POLARITY

Advancing Threat Hunting

Surveying the Landscape with Augmented Reality

Summary

Threat Hunters' primary initiative is to anticipate and mitigate potential malicious threat actors that lurk beneath organizational security systems. Successful threat hunting relies first on the difficult process of hypothesizing threat actor motives based on environmental knowledge. After a hypothesis is proven, threat hunters begin the process of analysis, an often time-consuming and inefficient process due to the number of queries and data entry required.

The key to improving this process is to enable threat hunters with tools that mitigate human error and compensate for human limitations like our inability to multi-task, retain large amounts of data, and share information instantaneously within teams. With tools like Polarity, threat hunters are able to use more advanced techniques like Glassing, which approaches threats with a broader view, taking contextual information into account.

Background

The cyber defense strategies adopted by organizations are often evolving to meet the challenges placed upon them by internal and external business drivers. Chief among those drivers are threat actors who have consistently demonstrated an ability to lurk within organizational networks for excessive periods before their detection and eventual expulsion.

These repeated occurrences of excessive adversarial dwell time have fueled the idea that to effectively manage an information security program, leaders must acknowledge that fundamental cyber solutions such as firewalls, endpoint protection, and anti-malware products are vulnerable to motivated threat actors.

This acknowledgment has contributed to the adoption of practices affectionately referred to as Cyber Threat Hunting, ("Threat Hunting" or "Hunting") within the cybersecurity community.

We can define threat hunting as a focused and iterative approach to searching out, identifying and understanding adversaries internal to the defender's networks (Lee, SANS).

This practice is usually performed based on a hypothesis, rather than triggered by an event or observations within the defender's network.

Cyber threat hunting can take lessons from techniques used by game hunters, in particular the lesser-known technique of glassing. Glassing is employed by more experienced game hunters and relies on the forgotten technique to survey a given landscape. From a vantage point, the entirety of the landscape is evaluated for the slightest indication of a target, without sacrificing the ability to narrow in the field of view with a scope.

Preparation

Successful threat hunting first depends on thoroughly understanding the operating environment (Long, 2016). Similar to traditional hunting, those responsible for cyber hunts will recognize greater successes if they understand their surroundings or operational environment. Further, it is not enough to know simply what exists within the environment, but how the orientation and interconnection of entities within the environment impact one another. Threat hunters need to be aware of systems available within the target environment, their functions, interconnections, intended configurations, and the value of those systems to a threat actor.

As personnel resources are limited, hunts are commonly targeted within environments where intelligence suggests potential threat action (e.g. actor targeting credit cards with specific tradecraft) or negative outcomes of environmental compromise are so immense, that investment into a hunt within that environment should be made without such intelligence (e.g. crown jewels). This can aid in the formulation of hypotheses prior to hunt execution.

Preparation Challenges
On a hunt, familiarization of the operational environment is key. Hunt teams often "parachute" into highly prioritized environments. While the organization may collect logs, feed analytical platforms, and utilize visualizations to support analysis, much of the value of such tools can be nullified if the members of the hunt team have no working understanding of the operating system producing the logs, or the business function and features of an application environment subject to the hunt exercise.

Formulating a Hypothesis

Cyber threat hunting is a relatively new concept within the framework of an information security program. Even the largest and most sophisticated organizations are still developing their hunt processes and lack the proper staffing to execute hunt operations to the degree that their leaders feel is necessary. As such, environments are targeted, leveraging a risk-based approach and threat intelligence to formulate a hypothesis about where bad actors may be lurking, and to what end or value, a compromise of the environment is for the bad actor.

Value to a threat actor may be one of the most difficult things to anticipate, theorize, or develop a hypothesis around.

This is because different threat actors recognize different aspects of the network as high-value targets. For example, in a scenario where an organization's analytical system is being targeted by a threat actor, one such actor may be interested in (1) stealing the data processed by the system, another in (2) disrupting the integrity of the system's output and another in (3) stealing the algorithm(s) running against the data set.

Should hunt team members fail to consider these (and other attributes), the hypothesis formulated by the team may be overly broad, under-resourced, and eventually yield little to no value to the organization compared to what could have been realized - had appropriate factors been considered.

Hypothesis Challenges
Tapping into intelligence sources, ensuring their currency, and making the best use of them in the formulation of hypotheses can be cumbersome. Often, this will include a review of historical incidents to determine realized or observed adversary activity and a review of historical vulnerability information, identifying the coverage of existing security controls that might mitigate certain threat activity or otherwise create detectability of such activity.

