BRICATA®

WHITE PAPER

IDS/IPS: The Most Useful Threat Detection Tool You Have

Why it is so powerful if used properly, and five key reasons why it frequently fails to deliver

By Justin Bajko

January 2017

The Twitter Paradox

The other day, I opened Twitter for the first time in a while (more on that later). I tapped on the icon on my phone, the little Twitter logo popped up, and there sat my timeline. And I said to myself, "You know, I clearly follow way too many people, because this thing just looks like a stream of noise to me. There is probably information here that I'd think is useful or entertaining in some way. Probably. Somewhere. Maybe." But I had no way to tell. There was just too much information — too many sources, too many categories, and most of it (because of the 140 character limit) lacked the context for me to do anything with, unless I clicked on a URL and read an accompanying article.

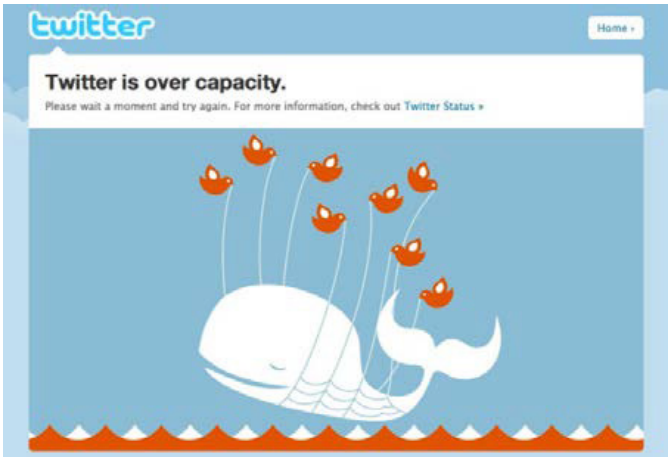

I decided I was going to whittle down the list of folks that I'm following to make it easier to digest my timeline. So I tap over to the list of accounts that I follow, and I'm given a list. I don't get any information about accounts that I've previously found most useful. It's impossible for me to easily manage my list of followed accounts. And that's when I remembered why it had been so long since I opened the Twitter app — information overload and difficult management. I don't dispute that it's probably good information, and there are undoubtedly things in there that I want to look at. But I just don't have the time or the patience to go through it all. It's simply too much. And trying to keep it curated is a hopeless task. So, technology that was designed to make something easy actually ends up making itself useless.

I've taken the time to tell you this story, because this is a microcosm of what your security analysts are likely going through on a daily basis when monitoring your intrusion prevention infrastructure. It is illustrative of why so many companies struggle with effectively deploying and monitoring IPS in their enterprise. Many people in the security industry would have you believe that "IDS/IPS is dead. It's a dinosaur technology." Well, I've spent a lot of time working with customers of all shapes and sizes and helping them find bad guys throughout their network, and one of the most consistently reliable ways that I've seen to find bad guys is a properly deployed, maintained, and monitored IPS infrastructure with a well-curated set of intelligence.

Network Monitoring is Essential for Security

IPS systems are invaluable as part of any security architecture. Just like you need door, window and motion sensors for physical security, your network sensors are the eyes and ears of your security system. They can also function as security cameras, capturing a chronological record of activity in the form of log files or even stored packets (PCAP), which give security analysts the ability to look back in time to see what has transpired.

One of the most consistently reliable ways that I've seen to find bad guys is a properly deployed, maintained, and monitored IPS infrastructure with a well-curated set of intelligence.

Whether deployed in IDS (passive) mode to inspect and alert, or IPS (active) mode to proactively block specific attacks, high quality sensors are very effective at identifying and alerting security personnel

about threats that require attention. As part of an integrated infrastructure, they deliver critical data needed for event correlation and advanced analytics. Without them, organizations are essentially blind to what is happening across their environment, and critical tools like SIEMs are substantially less effective at being able to piece together a chain of events.

