The 5 Levels of Autonomous Security:
What level are you?

Human decision-making can be sub-optimal at the best of times. We, most of us anyway, are not really "designed" to consider all the available data, factor probabilities and risks of each decision path and then make the optimal choice by balancing all those factors. However, humans are not expendable either. We have highly tuned abilities to recognize patterns[1], understand abstract relationships and generalize. For instance, we don't need hundreds or thousands of training samples to know that the user being targeted by this particular phishing attempt is the CFO for the organization. We recognize the name right away. The task at hand is to surface just the right information to enable operators to make optimal decisions that ultimately lead to positive business outcomes.

This is where the network comes in—it is foundational to producing and consuming all of that data, and intelligent network infrastructure provides the goldilocks balance of "just right" information. Delivering on this promise requires two components:

- A system to efficiently and in real-time collect the ground truth data.
- Intelligence to extract the information and context buried within the raw data.

Arista is uniquely positioned to deliver on both of these capabilities. Arista EOS® based on NetDL™, provides a multi-modal, multitenant-capable data lake that offers real-time network telemetry to other Arista solutions as well as those from our partners. Arista AVA uses an AI-driven approach to anticipate operator questions, extract answers from NetDL and deliver the insights necessary for effective human decision making. With this combination, Arista is providing networking for the data-driven enterprise.

AVA™, the world's first security decision support system, automatically pre-computes answers to investigative questions a highly skilled analyst would ask and surfaces the weak and early signals of an attack which delivers the context necessary to disrupt an adversary's objectives at the outset. AVA does this by leveraging the power of the cloud to scale without adding compute, storage and other rack requirements to the customer data center.

At the same time, the use of federated machine learning ensures customer-sensitive data never leaves their organizational infrastructure. Instead, the machine learning models train locally and only share normalized and anonymized data with the Arista cloud. Most importantly, Arista AVA is also directly accessible to end users and thus elevates the skills of an organization's existing security analysts. In this paper, we dive into the details of how this autonomous system works.

[1] https://arxiv.org/ftp/arxiv/papers/1801/1801.00631.pdf

An Analogy from the Highways

Automobiles today are faster, more fuel-efficient, fancier and safer than they have ever been. But one thing that has arguably gotten worse is the person in the driver's seat. It is no surprise that the National Highway Traffic Safety Administration's (NHTSA) analysis[2] shows that 94% of serious automobile crashes are the result of human error. So, automakers trying to improve safety are now focusing on eliminating the need for driver attention with innovations[3] that automate[4] a number of vehicular functions such as rain sensing wipers, automated headlights, blind-spot detection and collision avoidance systems. And as vehicles become more autonomous, not only does safety and reliability improve, but so does the driver experience.

Cruise control, for instance, was designed to eliminate the cumbersome act of keeping your foot on the accelerator. The problem with this "automation" is that the foot on the accelerator forces the driver to pay more attention. You take that away and rear end the car in front, going too fast into a bend, or hitting an unexpected obstacle are all potential consequences. These days adaptive cruise control (ACC) is becoming standard on more car models. Using radar sensors that detect other vehicles, dynamic speed setting and automated braking, ACC solves some of the challenges in Cruise Control 1.0 and reduces accidents.

It's helpful to think about the evolution in cruise control from something just automated to perhaps the beginnings of something autonomous. In fact, the Society of Automotive Engineers (SAE) has a standard for describing automation levels and places cruise control[5] at Level 0. ACC on the other hand is considered a Level 1 capability and Tesla's Autopilot or Cadillac Super Cruise are considered Level 2.