Log Collection and Consolidation

Following the development of a hypothesis, members of the hunt team will need to acquire the data necessary to prove or disprove the hypothetical scenario. Sources of data could include but are not limited to, the logs generated by network appliances, security appliances, native operating systems, and database applications.

Logs from system to system will vary in availability, reliability, and usability. Unavailability of logs is probably one of the most common obstacles encountered by hunt teams and incident responders. While in isolation, log unavailability can be an indicator of malicious activity, it is more commonly a direct result of failures in log management and IT/IS governance - leaving matters of reliability and usability to contend with. The reliability of log sources should be evaluated prior to placing reliance on such logs during a hunt or an investigation. In order to validate the reliability of the log sources, the hunt team may need to conduct inquiries into the log sources directly. This could include a review of historical accesses, changes to log files/repositories, and timeline analysis of logs for gaps in coverage. Depending on the size and scope of the hunt, validating the reliability of log sources could consume significant resources before the hunt effort is underway, detracting from the mission defined by the hunt team.

Experienced hunters know that their efforts are better served by augmenting their approach with the technology that allows them to be most efficient. In the later phases of a hunt, specifically during investigation/analysis phases, importing logs into a common platform for analysis may be paramount.

Collection / Consolidation Challenges
Many log sources, specifically legacy or custom environments, don't port easily to centralized log management or analysis platforms. This is compounded by the fact that many high-value targets within enterprise environments are either legacy or custom systems. Once again, hunters often lose valuable time when normalizing logs to be imported into analytic platforms, or formatting for manual analysis.

In extreme cases, log data may not be eligible for export. This condition will necessitate real-time log review or live forensic capture by an analyst.

Investigation / Analysis

The hunt team will evaluate collected logs and the outputs of analytical processes within the context of the initial hypothesis to determine if an actual threat has been realized. The analysis process should include the following steps:

- Confirm if a threat has been realized
- Evaluate the extent to which it has been realized (scope and magnitude)
- Establish a timeline of events
- Determine the overall impact

Often, a mandate to operate in accordance with a risk-based approach as well as within budgetary limits confines the execution of a hunt to a limited scope. This is not a desired condition. A limited scope could translate to the evaluation of information within such a narrowed view, or anchored in bias, that it results in failure to identify malicious activity as its nature was not a direct corollary to the scope of the hunt. Simply put, a team may be hunting for (1) Mrs. Silver, in the (2) study with the (3) candlestick, but within a narrow scope, the threat actor (Mrs. Silver) could go undetected because she's using a different weapon (dagger).

To be successful in analysis, hunt team members must collectively position the motivations of all manner of threat actor at the forefront of their mind, and establish a mechanism for understanding the various tools, techniques, and processes leveraged by these actors.

Investigation / Analysis Challenges
Reviewing at logs in isolation and relying on manual analysis alone can be cumbersome and ineffective (Lee, Lee 2016). Hours of monotonous lookups, queries, and data entry reduces the quality and speed of human decision making leading to mistakes of habit. Further, quality pattern recognition degrades to cognitive shortcuts to clear the queue of "false positives".

Enhanced Hunt Operations

Assuming more hunts are to be performed in the future, hunt team members should strive to enrich their existing data stores and technologies with the information and intelligence that they've gathered during hunt and incorporating lessons learned.

Successful hunts form the basis for informing and enriching automated analysis. A team should not waste valuable time doing the same hunt over and over. If an indicator or pattern is identified that could have the potential to recur in the target environment, teams should automate that indicator's detection so that they can continue to focus on the unknown(s). Information from hunts can be used to improve existing detection mechanisms, which might include updating analytical algorithms, SIEM rules, or detection signatures. The more a team knows about its own network, the better it can be defended. As such, it's invaluable to record and leverage new discoveries as they are made on a hunt.

Challenges in Enhancing the Hunt
One of the biggest challenges to enhancing operational delivery is disseminating valuable information learned from historical hunts (and cross-team operations).

Up until now, there has been no effective mechanism allowing for tactical intelligence to be annotated, enriched, and disseminated across functional teams in such a way that it is available in a real-time fashion to analysts on their desktops as they are conducting investigations.

Furthermore, once tactical intelligence has been developed, it is often stored in static repositories which must be manually queried by operational analysts. This manual retrieval and recall of intelligence has the unfortunate effect of slowing or stalling the process of attack identification even in situations where actionable intelligence already exists.