There are other practical and non-obvious advantages for IPS deployments, as well. If you're suddenly interested in a new type of traffic — whether as a part of an active investigation or because you're concerned about the potential of a new in-the-wild attack — an IPS deployment is a fantastic way to get real visibility into that traffic, very quickly. It's already out there watching the network. Make it work for you to make your life a little easier.

Understanding IPS Challenges

The primary challenge most organizations face is that their legacy IPS systems have not been properly deployed or configured. So, let's spend some time talking about why organizations can fail when attempting to deploy and monitor an IPS infrastructure, and some of the ways those failures can be avoided. There are five key areas that frequently present the most challenges:

1. Lack of Planning and Visibility

Many IPS deployments are doomed from the very start, because of poor planning. If you want to get the most out of your IPS deployment, you need a solid understanding of the network infrastructure into which it will be deployed. In many cases, organizations take a network map that was written years ago by someone who is no longer with the company, and they plan their deployments that way. Well, networks are organic and they change fairly rapidly. Having a program in place that regularly assesses the state of the network and using that information to inform the deployment of the IDS is critical. And if you don't have a program like that in place, it's best to spend some time understanding the nature of the network before you start deploying IDS devices. Spend some time with your counterparts who run the network. Understand how the network is segmented, and which devices are responsible for those networks. Once you place your IPS device, confirm with your network engineers that the traffic the device is seeing is the type and amount of traffic that is to be expected for that network segment. Maybe even schedule quarterly review meetings with them to make sure the network hasn't fundamentally changed in a way that now limits your visibility.

Unfortunately, a regular conversation that organizations have with incident response teams after a breach is the one where they find out there are large swathes of the network that, at best, have no security visibility, and at worst, no one even knew existed. Spend the time to understand your network — it'll pay dividends over the years, and not just as it relates to your IPS deployment.

2. Set It and Forget It

Once the devices are deployed, IPS sensors tend to fall victim to the false assumption that they are "set it and forget it" devices. They aren't. A lot of the routine management functions can of course be automated, such as new signature deployment. But every organization and network is different — the signature set that works for one environment is likely not the ideal signature set for another environment. IPS signature sets that aren't well managed can easily become a signal: noise nightmare for your analysts. And when that happens, things get missed.

IPS signature sets that aren't well managed can easily become a signal to noise nightmare for your analysts. And when that happens, things get missed.

Just look at some of the recent case studies of high-profile breaches. In many cases, the warning signs were there — the security architecture found something and alerted on it. But because there were just so many alerts that analysts were responsible for, it became impossible to determine why this alert might have been more important than any of the others. If your analysts get to a point where they can't understand which alerts are most important and can't get through the important ones, that's a red flag (if you're in that boat, though, don't worry — a lot of organizations are), and it's time to start figuring out how to dig yourself out of that hole.

One way to start down this path is to have a process by which signatures that are routinely generating false positives or alerts that aren't interesting to the organization are flagged for further review. During that review process, the signature can be tuned or completely disabled. Reducing the signal to noise ratio is crucial for an effective IPS deployment.

3. Network Changes and Misconfigurations

And let's not forget an earlier point: networks are organic. They change. And if you set and forget your IPS, you may not notice that someone unconfigured that one span port. You may not notice that some network engineer made a routing change and now you're exceeding the bandwidth the device can handle, and you're dropping packets from inspection and potentially letting attackers sail right by.

However, there are some things that you can do proactively. Set up an alerting infrastructure that warns your team when the IPS device is reaching its capacity, whether from a bandwidth, memory, or processing perspective. The last thing you want is for an attacker's traffic to be what the IPS has to drop from inspection in order to keep its head above water. If your device is configured in passive mode, routinely check those devices to ensure that all of the monitoring ports are still receiving traffic, and that the level of traffic is what you would expect for that network. If you see dramatic differences, spend some time figuring out why. Like any piece of network or security equipment, it's important to have a regimen in place for maintaining the devices to make sure they're performing at their best. If you take care of your IPS, it'll take care of you. Unfortunately, however, that's easier said than done in many cases.