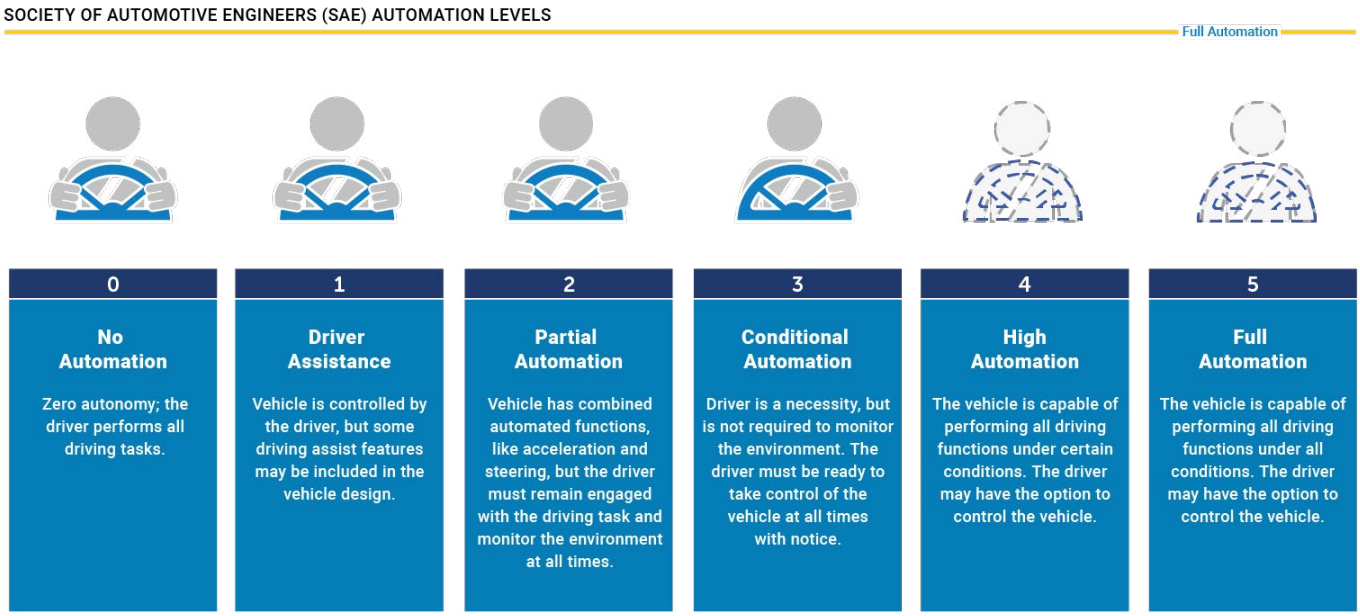

Figure 1: Standards prescribed by Society of Automotive Engineers (SAE)

[2] https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety

[3] https://www.usatoday.com/story/money/cars/2019/06/10/traffic-accidents-decline-because-of-car-features/1407386001/

[4] https://blog.nationwide.com/car-safety-timeline-infographic/

[5] https://www.caranddriver.com/features/a15079828/autonomous-self-driving-car-levels-car-level/

The Levels of Autonomous security

One of the basic forms of automation we consider standard today is the correlation performed by SIEMs and network security tools – e.g. collate all the "alerts" associated with an IP address together onto one screen or identify an "attack campaign" by grouping alerts that share a source or a destination. Some tools are a bit smarter and use additional sources of context such as Active Directory (AD) or threat intelligence, or filter out the "known good." But much like cruise control, there are a lot of unintended consequences that manifest in the security world primarily through false positives and negatives. For instance, we find the average device has more than half a dozen IP addresses over the course of the week and conversely the average IP address is associated with more than half a dozen devices in that same time frame. The point being any analysis that is based on an IP address is flawed from the get-go. And so, if cruise control is considered Level 0, would we consider this correlation Level 0 too? In fact, if you look at the degree of automation in security today, we might argue that the industry average is at best at the beginning of Level 1. The Security Orchestration, Automation and Response (SOAR) category could perhaps have the best claim to Level 2 – Partial Automation. These technologies automate several low-impact response and remediation tasks, for instance creating support tickets for the IT helpdesk, automatically correlating between the multiple security tools, grabbing evidence into an incident data store, computing metrics, triggering out of band notifications etc.

Back to Scheduled Programming

Using the SAE automation levels and replacing driver with analyst, how might we define a sliding scale for autonomous security?

0. No Automation

Zero autonomy; the analyst performs all triage, hunting & investigations.

1. Analyst Assistance

Most security is manual, but some analyst assist features maybe included in the toolset.

2. Partial Automation

The security program automates functions like response actions and policy enforcement, but the analyst must remain engaged due to the prevalence of false positives and negatives.

3. Conditional Automation

Analyst is a necessity and can take control at any time but high-fidelity detections, autonomous hunting, triage, investigations and response result in improved security and efficiency.

