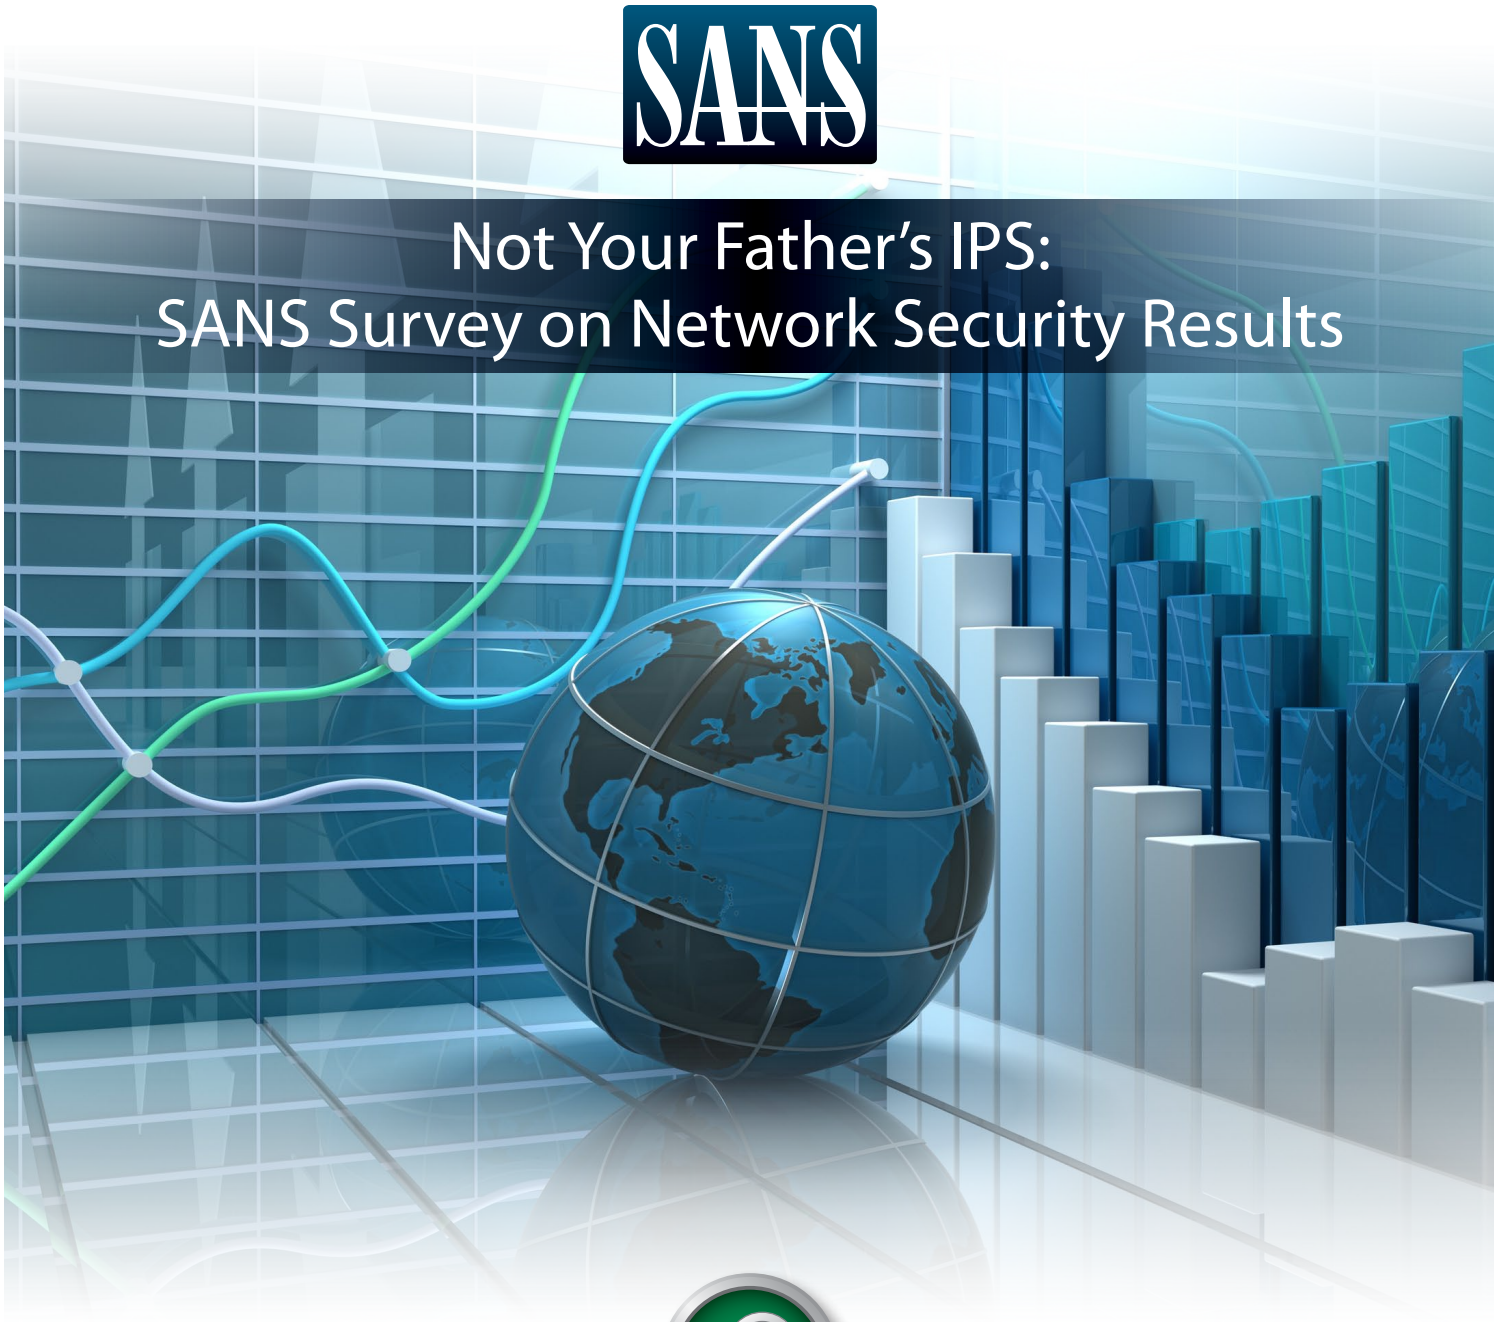

Not Your Father's IPS: SANS Survey on Network Security Results

A SANS Analyst Survey

Written by Rob VandenBrink

October 2013

Sponsored by
Hewlett-Packard

Introduction

In today's security landscape, IPSs are in the process of radical change. Established IPS companies are rapidly improving the look and feel of their IPS devices, making them much more management friendly. IPS inputs were once confined to direct packet capture, but now the packet capture function, while still critical, is in many cases being moved to a sensor platform. The IPS proper is now moving toward what many would call a "Next-Generation IPS" (NG-IPS), complete with a central IPS console and inputs from many sensors. Sensors might include multiple packet capture sensors, but they might also include endpoint protection systems, syslog, logs from Windows servers or a wide variety of logging services.

With all this data to play with, many IPS platforms are rapidly morphing into a more "SIEM-like" product. Some IPS companies have arrived here by purchasing security information and event management (SIEM) companies; others have gotten here by developing new products and extensions to their existing products.

Because of these changes in network security, SANS asked our community of security professionals to let us know about their network security practices today, their use of IPS, their technical and management capacities, and how they expect their IPS to be integrated into their overall security strategy in the years to come.

This SANS survey also asked what features these security professionals would find most useful and how they would use them for improved visibility and more reliable blocking. Finally, we asked questions about IT security budgets and support for the security effort from management.

This document, which analyzes the network security survey responses provided by 439 security professionals in a variety of industries, contains the answer to these and other questions.

About the Respondents

During the months of August and September 2013, SANS surveyed professionals whose organizations use network IPS systems as part of their security infrastructure to learn how IPS systems were deployed and used today and how respondents felt IPS products could be made more useful over the next few years. The survey did not target any specific industry or market segment. In fact, the 439 respondents come from a wide range of industries, with the financial (19%) and government (16%) sectors being the most commonly represented, as illustrated in Figure 1.