4. Device Management and Threat Intelligence

This brings us to the next point: IPS infrastructure management can be painful. The reason that so many of them are left to rot is because it's a cumbersome and difficult process, and eventually many organizations just give up because they think it's a losing battle. It becomes just another example of security shelfware. Not only is it difficult to maintain these devices, but intelligence sources are often not straightforward to integrate. This can lead to the monotonous task of, in some cases, copying and pasting lines from a spreadsheet into an IPS management console. And with the rate that security intelligence is produced today, that can quickly become a full-time job. Or two. Or three.

5. Context Around Security Events and Alerts

But let's assume for a minute that you've conquered all the challenges related to getting and keeping your IPS deployment in a good place: You know your network, you're confident that you've got the devices placed properly and seeing the right traffic. You have a signature set that is curated for your organization that makes your heart fill with pride when you think about it, and you're regularly integrating new intelligence feeds and updating existing feeds. Life is good, right? Maybe not.

If your analyst is looking at an alert, does she have the right level of context to actually act on that alert? In too many cases, the analyst is told that a signature with a cryptic name fired, and the only description that they're offered is too bare. As shown in Figure 1 it'll say something like, "This signature is designed to detect exploits in the wild."

Okay, well now what? What does the analyst do with that? She can maybe look at the associated packet data, but that's a snapshot in time. It tells the analyst

```
hojung@ubuntu-171:~/temp$ cat alert
[**] [1:28911:1] EXPLOIT-KIT Neutrino exploit kit initial outbound request - generic detection [**]
[Classification: A Network Trojan was Detected] [Priority: 1]
03/01-12:45:03.447872 192.168.204.174:49462 -> 23.227.189.17:8000
TCP TTL:128 TOS:0x0 ID:13722 IpLen:20 DgmLen:544
***A**** Seq: 0xC41BFE7  Ack: 0x4092166A  Win: 0xFAF0  TcpLen: 20
[Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=2013-2465][Xref => http://cve.mitre.org/cgi-bin/cvename.c
gi?name=2013-2423][Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=2013-1493][Xref => http://cve.mitre.org
/cgi-bin/cvename.cgi?name=2013-0431]

[**] [1:28474:2] EXPLOIT-KIT Neutrino exploit kit outbound plugin detection response - generic detection [**]
[Classification: A Network Trojan was Detected] [Priority: 1]
03/01-12:45:10.086312 192.168.204.174:49462 -> 23.227.189.17:8000
TCP TTL:128 TOS:0x0 ID:13967 IpLen:20 DgmLen:577
***A**** Seq: 0xC41C51B  Ack: 0x40953128  Win: 0xFAF0  TcpLen: 20
[Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=2013-2465][Xref => http://cve.mitre.org/cgi-bin/cvename.c
gi?name=2013-2423][Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=2013-1493][Xref => http://cve.mitre.org
/cgi-bin/cvename.cgi?name=2013-0431]
```

Figure 1. Generic IDS Alerts

what was in the packet that caused the signature to fire. But the analyst doesn't know if there was anything interesting that happened leading up to that alert firing, nor does she know if the exploit attempt was successful, and if so, was there follow-on activity that indicates an active breach?

Context is king when triaging potential security breaches and then responding to them. If you can't give an analyst a picture of what's going on — not just what happened — then she's left to pivot into other sources which takes time, at best. At worst, she's guessing. And guessing is not a great practice when it comes to finding bad guys in your environment. Legacy IPS systems have left organizations to fend for themselves when it comes to validating an alert. You know the alert fired, but you lack the context to figure out whether that alert translated into a real intrusion. Luckily, modern IPS systems are beginning to provide security teams with exactly that level of context, sometimes through open source technologies like Bro. Leveraging a technology like Bro provides the analyst the necessary context — visibility into what happened before and immediately after the signature fired, for example — that she needs to determine whether that alert warrants further investigation.