4. High Automation

All security tools can operate in an autonomous manner under certain conditions. The analyst focuses on defining and controlling policies that are then enforced by the technology.

5. Full Automation

All security tools operate autonomously under all conditions. The technology defines and enforces policies. The analyst can override these autonomous policies.

What Does Autonomous Security Look Like?

So, how do we enable "self-driving" security that empowers the analyst to focus on things best done by humans? Admittedly, getting to Level 4 and Level 5 will take the entire industry to raise its game. While that in itself is going to be challenging, perhaps the bigger complexity would be changing mindsets of people to be "ok" with that level of automation in a critical function like security.

So, for now, how about we just focus on getting to Level 3 – Conditional Automation?

Before we address how to do that it is necessary to understand the process an analyst goes through today, since that is what would need to be automated. When served up with an alert, the typical process includes:

1. Translate the IP address to a device and / or user – consult DHCP and DNS logs, CMDB, AD, etc. with the goal of identifying the impacted device and user.
2. Identify command and control / data exfiltration – review endpoint and network logs to identify domains / external IP addresses involved in the alert. Perform session and passive DNS analysis to identify related domains that may also be involved. Enhance context with threat and open source intelligence on the domains / IP addresses. The goal here is to identify an initial view of the attacker infrastructure.
3. Identify additional suspect external activity. For instance, identify other unique domains that appear to be accessed within the same time frame and disassociate domains that analysis shows are not actually related to this attack. Then enhance context with threat and open source intelligence on the resulting domains / IP addresses. The goal of this step is to get a broader understanding of the attacker infrastructure.
4. Identify lateral movement from the device to other hosts on the network using credential abuse (e.g. brute force or stolen credential usage) and protocols such as SMB, RDP, etc.
5. Identify other internal hosts that are accessing the identified attacker infrastructure. For each new host identified in steps 4 and 5 repeat steps 1 through 4.

We could go on, but you get the picture. In fact, if you plot out the decision tree an analyst must work through to simply understand the full scope of one suspected phishing threat, you could end up with the unreadable flow chart below.

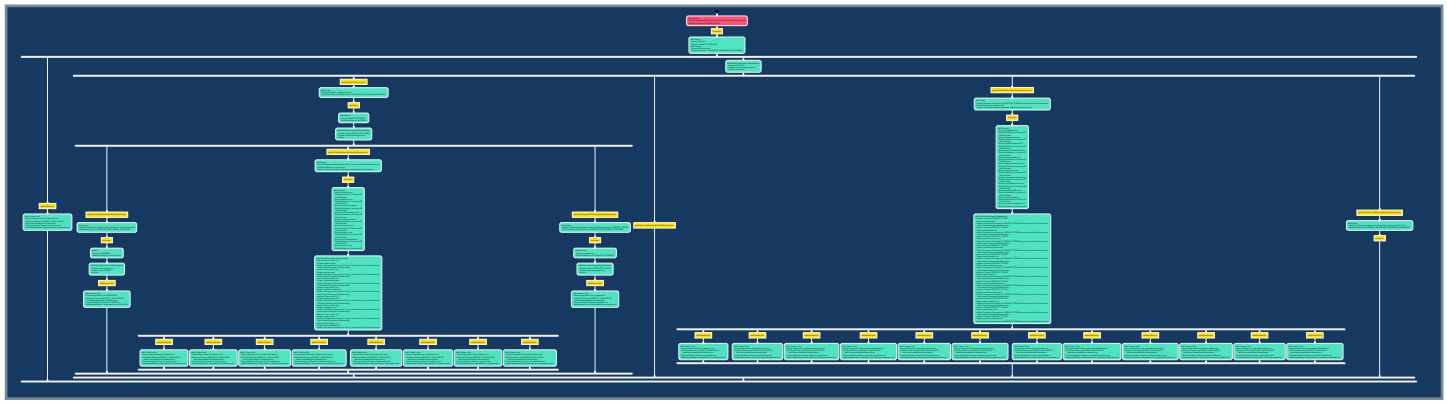

Figure 3: Phishing Investigation Decision Tree[6]

[6] *If you do want to read the unreadable flow chart you can find it at https://solutions.arista.com/Phishing-Decision-Tree.svg*

Could you achieve level 3 automation by scripting and plumbing your various security tools together? That's a fairly common question, and the answer is a qualified yes. But there are two problems with this approach.