What is your company's primary industry?

Select the best answer.

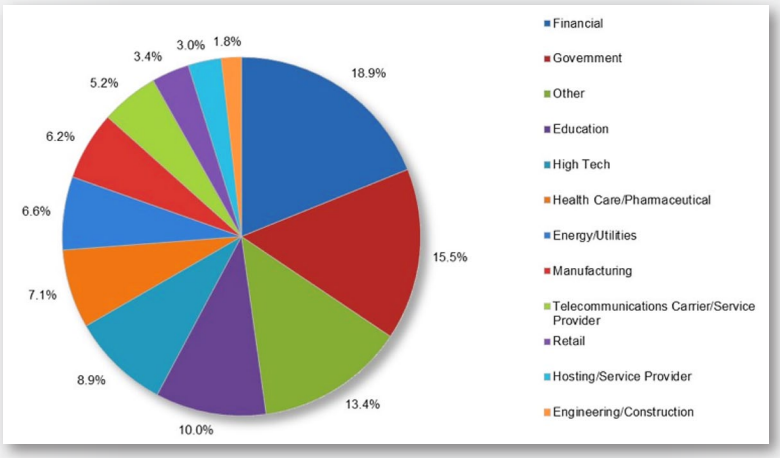

Figure 1. Industries Represented in This Survey

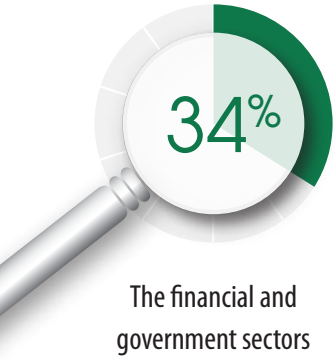

The financial and government sectors account for more than one-third of the responding organizations.

Responses also reveal representation from a wide variety of organization sizes. Most are from large or very large enterprises of 2,000 or more (domestic and international), which represented 47% of survey takers. An additional 11% come from Global 200 organizations, and another 11% are from the smallest organizations of fewer than 100 employees, as illustrated in Figure 2.

How large is your organization?

Select the best answer.

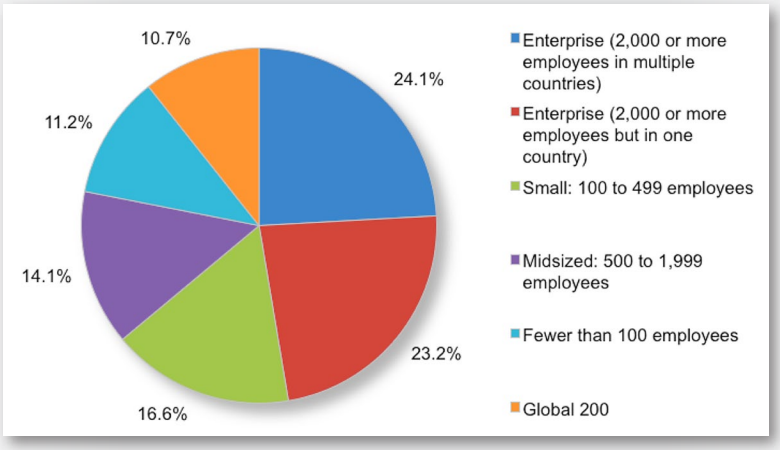

Figure 2. Sizes of Organizations

Roles Represented

The largest single category selected was security administrator/analysts, followed by IT security managers and then network or system engineers. In terms of job titles, the management side of the team (CSOs, CIOs, auditors and business unit managers) was well represented, as shown in Figure 3.

What is (are) your role(s) in the organization, whether as staff or consultant?

Select all that apply.