Surveying the Landscape with Polarity

Glassing is a lesser-known technique employed by more experienced game hunters that relies on forgotten tradecraft to survey a given landscape. From a vantage point, the entirety of the landscape can be evaluated for the slightest indication of the target, without sacrificing range from the narrow view of scope as preferred by many game hunters.

The expanded field of view offered by glassing is made possible via traditional tools such as binoculars, tripods, and simple awareness of environmental conditions.

Within the context of a cyber hunt, glassing involves stepping away from narrowly scoped investigative analysis and evaluating information from a longer and wider view.

Cyber analysts that employ this technique can extract the most value out of information at their disposal and as a result, delivering maximum value from each hunt. However, in order to leverage this technique without sacrificing focus, time and/or resources, an abundance of contextual information must be available to the hunt team instantaneously, so long as it is relevant and valuable to the hunt.

The availability of this contextual information is enabled through Augmented Reality. In the methodology highlighted in this document, there were several core challenges identified that spanned across the five phases of hunt. The following is an overview of how augmented reality can address these challenges.

Preparation
Before launching into a hunt operation, Polarity enables teams to instantaneously have working knowledge of all systems within scope for a hunt operation. Such knowledge can be accessed from historical hunts, human annotations, asset management solutions or Configuration Management Databases.

Further, these teams can operate with immediate knowledge of keywords, functions, usernames/user associations, service accounts, syntax descriptors, event ID values and translations etc. without having to leave their screen, break from the operation to engage the business or alter critical path to reference internal wikis and data sources.

Almost as soon as a hunt operation is approved, a hunt team using Polarity can integrate operational intelligence into the platform, and immediately draw upon its value by way of near-real-time situational awareness.

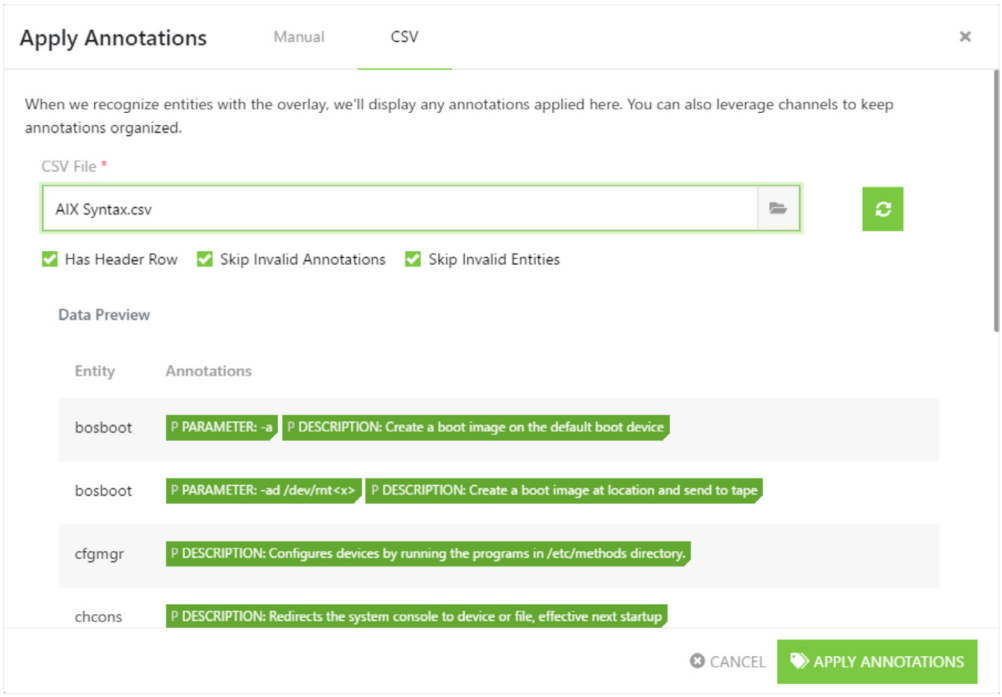

In the screen capture (left) the threat hunter has imported all AIX syntax into Polarity. Whenever the syntax is observed by the threat hunter or members of the hunt team in the future, the description of the syntax will be overlaid on their screen in real-time.