First, you need to have the playbooks, automation scripts and integrations per decision tree, which quickly becomes unwieldy. But second, and perhaps more importantly, even if you managed that part of the process, implementing it requires connecting and maintaining integrations across multiple systems from disparate vendors.

It is a bit like buying an aftermarket adaptive cruise control capability for your existing car. Can it be done? Sure. Are there instructions and YouTube videos on how to do it? Sure. Is it reliable, cost effective and worth the trouble? Hmm.

Introducing Level 3 Autonomous Security with Arista

Let's go back to our analogy for a little bit. For something like Tesla Autopilot to work effectively, it needs to understand the vehicle (speed, which lane its driving in, navigation, braking and acceleration parameters, etc.) in the context of other vehicles sharing the road. It relies on a number of sensors with the goal of surfacing the data the driver needs to make a decision.

For example, what's the distance and relative speed of the vehicle ahead of you, is the vehicle ahead of you a Vespa or an 18-wheeler, is there a vehicle in your blind spot and if so what kind of vehicle is it, etc. Based on this, the driver and indeed the autonomous car can take response actions such as switching lanes, accelerating, slowing down, taking an exit, etc.

While autonomous vehicles are clearly a world apart from security operation, the analogy is again helpful in understanding three basic requirements for an autonomous system. We will discuss these next.

Reduce Cognitive Load

Cognitive load[7] is the amount of working memory being used. Clearly most humans don't have an infinite or extensible supply of working memory, so for a person to be effective at any task, the cognitive load of performing the task must be less (and ideally a fair amount less, since the person has a lot of other things going on) than the working memory. This is true for drivers of an automobile or security operations workers sitting in front of a console.

When presented with too much or ambiguous data, human analysts are far more likely to make decisions based on intuition rather than sound reasoning. These intuitions are called judgmental heuristics[8], and are our way of reducing cognitive load. But more often than not these lead to suboptimal decisions.

The question then becomes one of how can you safely and optimally reduce cognitive load. In the context of the process laid out above, this comes down to eliminating some of the manual and tedious tasks in the process.

For instance, EntityIQ™is the world's only AI-based security knowledge graph that identifies, profiles and tracks all the devices, users and applications with just a network connection. This eliminates the need for analysts to manually consult with DHCP, DNS, AD, etc. to determine what device or user an IP address represents—something that can easily take 30-45 minutes adding to the cognitive load. Instead, the Arista NDR platform performs this set of tasks autonomously for you.

The automation also has some valuable side benefits. First, it delivers an instant skills upgrade. In other words, with the Arista NDR platform every member of the security team is now proficient at all the ways to track down a device or user, even if they aren't fully aware of what is happening behind the scenes.

The human can now focus on making a decision: does the device in question represent a critical asset to the organization? Is it in a sensitive part of the network? Is the user leaving the organization based on a list from human resources?

---

[7] https://en.wikipedia.org/wiki/Cognitive_load

[8] https://www.verywellmind.com/what-is-a-heuristic-2795235

The automation also does not leave the experts on the team behind. In fact, somewhat surprisingly, research[9] shows that while experts in a field are more accurate than a novice, both are liable to judgmental biases that result in errors and less than optimal decisions. In the context of security operations, autonomous security benefits your most senior analyst by decreasing their errors too. But it also helps ensure fewer alerts are escalated because junior analysts just don't know what to do with them.

Finally, and perhaps most importantly, all the automation frees up time for the expert analyst to hunt for more sophisticated threats, perform root cause analysis, improve processes and work on other more strategic initiatives.

Ultimately, the autonomous skills can run in parallel feeding off and into each other. This gets to answers quicker and ultimately lowers risk by reducing the time to remediation, without adding to the headcount needs for the organization.

Eliminate Stressors

Humans don't do well with monotonous tasks like responding to alerts. Consistently and accurately performing analysis is not our strong suit, especially in security where each new alert can seem like groundhog day. In fact, that repeating pattern is in itself stress inducing[10] which in turn negatively affects[11] the ability for a security analyst to perform the task in the first place. You can clearly see how this vicious cycle plays out.