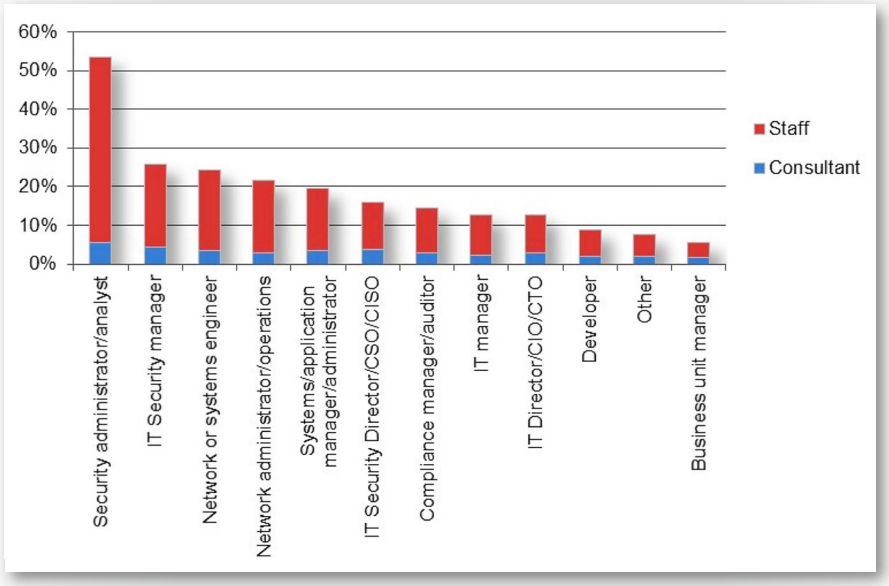

Figure 3. Respondent Roles

The vast majority of respondents indicated that they perform as a staff member in their various roles (e.g., employee) as opposed to as a consultant.

Those who get their hands dirty in managing security are the majority in this survey, when all responses are considered.

Responsibility for Network Security

Not surprisingly, an overwhelming number of respondents indicated that the ultimate responsibility for security in their organization rests with their IT security or IT operations groups. This response is certainly understandable when considering the importance of management of security events and associated log and event data, as well as how data within the organization should be secured.

An overwhelming number of respondents indicated that the ultimate responsibility for security in their organization rests with their IT security or IT operations groups.

While day-to-day security operations should rest with the technical teams closest to the problem, it's surprising that more respondents didn't indicate that the ultimate responsibility for security in their organization rests higher in the organization, with directors, the CSO/CISO/CEO or their board of directors. The results are presented in Figure 4.

What group in your organization is ultimately responsible for secure network operations?

Select the best answer.

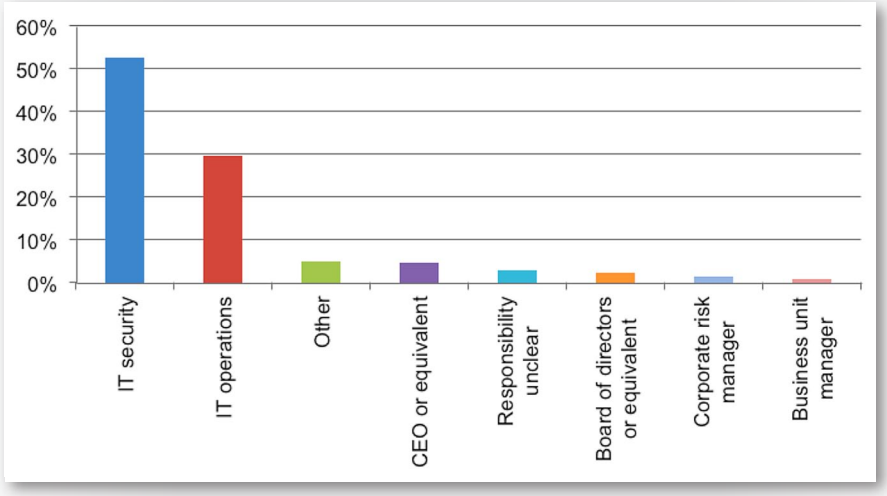

Figure 4. Responsibility for Network Security

This disparity indicates that we are still seeing a large divergence between legal precedent and real-life implementation within an organization. The courts hold the senior management team and board of directors responsible for the security. They are not responsible so much for the specifics of any one solution, but they must know that due diligence was applied to prevent breaches of regulated and confidential data, such as customer financial or personal information that falls under any privacy legislation. However, within many organizations, management still considers IT security as a single check-box item and defers responsibility for it to technical staff. This deferment is often exacerbated by the fact that improving security is too often seen as a "cost center" activity—something that is necessary but does not contribute to the revenue of the organization.