Hypothesis Development
A strong hypothesis cannot be developed without attempting to include known or anticipated independent and dependent variables. In the context of a cyber hunt, Polarity helps teams recover from inefficiencies associated with historical information gathering processes such as accessing ticketing systems, case/incident management platforms, obtaining historical vulnerability data, observing network diagrams, etc.

Further, Polarity helps to avoid breakdowns and intelligence failures that manual processes are prone to.

Once any member of the hunt team observes in-scope systems on their screen, environmental variables will become clear, collaboration will be enabled via the platform, and a strong hypothesis will follow.

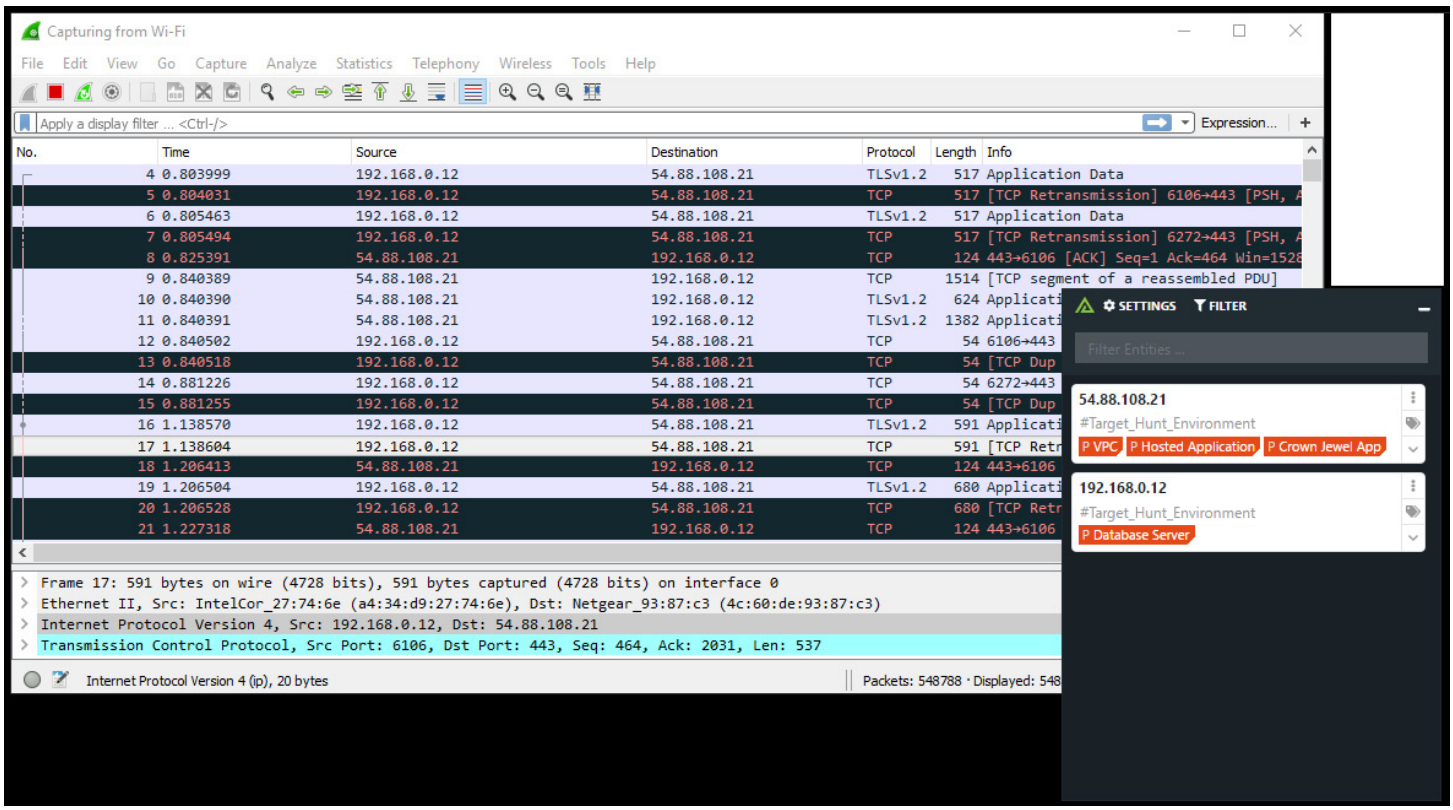

In this screen capture Wireshark information is overlaid from internal annotations or asset repositories, allowing the threat hunter to identify possible High Value Targets (HVT) within the environment.