Software can help here immensely. Rather than simply throwing more data at the human with the assumption that is simply what is needed to make a decision, what if we pre-computed the answers to the questions we laid out in the section above. The analyst doesn't even need to remember the questions, let alone ask them. Instead, the information they need is right there at their fingertips.

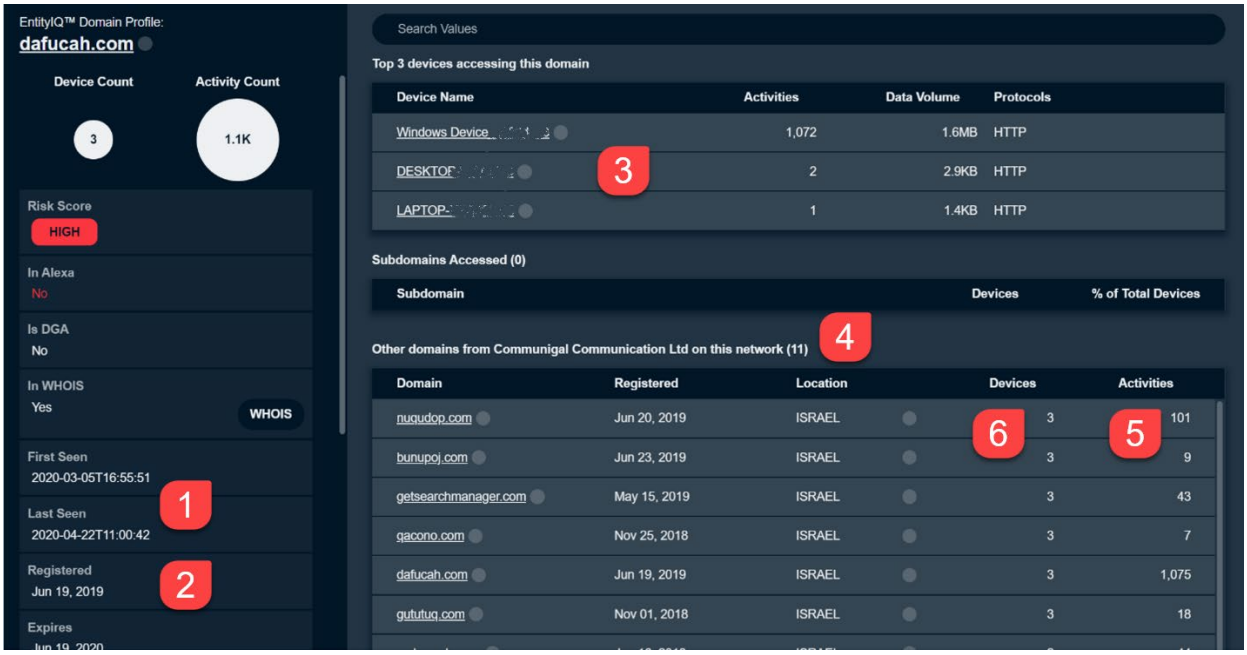

Figure 4: Arista NDR Decision Support System

[9] http://www.pitt.edu/~druzdzel/psfiles/dss.pdf

[10] https://securityintelligence.com/articles/9-reasons-why-cybersecurity-stress-is-an-industry-epidemic/

[11] http://news.mit.edu/2017/stress-can-lead-risky-decisions-1116

Expanding this concept further you arrive at what's referred to as a decision support system. AVA uses a combination of hybrid-cloud components to deliver such a knowledge-based system. As researchers[12] from the University of Pittsburgh so eloquently put it:

> *"They [decision support systems] are especially valuable in situations in which the amount of available information is prohibitive for the intuition of an unaided human decision maker and in which precision and optimality are of importance. Decision support systems can aid human cognitive deficiencies by integrating various sources of information, providing intelligent access to relevant knowledge, and aiding the process of structuring decisions."*

Because it has access to an elastic set of resources in the cloud, AVA systematically analyzes the entire decision tree like the type laid out in the phishing example above. It makes threat hunting decisions at each choice node[13] and then considers the different possible next steps or actions at chance nodes. A simplified snapshot of this process is illustrated in the diagram below where a choice node is represented as a square and a chance node is represented as a circle.