... improving security is too often seen as a "cost center" activity—something that is necessary but does not contribute to the revenue of the organization.

Investment in Security

What is interesting is the variations in security investments, in which the largest group (24%) responded that security made up 6–10% of their budgets. This is a higher portion than the 1–3% organizations commonly spend on information security (based on multiple reports). However, when combined, the 1–5% group accounted for nearly 40% of responses (see Figure 5).

What percent of your organizational IT budget is spent on information security management, compliance and response?

Select the best answer.

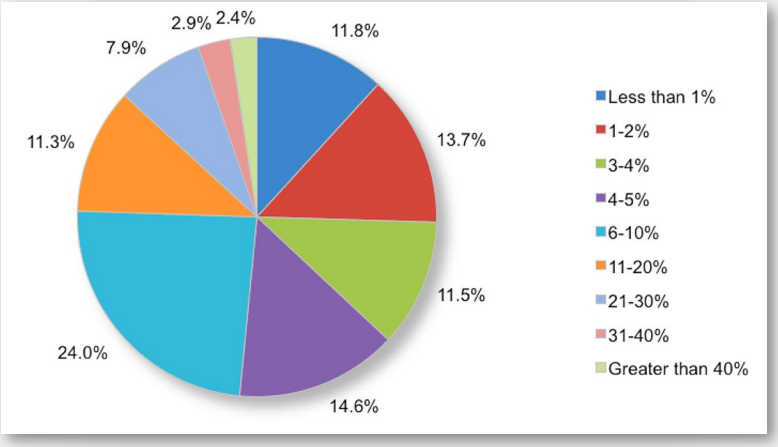

Figure 5. Percent of Budget by Organizational Workforce Size

The 2–4% range is more in keeping with outside statistics indicating the average security budgets tend to be around 2–3% of overall security budgets.

However, nearly 12% didn't have a reasonable budget for security (having responded with less than 1%).

Organizations are driven by both the business requirements for management-level reporting and compliance, and the technical requirements for better protections, technical reports and incident response tools in IPS products. We see these complementary needs reflected in our survey, where business requirements and technical needs seem fairly well balanced. In order, our respondents are concerned about the following:

- Sophisticated attacks that are difficult to detect
- Regulatory compliance
- Business-driven requirements for secure operations
- Increased attacks against the organization
- Increased successful exploits
- Mobility/BYOD initiatives

Percentage receiving full support (or full support with constraints) for security efforts

Levels of Support for Security

Business decision makers have a hard time with intrusion prevention, which tends to get lumped together with firewalls—often with the comment, "There, now we're secure." For managers not directly involved with security, the distinctions between overlapping services such as intrusion prevention, passive reconnaissance (with tools such as P0F), firewall services and advanced log management, which might be found in a SIEM product, are not clear. This lack of clarity sometimes makes it more difficult to garner management support for initiatives.

In our survey results, we see excellent support of network security efforts, with fewer than 9% of respondents indicating that they have no support or only enough to satisfy audit requirements. Almost 70% indicate receiving full support or full support with some constraints (normally in the form of budget constraints). Roughly 22% indicate support from management, but limited or no budget for tools and training (see Figure 6).

Does the management in your organization fully support your network security efforts, with appropriate budget, staffing and tools?

