How to detect hackers who know we're on to them

July 18 2016

Cyber attackers' methods evolve rapidly, and software that worked to detect network attacks yesterday might be ineffective tomorrow.

The best detectors don't focus solely on keeping intruders out: they also help identify intruders that have already broken in, as through a malicious link or email attachment.

"Attackers go from one computer to another and mine for information, looking for system credentials and elevating their privileges," says Justin Grana, an SFI Postdoctoral Fellow.

Security tools designed to find these attackers typically scan a network and search for an anomaly – activity that differs significantly from

"normal" behavior. These statistical approaches assume that an attacker will stick out as they move through the network.

That strategy is problematic, says Grana, because it's difficult to know what behaviors are normal. What appears to be an intruder prowling the system might be the benign activities of new employees getting their bearings in a network. "There are a ton of false alarms," says Grana.

In a new paper in the Journal of Network and Computer Applications, Grana and his collaborators – including SFI professors David Wolpert and Tanmoy Bhattacharya – take a different approach. Calling on the tools of game theory, they suggest that a better way to stop an attacker might be to think like one.

Rather than assuming it knows how an attacker behaves, the proposed detector assumes attackers will follow near-optimal strategies given their knowledge that defenders are looking for them. This allows the detector to compare the probability that certain activities were generated by normal network behavior to the probability that it originated with an attacker. This ratio – not just the probability that the activities reflect normal network behavior – is used to determine whether or not to sound the alarm.

This better solves for what a smart attacker would do, says Grana.

"We want to use that information to refine our detector without assuming we know how the attacker will achieve their goals, only what those goals are," adds Wolpert.

The researchers' model has outperformed simple anomaly detectors under many network scenarios. To ensure their results scale up to real-world conditions, the team tested the model on network data from Los Alamos National Laboratory.

Grana says the paper represents a first step toward integrating game theory ideas into smarter detectors.

More information: Justin Grana et al. A likelihood ratio anomaly detector for identifying within-perimeter computer network attacks, *Journal of Network and Computer Applications* (2016). DOI: 10.1016/j.jnca.2016.03.008

Provided by Santa Fe Institute