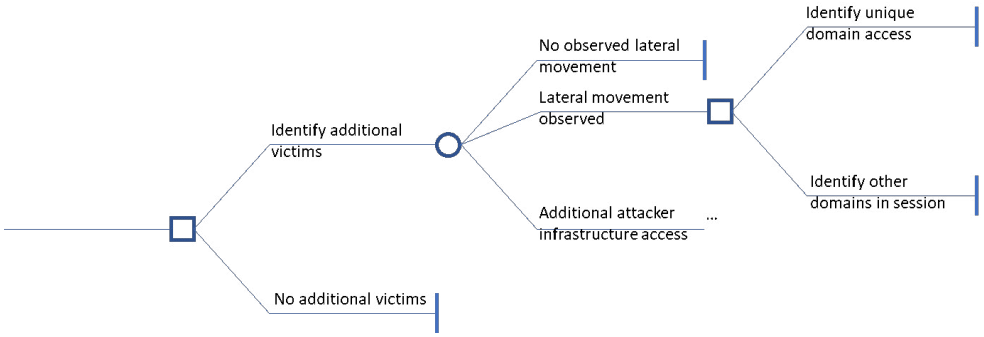

Figure 5: Graphical representation of Arista NDR Decision Tree

As the graphical view of the decision tree illustrates, it is not that a human analyst couldn't manage the same process. The issue is can you really rely on this happening consistently and in an error-free manner for every security event, especially while operating under stress?

AVA also recognizes when exploring a particular path consistently returns no value, or when explicit user feedback shows that the information surfaced is not relevant to the threat being detected and investigated. These kinds of feedback loops are processed through federated learning and help improve AVA's intelligence and refine the decision tree on a continuous basis.

AVA's analysis also has another stress-relieving advantage: it can ask very precise questions of the data that exists within the system and surface just that information for the analyst to consume.

While a human can potentially collect all that information, interrogating with the same precision is hard to replicate since it adds to the tedium. Instead, faced with this situation, most experts would ask broader questions and then eyeball the results for the specific data they are expecting to find. Not only is this approach likely to be error-prone, it also adds stress both to the system and to the analyst looking for the needle in the haystack.

[12] http://www.pitt.edu/~druzdzel/psfiles/dss.pdf

[13] https://people.ok.ubc.ca/bowenhui/iui/introduction-to-decision-trees.pdf

Focus on User Experience

One human quality that is especially useful in security professionals is healthy skepticism—the desire to understand why something occurred. While this skill is useful in breaking down and analyzing security attacks, it can hurt the process when those same analysts are saddled with software systems that come across as black boxes.

Explainability is therefore a key requirement for a security decision support system. The system must provide a user experience that documents the decision paths, or the choice nodes explored, and which ones resulted in new findings (or chance nodes) leading to further exploration paths.

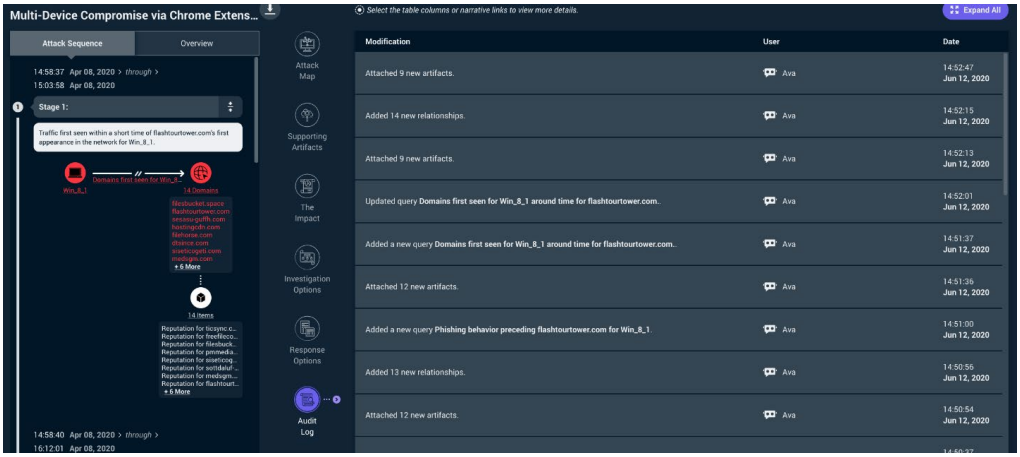

Figure 6: Enabling deduction through documented decision paths

In fact, our research found that to inspire confidence in the autonomous system it is just as important to list out productive paths as it is to identify paths that were explored only to find they were dead ends or terminal nodes in DSS parlance.