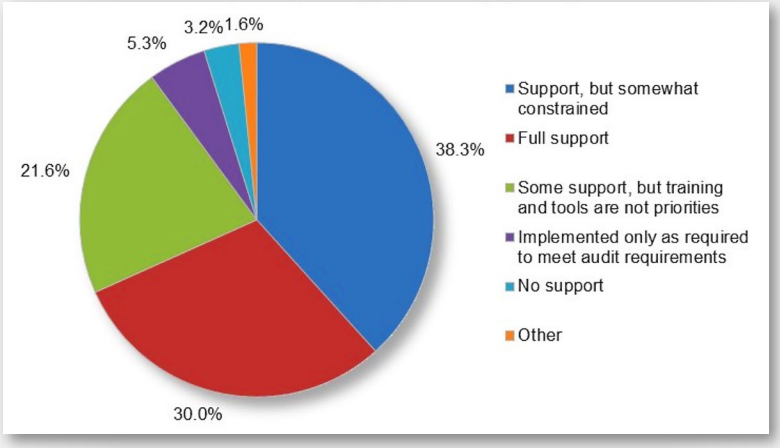

Figure 6. Management Support for Network Security Efforts

Support for security efforts was further reinforced when we asked what would happen if the IPS and associated infrastructure suddenly got better—if there were more inputs and more data—and the analytic tools also got better and easier to use. Overwhelmingly, the organizations that participated in our survey indicated this would mean they would grow their team's skills and move more of the operational aspects of the security team into other groups. In other words, they would institute much more training. (More on this result set later in our paper.)

Current Use of IPS

How are IPS systems used today? In almost all cases, IPS systems are deployed at the perimeter of a logical network, and its outputs are considered in isolation from other security data. So, during day-to-day operations, the analyst's view of attacks or other security incidents is shaped almost entirely by the IPS view of network traffic. It is only during an incident response for a security event that system logs, syslog data, central antivirus logs and perhaps consolidated data from a SIEM or log aggregator might all be used together to define timelines and determine what really happened.

In almost all cases, IPS systems are deployed at the perimeter of a logical network, and its outputs are considered in isolation from other security data.

While it's possible for an IPS to be used to monitor all traffic for a targeted VLAN, it's almost always deemed "good enough" to deploy the IPS at the perimeter of a targeted zone, analyzing traffic only on the way in and out of that zone, as illustrated in Figure 7.

How is your IPS infrastructure currently deployed?

Select all that apply.

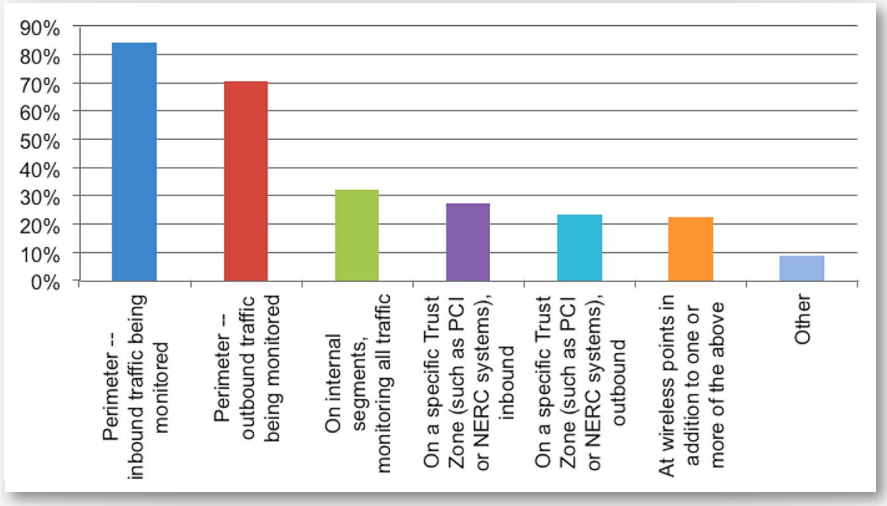

Figure 7. IPS Deployment Scenarios

The fact that they are monitoring outbound traffic almost as much as inbound traffic is important in prevention of data leakage and command and control channels between infected devices and their masters. There is less interest in using IPS on internal network segments, wireless access points and trust zones; however, 20–30% do deploy at these levels.

What They're Detecting

So how are these current strategies for IPS working? This is a tough question because we are, in effect, asking about information that people do not yet know—we're asking for numbers of successful attacks and how long a successful attack might remain undetected. These results indicate a failure of a product, implementation, training or processes. If anywhere, this is where you might expect more "spin" in responses. Our results seemed very honest, however.

While roughly one-third of respondents (33%) indicated no compromises in the past two years, almost 27% had 11 compromises or more that required manual intervention to resolve, as shown in Figure 8.