Light at the End of the Tunnel

There is no silver bullet to managing the challenges discussed above, but a well-defined program that combines people, process and technology can. And it isn't as difficult as it sounds — it just takes a little planning, a solid understanding of your environment, defining a process for how your organization is going to monitor and respond to incidents, and applying the right technology to enable everything. On the technology side, new advancements in IPS solutions can help in several ways.

To start, newer open source technologies such as Suricata and Bro give security teams the power and flexibility they need to identify attacks with greater fidelity and reduce the barrage of false positives plaguing most organizations. On the Suricata side, intelligent signatures and multi-threaded processing make line rate detection more efficient and cost effective. Protocol recognition allows rules to be created that are protocol specific but not restricted to specific ports. It can also run IP reputation and MD5 hash analysis in memory. The improved processing power allows organizations to run a lot more rules — meaning you don't have to sacrifice coverage to meet throughput requirements and avoid dropping packets.

Bro delivers network behavioral analysis with powerful scripting capabilities that can detect more complex indicators of compromise (IOCs). It's also incredibly valuable at delivering rich metadata that helps analysts with badly needed context around security incidents. If you use sandboxing technology for advanced threat defense, Bro is also very efficient at file extraction, and you can use Bro sensors to forward unknown files to a centralized sandbox — saving you considerable time, money and headaches deploying multiple sandboxes across the network.

Combined, these two engines make it possible for organizations to detect more threats, more accurately. They're also fantastic at recognizing post-foothold activity, such as command and control communications. This helps give analysts important context by identifying what happened after an alert fired that may escalate its priority.

The biggest challenge with these new open source technologies is that they're, well, open source. Which means they're typically unsupported and lightly documented, which can make building and deploying systems a sizable project. For organizations that need a lot of sensors, the prospect of building and managing that infrastructure is daunting, knowing that security personnel will spend a lot of time managing boxes instead of investigating and responding to incidents.

IPS signature sets that aren't well managed can easily become a signal to noise nightmare for your analysts. And when that happens, things get missed.

Fortunately, companies are beginning to emerge that provide commercial-grade solutions built on these innovative open source technologies. Bricata's line of ProAccel appliances (www.bricata.com) offer the industry's first hybrid IPS that includes multi-engine detection with Bro and Suricata integrated on a single device and managed through a central GUI console. It combines signature and network behavioral detection methods to improve threat identification and efficiency — allowing each engine to process what it's best at detecting. Centralized management with built-in signature, script and policy management features makes large-scale deployments a lot easier than building and maintaining it yourself. Bricata also includes numerous usability enhancements and technology integrations that make it easier to fit into your existing infrastructure. And since it's a commercial solution, it has a full complement of support and warranty options.

Network visibility is critical. Organizations are largely frustrated with IPS technologies due to their historical difficulties. But few, if any, are willing to abandon them and risk missing a critical attack. Your network monitoring sensors can be the foundation of a strong security infrastructure if deployed and managed properly. New technologies make it easier to accomplish those goals, and at the same time, deliver enhanced value to improve threat detection, event analysis and incident response capabilities. Even if you choose to go it with your existing technologies, spending a little time to align the right people, plan your processes, and optimize technology will pay huge dividends in creating sustainable, effective security for your organization.

About Bricata

Bricata is a network cybersecurity solution supplier helping organizations harness the power of complete network visibility to detect, hunt, and prevent threats with the only commercialized Open Source and partner developed malware conviction engine. A specialized component-based approach to today's advanced, persistent, and coordinated attacks leaves organizations with a stack of tools to manage, lack of visibility across the network, and inconsistent security policies. Bricata's platform for federating security technology and console provides organizations with process automation, streamlining operations with the most effective, affordable solution for situational awareness and proactive threat defense, reducing complexity, dwell time, and time to containment.