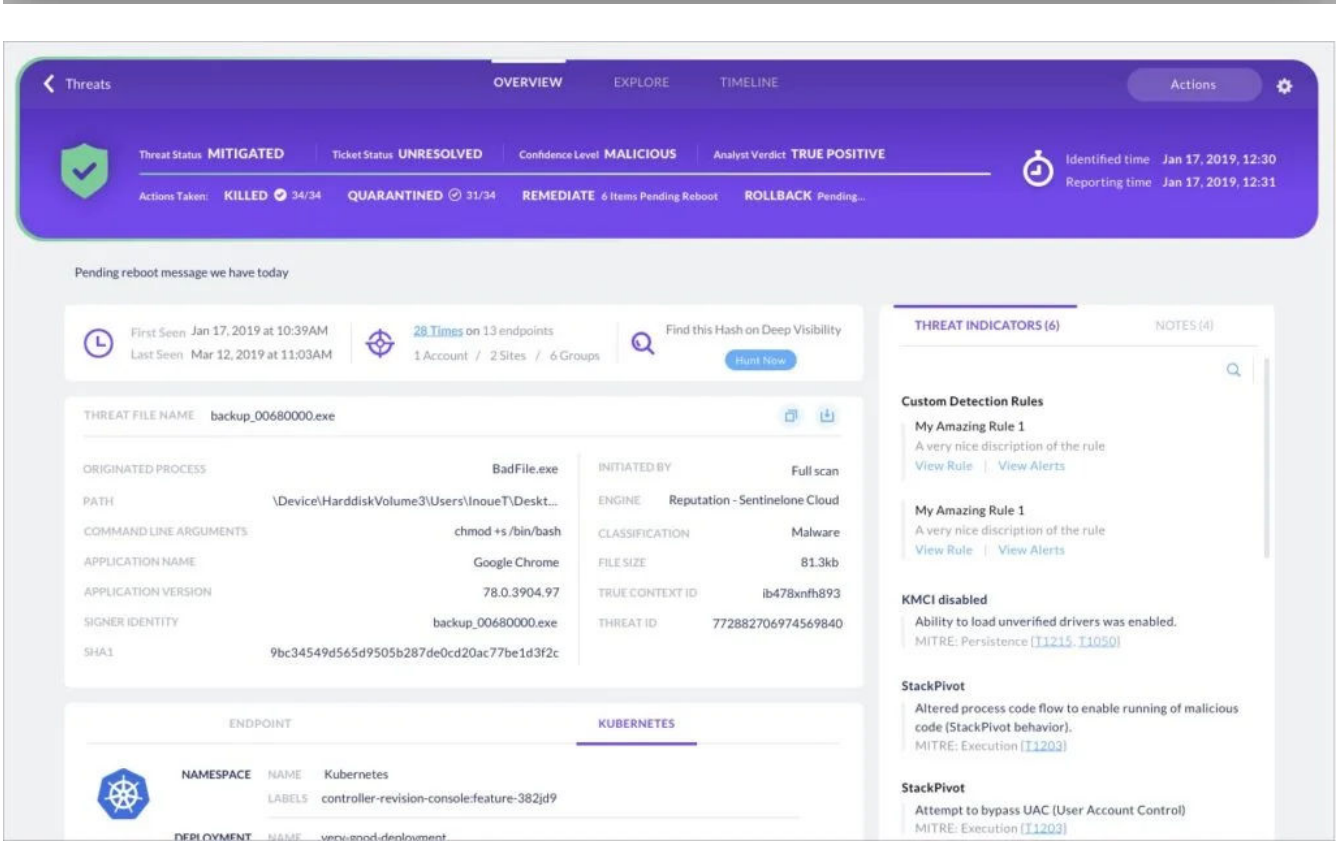

SentinelOne can also automatically mitigate detections based on the policy for suspicious threats or the policy for malicious threats or can put endpoints in Network Quarantine. Alerts are triggered in near-real-time and show in the Activity log in the Management Console. You can enable alerts in Syslog that can be used for triage and SIEM integration.

After running the query in Deep Visibility and investigating, you can select an Auto-Response for the rule to automatically mitigate the rule detections. With that, you have set your SentinelOne solution to automatically protect your environment, according to your needs, from every threat, every second of every day. Modern adversaries are automating their techniques, tactics, and procedures to evade preventative defenses, so it makes sense that enterprise security teams can better keep up with attacks by automating their manual workloads.

Closing Thoughts

Implementing a threat hunting program can reap many benefits to the organization, including proactively uncovering security incidents, faster Incident Response times, and a more robust security posture. Effective threat hunting needs to result in less work for your busy analysts while at the same time future-proofing your SOC from a variety of known and unknown adversaries. SentinelOne gives you visibility, ease of use, speed, and context to make threat hunting more effective than ever before. Please contact us or request a demo to see how SentinelOne can help you develop an efficient hunting program.

Additional Resources

Deep Dive – Hunting with MITRE ATT&CK
Use the S1QL Cheatsheet For Security Analysis
Learn more about Rapid Threat Hunting with Storyline
Visit SentinelOne Platform page
Visit Sans Threat Hunting Report – Automating Hunt
Read Gartner Report about Using Threat Hunting for Proactive Threat Detection

---

Read more about Cyber Security

- Stopping Cyberattacks on Remote Workers Starts at the Endpoint
- 7 Ways Hackers Steal Your Passwords
- Insights from the CyberLaw Forum | Intersecting Cybersecurity, Insurance & Regulation
- Unify the Analyst Experience with Singularity Operations Center
- EDR vs Antivirus: What's the Difference?
- What Are TTPs? Tactics, Techniques & Procedures – Inside the Mind of a Cyber Attacker