Percentage who had 11 or more compromises in two years

In the past two years, how many attacks has your IPS detected that were actual compromises that needed your attention?

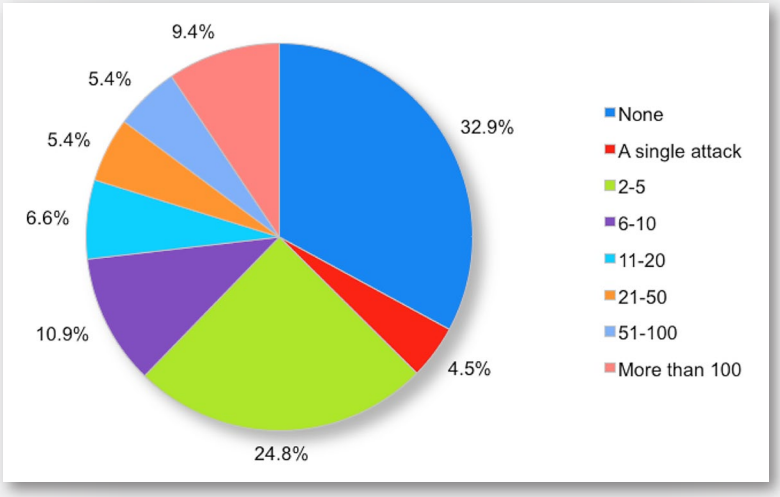

Figure 8. Attacks Detected by IPS

In such instances, the IPS system detected a compromise but didn't prevent it. While on the face of it this might seem like a negative statistic, this is actually a positive result. The prevailing wisdom these days is to assume that you will be compromised—and plan to handle the compromise correctly when it happens.

Detection Times and Thresholds

Digging deeper into this, we found more positive results when we asked how long respondents take to detect compromises. As expected, an overwhelming number of compromises are detected immediately or within the first day. However, the number of compromises that are measured in months or years of duration certainly gives us pause. Our respondents indicated almost 7% of incidents go undetected for six months or longer, as shown in Table 1.

Table 1. Compromise Detection Time

Time	Shortest time to detect	Longest time to detect
Immediate	45%	5%
1 day or less	28%	13%
3 days or less	4%	11%
A week or less	5%	14%
A month or less	0%	13%
3 months or less	1%	5%
5 months or less	0%	5%
12 months or less	0%	3%
24 months or less	0%	1%
More than 24 months	0%	2%
Unknown	16%	24%
Response count	100%	98%

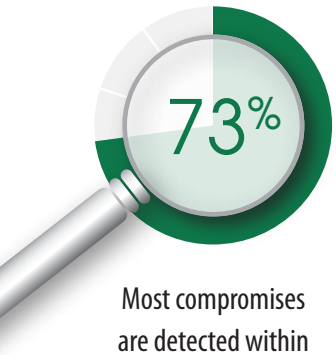

Most compromises are detected within the first day.

However, when asked about how their IPS is tuned, our respondents indicate that they still have a fairly high threshold before automated blocking occurs, with only 11% feeling confident enough to fully turn on their automatic blocking. Their responses are illustrated in Figure 9.

What percentage of attacks does your IPS system block automatically?

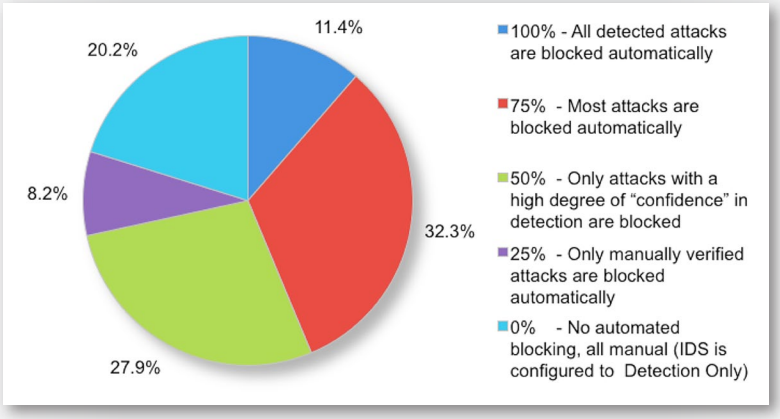

Figure 9. IPS Configuration

The good news is that most organizations are turning on some form of automatic blocking, with only 20% of respondents turning off automated blocking completely. These statistics indicate a fear, likely justified, that any false positives in IPS operation have a high probability of affecting critical operations within the organization. In the next section, we discuss what improvements organizations are hoping for in terms of network security that would enable them to more confidently use the features of their IPS systems.

