Automating Enterprise Cybersecurity With Clusters

No industry is more desperate for automation than enterprise cybersecurity. Humans can successfully drive vehicles, and autonomous mobility is merely an advancement. But no sheer human effort can possibly cope with the practical demands of enterprise cybersecurity without some serious machine assistance.

The reasons are obvious once you grok the nature of the task. At a simple and entirely inadequate level, an enterprise cybersecurity team runs down alerts generated by its network proxy servers, firewalls and gateways. These are typically attempts to access IP addresses that have been specifically blocked by policy. But the number of these alerts can rapidly overwhelm the human security staff, and it's easy to disregard transactions that have already been blocked.

The ideal function of the cybersecurity team (often called the Security Operations Center, SOC) is to sift through the entire astronomical volume of enterprise network traffic looking for indications of compromise in near real time. If this were in any way practicable, it would involve human experts identifying and collating all of the perhaps very numerous network transactions that constitute the substance of a given attack or incident. For example, as part of a single attack plan, a number of enterprise hosts may have downloaded and installed malware, such that they periodically contact remote servers that may instruct each host to upload (exfiltrate) files to specified drop points. This seemingly complex scenario is routine and even unambitious by today's stealthy standards.

Imagine working by hand through perhaps terabytes of system logs generated all over the network, putting together the pieces in what amounts to the greatest jigsaw puzzle imaginable, or rather of searching for a set of related needles dispersed in a haystack. And there's an enormous difference between finding only part of the problem and finding all of it. In our example, to remediate and close the case means identifying every enterprise host and user compromised in the incident, and every upload of exfiltrated data. Leaving any living roots in the ground is untenable.

The only solution is to automate the process so that, using a machine equivalent of the applicable human investigation and sorting, all of the network traffic metadata are combed for groups of transactions that are logically interconnected and can be understood as a single case unit or incident, i.e., all the transactions made pursuant to a specific plan of a specific attacker. The tool for creating these functional and comprehensive groups that constitute the full scope of an attack is a mathematical discipline called ***cluster analysis***.

To get the roughest possible sense, imagine dots scattered on a flat surface. Cluster analysis can be used to automatically discriminate the dots into distinct clusters to the extent they coalesce around functional centers of attraction to create locally dense regions, like galaxies in space. Different algorithms and assumptions will obviously impact the results, so there's a certain amount of judgment, and even art, required in practical application. But the principle is solid and very well-established in machine learning science.

In our primitive example of locations on a simple Euclidean plane, the meaning of "distance" between points is self-evident. It requires some profound sophistication and deep thinking to apply clustering algorithms to the more abstract domain of network transaction events, by assigning them values that can be treated as functionally "nearer" or "farther" from each other. But assuming this problem to be tractable, we should now be staring down our Holy Grail. Instead of mountains of raw transaction data, the security professional is handed clusters of functionally integrated transactions that, ideally, represent the entire scope of an incident in a time-ordered narrative that identifies all of the implicated enterprise users and hosts. Even better, we should hope that these event clusters would be presented pre-characterized for comprehension, on the basis of the inferred malicious purpose and tactics. This could lead to automated or semi-automated remediation. Our sample incident could be presented to the professional in a UI in terms that correctly indicate the nature of the attack, with an option to immediately remediate based on this understanding of the underlying behavior. The sheer pressure to stop compromises as rapidly as possible, and to tick quickly through a list of open cases, makes default remediations almost imperative.

So we know what we want, or rather what the enterprise desperately needs, especially with an acute shortage of qualified human analysts. Can we get what we want? Well, a startup based in Tel Aviv seems well ahead of the curve in bringing this level of AI horsepower to the industry. I had some trouble remembering their name -- SecBI -- until I came up with the quip that "in cybersec, AI is BI." But the company would actually prefer that you understand AI to stand for their trademark "Autonomous Investigation."

Deployment is straightforward because you're simply streaming your network log data into the cloud-based service. (Using only transaction metadata saves the expense, hassle and risk of capturing data packets or decrypting them.) In simple terms, the SecBI tech analyses the log data to identify clusters representing the entire scope of an attack in its current state of maturity. These clusters are then themselves analyzed to infer behaviors (the malicious purpose) that can be labelled for clear identification and for which default remediation can be suggested. Outside threat intelligence is incorporated in this step to better confirm malicious behaviors.

Bottom line, the analyst consulting the SecBI dashboard is presented with the theoretical ideal -- a list of clusters worthy of attention, representing the full-scope of each incident (all component transactions), and ready for immediate remediation or further analysis. Does it work? The short answer is simply to give it a quick try by feeding it some of your own log data and see what your security team thinks -- not only of the clustered output and default remediations, but also of the powerful query engine for that more pro-active species of investigation called "threat hunting." And if you forget the name, remember "AI is BI."

Robert Polevoi writes about technology.