

Correlation Threat Technologies Sense Security Problems

By Elizabeth M. Ferrarini

Patch management is a widely used but rarely discussed method of shoring up servers with operating system fixes as soon as vulnerabilities are discovered. Recent worms, however, show that patching is a failing strategy. Some of the patches don't work; some are unacceptable. One solution is to integrate a company's IT network sensors to get an enterprise view of attacks and manage them appropriately, focusing patches on mission-critical servers and proactively dealing with threats as they emerge.

Having security devices in place to manage threats is a first step toward establishing a secure IT infrastructure, but IT also needs to effectively manage its numerous security devices to make sense of the mountains of data they produce. Firewall, antivirus, VPN and intrusion-detection devices compile large and detailed event log files or equivalent data streams of all activities across the network, regardless of the level of threat.

Security technologies such as firewalls and antivirus are now among the most widely deployed systems; however, antivirus solutions are potentially vulnerable when their signatures are not up to date (a problem they share with intrusion-detection systems), and the value of firewalls as a defensive tool starts to diminish as more and more attacks become Web- or mail-based, exploiting the ports that are traditionally wide open on most corporate perimeters.

Like any application—and firewalls, intrusion detection (IDS), antivirus and their brethren are ultimately just software applications, even if running on a hardened appliance—these security systems must be kept up to date, as they otherwise become decreasingly effective. Just as importantly, however, data streams—the logs that these devices produce—need attention. These logs represent a virtual paper trail that can yield a wealth of information about what attacks are being mounted.

Herein lies the security manager's dilemma. If all the checks on the IDS are enabled and the firewalls' log files are reviewed, there is simply too much data. These systems generate too many false alarms, known as false positives. If the security teams do not look at the data, or tune it so that only the most egregious threats are alerted, then some real threats will not be reported. Unfortunately, the only way a false negative can be discovered is once the network has been compromised by it.

Which Traffic to Watch

Compounding the false-positive problems from IDS is that they typically identify all suspect traffic, even when those exploited do not apply to the infrastructure. There is no value to being alerted to an attempt to exploit an intelligent server switch (ISS) vulnerability if, for example, only Apache is run.

Turning those signatures off is tempting, which is fine if there are no Internet information servers (IIS) in the organization. The default installation of Windows includes ISS, which means there are probably unmanaged, forgotten

(or end-user) machines that are vulnerable. These machines can be used as a base to spread the relevant virus, or as a base to steal corporate data.

Vulnerability scans are part of the solution to this problem, but they, too, need to be kept up to date. Moreover, like IDS, they may produce large numbers of false positives, so they need to be treated carefully.

Clearly, what is needed is the ability to accept as much information from sensors in real time to reduce false negatives, and then use threat-management technology to filter out the false positives, integrate data together to track professional-grade "blended" threats, and map the output to vulnerability databases. Ideally, these technologies would also deliver forensic and management reporting capabilities to enable post-mortem analyses of individual events.

Threat managers exist today, delivering security management by sorting through sensor data streams, some in real time, to filter out irrelevant information. With many of the false positives identified and removed, these products then normalize the data into standard representations so they can be correlated.

Correlation is the process that enables security managers to focus on only the most important and relevant threats. Correlation should serve three major purposes: pulling out remaining false positives, escalating false negatives and enriching threat data.

One way to remove false positives is to eliminate them using smart data-collection technology, reducing the volume of likely threat data to the correlation engine. The correlation engine itself can then reduce false positives further in a number of ways.

First, by failing to find any other threatening behavior for that particular attacker or target (by not correlating the event with anything else), the correlation engine indicates that, in all probability, the attacker is not active and has turned his attentions elsewhere. That attack is therefore no longer important, and if anything needs to happen it can be deferred to later (e.g., by patching the target host appropriately).

Flagging Real Threats

Additionally, the correlation engine can flag threats that are inapplicable by cross-referencing event data with the vulnerability assessment database. If the attack is an attempt to use an IIS exploit on a Solaris Apache server, that threat can be ignored from a risk perspective. So, the false positive is never surfaced to the security response team. A correlation engine is, in fact, the only place where tying IDS data streams to vulnerability data about the network makes sense, because it is the only place where the data from all IDS implementations has been normalized into a single coherent dataset.


Although the correlation server does not escalate false positives, they still need to be tracked. This is because they indicate malicious intent, if

nothing else, and the staff needs to be aware of intruders attempting to compromise systems or steal intellectual property.

Correlation technology tracks all attacks and can link events from multiple sources to identify persistent and related attempts to breach security. The correlation engine should escalate threats that are somehow linked according to its rules, and escalate them quickly to the security management team.

The correlation technology needs to find attacks regardless of the vendor or technology deployed, so a cross-vendor, cross-platform solution is vital. Correlation systems can pull the data in from around the network and link it together to derive information about active threats that staff simply could not find out about otherwise.

The best correlation engines will also go one step further than simply filtering threats in or out of the security-management console. Some correlation engines will add links to the vendor's Web site to threat data, allowing the operator to quickly understand the risk and remedy of each individual event that the correlation center delivers to the console. This enrichment simplifies and automates the process of threat containment.

Evolving from log analysis automation solutions, these threat-management technologies relieve the burden of manual or post-mortem log file analysis, and reduce the risks of false negatives, while reducing the burden of responding to false positives. 

Elizabeth M. Ferrarini is an IT consultant from Boston, Massachusetts.