Data Collection

Not all datasets are considered equal. True to the principles of glassing, an analyst may focus on areas that a target is most likely to occupy, but the whole of the landscape is evaluated for outliers. In some scenarios, analysts may opt to exclude data sets from more targeted analysis or opt to exclude the data sets in any analytical functions applied to the data in the context of a wider body of information.

These exclusions may be done as a result of consent, or as a byproduct of bias. Whatever the reason, Polarity can help give analysts assurance that data points displayed on-screen during manual review or analyzed in isolation can be compared to broader datasets, on the fly, during analysis.

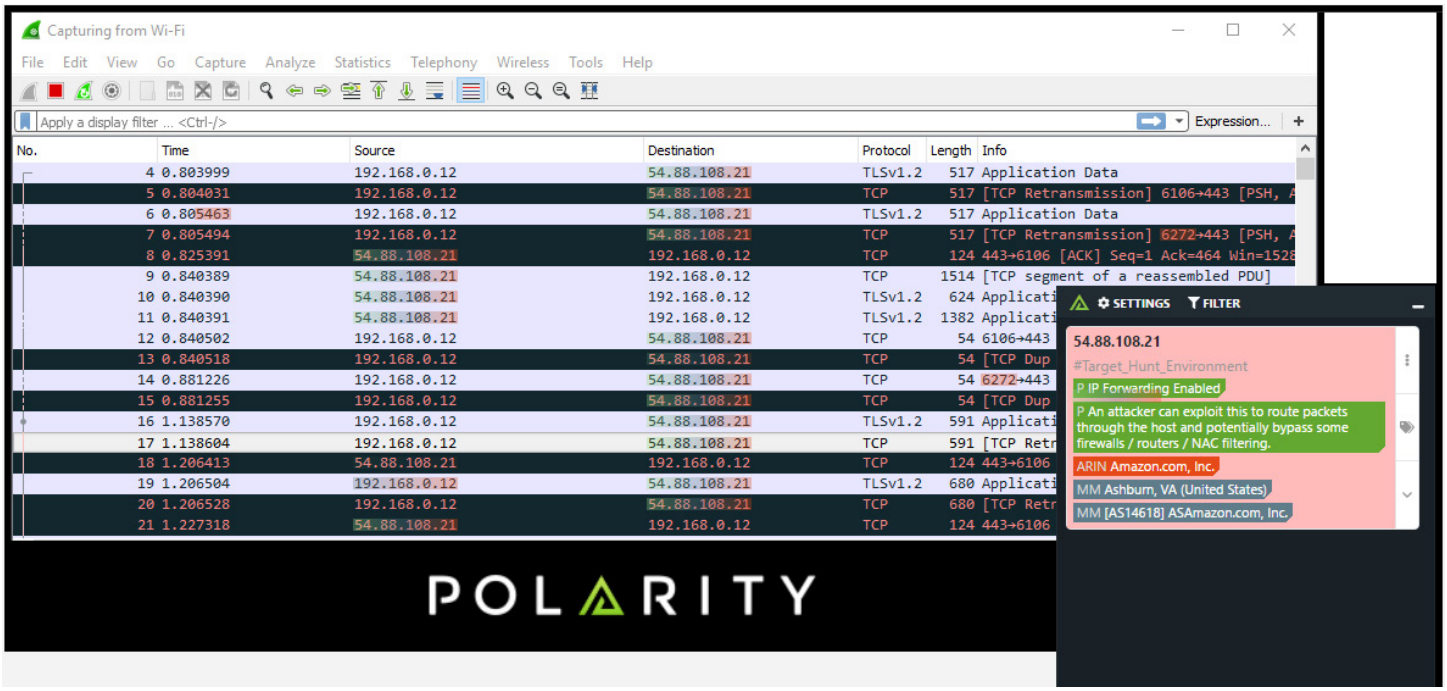

Investigation

Polarity automatically searches for and delivers relevant context to analysts as they are working. Analysts are less likely to miss critical intelligence because Polarity removes the burden of finding relevant contextual information. Since Polarity operates at the screen level, it enables collaboration across multiple applications, toolsets, and workflows. Analysts no longer need to choose between working fast and working thoroughly.

