Demystifying Threat Hunting Concepts

January 5, 2017 by
Josh Liburdi

Demystifying Threat Hunting Concepts

This post is about demystifying threat hunting concepts that seem to trip up practitioners and outsiders. If the summary in the TLDR below seems appealing, then please continue to the meat of the post.

TLDR?

- Threat hunting doesn't have to be complex, but it's not for everyone
- Knowing how to begin and end a hunt is more important than knowing how to carry out a hunt
- If you need a place to start, look at trends in the threat landscape and focus on threats that you do not have automated alerts/detections for
- Hunting is a creative process that rewards those who take chances
- Finish with something, anything actionable — so long as it provides value

All set?

For a moment, forget what you know about established threat hunting procedures, processes, or models and consider it from a minimalist perspective: like any activity, threat hunting has a beginning, a middle, and an end. The beginning describes what you intend to hunt, the middle describes how you do it, and the end describes what you do with the results. Although this may be counter-intuitive to digital investigators, in my experience, effective hunting is not determined by the act of investigating (the middle), but by how one chooses to start (the beginning) and how one reacts to newly learned information (the end).

With respect to the maturity of security operations organizations and experience of practitioners who want to perform threat hunting, there is a concept of 'you must be this tall' (YMBTT) to ride the threat hunting rollercoaster. (Threat hunting is not for everyone *right now* , though it maybe in the future.) YMBTT is tied directly to the ability of an analyst to plan, conduct, and process the outcome of a hunt. It should be a goal of any threat hunting advocate to lower the YMBTT bar so inexperienced analysts can participate.

Where to start?

The act of beginning is possibly the most intricate step of an effective hunt. It requires an understanding of the target network, the capabilities of and tools used by a security operations team, and the capabilities of threats to the target. While each of these can individually be a catalyst for the beginning of a hunt, it's best to have a holistic understanding of each. Here are sets of questions that can get you started on establishing an understanding of each component.

- What is the layout of the network? What operating systems are running in the network? What tools and services are running on the operating systems? What (or where) are the critical assets in the network? (Use questions like these to determine what is normal and abnormal in the network.)
- What is the security operations team already looking for? What automated detection is in place and how precise is it ? (Use questions like these to determine what you can already find — **don't hunt for things you can already find** .)
- Which assets do threats target? What tools and tactics do threats use? How have threats behaved in the past? (Use questions like these to determine what an attacker might do in your network.)

Of these, the most useful to me are questions and understandings derived from threat intelligence. Breaking threat intelligence reporting (either external or internal) into various levels of detail can reduce the YMBTT requirements of using threat intelligence to hunt; one way to achieve this is to not focus entirely on technical indicators described in a report, but instead to derive a broader understanding of trends in the threat landscape. For example, in the context of establishing a hunt, knowing that…

> threat actor AAA used a PowerShell script with the filename BBB against a target in industry CCC on date EEE launched from IP address FFF

…may be less useful to me than the broader understanding of 'attackers continue to use PowerShell to conduct attacks.'

Carrying out a hunt

When analysts talk about threat hunting, they tend to focus on the middle of a hunt — the stage where one retrieves data, processes it, and investigates (see: this). Unsurprisingly, this is also the most well-documented part of hunting (see: this). I don't think there's much more to add here; if you are looking for ideas that drop you right into the hunting process, then #threathunting conversations on Twitter and the threathunting.net repository are great resources.

That mentioned, there are two concepts that aid me when carrying out a hunt that apply equally to borrowed ideas and new ones.

The first focuses on retrieving data for carrying out a hunt — this step is critically important because it directly impacts everything following it. Improper data retrieval can lead to poor data processing, can lead to incorrect results, can lead to misguided interpretation, can lead to mistaken corrective action; if appropriate data was not retrieved, then the entire hunting process is brought into question. This can be avoided through intelligent filtering and categorization of data — only retrieve the data required to complete the hunt. This is a step that one usually does not get right the first time and may require a cycle of retrieving data and triaging results to confirm the data is appropriate.

The second is less logical and possibly divisive — hunting is a creative process. One's abilities to think abstractly, challenge ideas, and be unafraid of failure lead to more knowledge and breakthroughs than someone who does everything the same way every time. The idea of creativity becomes hugely important when applying hunting techniques — with so many techniques to choose from (text-based searching, dozens of visualizations, endless permutations of machine learning algorithms), how do you learn which technique is most effective without trying out new ideas? Returning to the concept of YMBTT, creativity is not a trait inherent to every person (if you need some, ask Brian Eno and Peter Schmidt); it's important that techniques are well-documented and shared among a group so that everyone can benefit.

Now what?

Ending a hunt is the epitome of 'simple to understand, difficult to execute.' The goal is to expand threat detection coverage (either through the identification of new opportunities or improvements to existing ones) — without this, it's likely that hunting efforts will have no lasting impact on the organization. Difficulty in achieving this goal assumes that the utilized hunting technique did not provide suitable precision (because, by design, they almost never do). Generally, one has to deconstruct the results, generalize the procedure, and automate it into new or existing tools. Sounds hard, right? Not if you start with what you know and increase complexity as needed — hunts that lead to daily reports of vulnerable systems, identification of indicators of compromise, or new/improved intrusion detection system signatures are just as valid as hunts that lead to the creation of innovative threat detection tools, so long as the outcome provides value to the organization.