Which is the Right IDS for You?

March 15, 2021

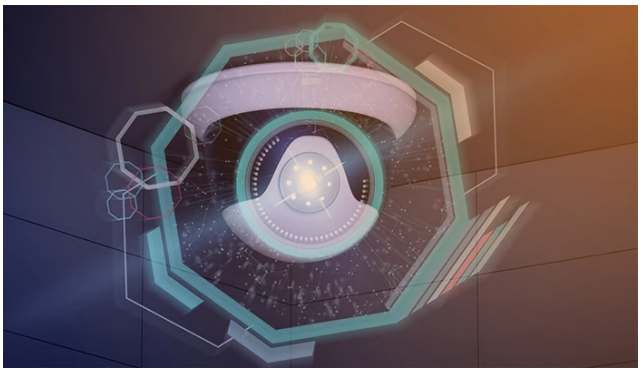

Where to place an intrusion detection system (IDS) is an important first question, but it's far from the last decision you'll make as you work to protect your enterprise from zero-day attacks like SolarWinds.

Once you've selected a location — on the network, on your devices, or, at a more granular level, on your applications, — you must next choose which kind of detection your IDS will use. What kinds of information should your IDS monitor?

One popular type of IDS is **rule-based**, which means the system identifies intruders based on rules that someone has written. Blacklisting is an example: Someone writes a rule flagging specific IP addresses as malicious or fake. When the IDS sees information packets coming from any of those addresses, it sends an alert.

Signature-based detection, a type of rule-based detection, involves monitoring for specific identifiers such as signatures, hashes, or code patterns.

With rule- and signature-based detection, the IDS is looking for known threats. New or otherwise unknown threats can go unseen. Since attackers tend to switch tactics often, the chances are high that your rule- or signature-based IDS will miss something.

A promising, but premature, solution

Anomaly-based detection, which takes note of the network, endpoint, or application's baseline and sends alerts for unusual activity, is much touted as the answer to the problems that rule- and signature-based IDS encounter.

But this type of detection is a long way from being truly useful. In fact, it requires immense amounts of data to work well, and it needs to be able to improve itself continually using machine learning.

In spite of claims to the contrary, no anomaly-based IDS yet exists that doesn't flag a huge number of false positives or overlook certain anomalous behavior altogether ("false negatives").

False positives can cause alert fatigue, leaving room for real threats to get overlooked.

False negatives, on the other hand, give a free pass to intruders.

What anomaly-based IDS does do well, especially at the application

level, is spot aberrant use patterns. Is traffic leaving the application for some strange location? Is the user logging on at a time when they would normally be asleep?

At the application level, changes in use are easier to spot because the scale of application use and the amount of data it generates are smaller than with the use of an entire device or network.

Highly anticipated, application-based, anomaly-based IDS holds particular promise for cybersecurity and the thwarting of zero-day attacks.

In the meantime, just choosing to place your IDS on the host, at the application level, and to use as many detection techniques as possible — perhaps a hybrid of rule-, signature-, and anomaly-based — can greatly enhance your visibility.

Then, attackers will be less likely to lurk in your networks, undetected for weeks or even months, endangering your organization and your clients and customers.

To spot and stop stealthy, persistent attacks, however, you'll need to tailor your IDS to your entity's specific situation and needs. The next step in your IDS program: designing a solution and making it work for you.

BACK