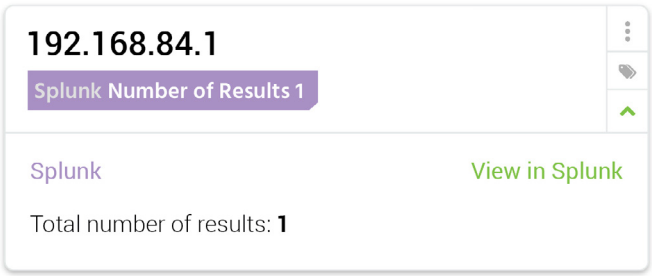

Polarity combats analyst fatigue by automating the most repetitive and time-consuming components of an analyst's daily workflow. **Reduced lookups and automatically delivered contextual data speeds up the decision-making process, letting analysts do analysis.**

Enhancement

Polarity allows for the efforts applied on a single hunt to be applied to future hunt operations. The experience and tradecraft collected by a team of skilled professionals can be leveraged to augment a modified team in the future, or a completely different team operating in another hemisphere.

The future of Threat Hunting

Imagine hunt teams capable of superhuman memory and armed with practical augmented reality. These teams would be able to apply historical information seamlessly to an operation, instantaneously access valuable intelligence sourced from connected datasets, share and collaborate across hunt teams, and effortlessly collaborate with front line network defenders. These teams could share high-value indicators or guidance on how to properly triage and/or escalate certain types/classes of alerts without performing manual lookups against datasets.

With Polarity's AR for your desktop, this imagined scenario is now a reality.

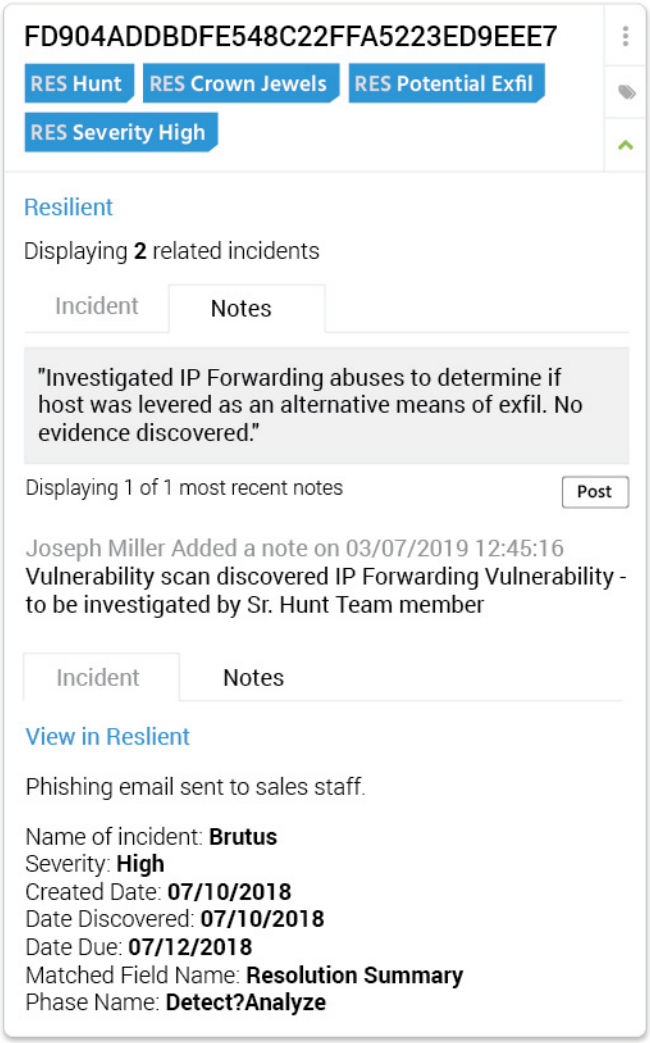

In this screenshot example, investigation notes posted to the Resilient platform ensures investigative teams have historical information easily at hand and enables them to update investigations without ever leaving the platform or tools they are working in

POLARITY

Augmented Reality for Your Desktop

References

Lee, Rob and Lee, Robert. M. (2017, April). The Hunter Strikes Back: The SANS 2017 Threat Hunting Survey https://www.malwarebytes.com/pdf/white-papers/SANS_Report-The_Hunter_Strikes_Back_2017.pdf

Long II, Michael. C. (2016, July). Scalable Methods for Conducting Cyber Threat Hunt Operations https://www.giac.org/paper/gsec/38852/scalable-methods-conducting-cyber-threat-hunt-operations/1527444