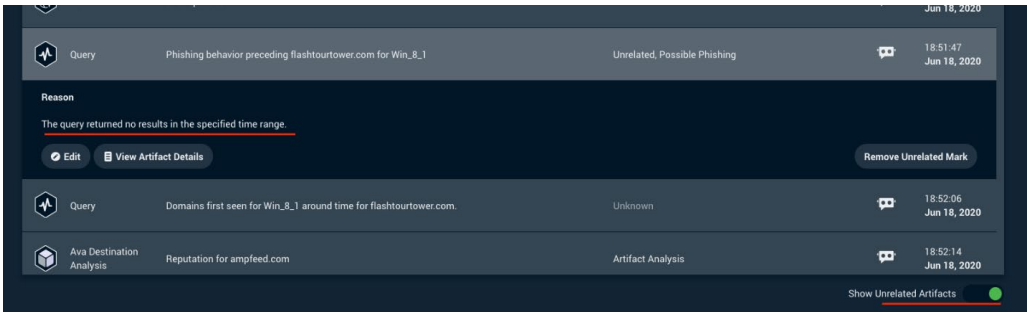

Figure 7: Exploring the Various Decision Paths of the Autonomous System

The user experience is critical in ensuring data is not just thrown over the fence at the analyst. Instead, the goal must always be to provide the right level of information for the task at hand with the option to dive deeper if necessary. This requires not just an understanding of where we are in the process but also who the user is and what their relative skill level is.

Finally, an often-ignored aspect of the user experience in automated systems is ensuring the human being is in control. Much like the human can override where the autonomous vehicle is headed or take a detour if they choose, autonomous security must let the human influence the process at any point. For instance, AVA allows the analyst to embellish the unfolding situation with information they may deem relevant. As an analyst you might determine or receive intel that another domain or artifact is relevant to the situation. AVA not only ingests that data but autonomously explores the analyst's hypothesis by applying all the relevant skills to embellish and add context to this human-sourced data. The analyst can then decide if this new artifact is still relevant or not.

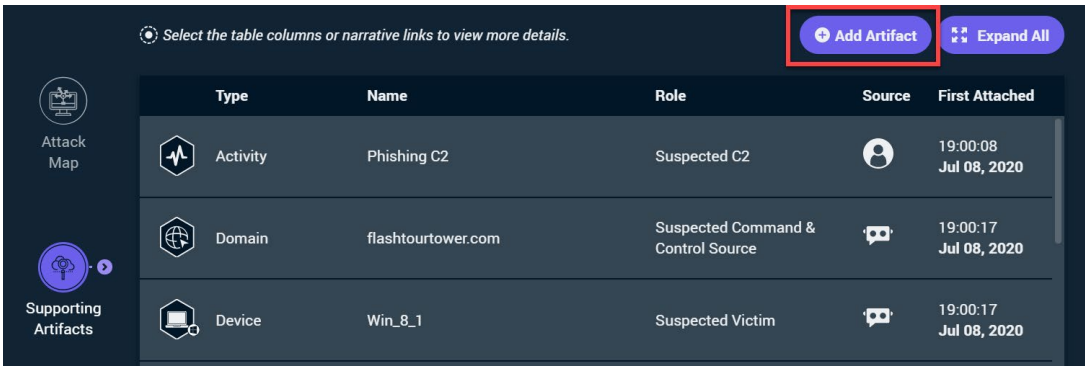

Figure 8: Exploring Various Supporting Artifactions with AVA

Summary

Human analysts do and will likely, for years to come, continue to play a significant role in the security operations process. With that said, the available human skills can be elevated to a higher level by eliminating both the tribal knowledge as well as the rigor needed to surface the information they need to make optimal security decisions. That is the role that a decision support system plays, and it puts the organization firmly on the path towards autonomous security. The combination of Arista's EntityIQ and AVA technologies, built on a foundation of cloud-scale federated machine learning, automate the decision tree analysis that security professionals struggle to perform manually today.

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office
1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062