What's Coming for IPS?

We've covered how security operations are viewed within the organization and how IPS systems are deployed in support of the overall security effort. In this section we discuss what the industry wants to see from tomorrow's NG-IPS systems and what we can expect in response to that.

What Users Want to See

By far, our respondents indicate that a next-generation IPS must include more application awareness (79%), followed by context awareness (67%), content awareness (57%) and full stack inspection (56%). This question allowed multiple responses, and this ranking indicates that, above all, respondents want smarter IPS devices that work with a variety of needs (see Figure 10).

... respondents want smarter IPS devices that work with a variety of needs.

What features would you like to see in a "next-generation" IPS, firewall or other security gear?

Select the best answer.

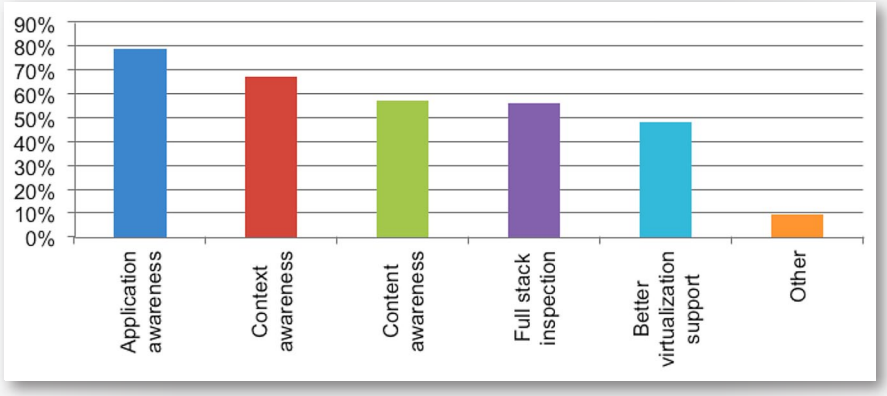

Figure 10. Preferred Features in Next-Gen Network Security

More Awareness

There are four types of awareness required by IPSs:

- **Application awareness.** When sourced internally, are different applications on the same protocol properly identified? For instance, is Firefox different from IE or Chrome or Safari? And is a Java app making the same request identified as a different application?
- **Context awareness.** Are different uses of the same protocol treated differently? For instance, is weekend work treated as more suspicious? Is guest traffic properly categorized? Is guest traffic from production networks properly identified and handled?
- **Content awareness.** Can different targets be treated as different applications? For instance, are Facebook, YouTube and LinkedIn all treated as separate applications?
- **Full stack inspection.** Is the entire packet or data stream properly parsed so that protocol masquerading, tunneling and other hiding mechanisms are all properly discovered?

Apps Are Not the Same

Not all applications are the same, nor should your intrusion management system treat them as such. In short, you should consider separate IPS policies around common applications, such as:

- **Distinctions between web destinations.** LinkedIn, for instance, is a different application from Facebook, which is different from YouTube, and each might have different rules applied.
- **Protocol distinctions.** Different applications that use the same protocol will be considered distinct events. A tablet or phone that uses a browser should be considered a different application and distinct from a PC browser.
- **Application platform distinctions.** Java and Flash should be addressable as different from generic web browsing.
- **Masquerading.** Tunneling http over DNS, or masquerading SSH inside port 80 or 443, should all be detected and controlled by inspecting packets and checking protocol characteristics.

In the comments, several participants identified an additional aspect that should be covered: identity awareness. A next-generation IPS should know who the person, service or machine is that is generating traffic. If they're part of the organization, there may be specific rules applied, depending on the group in which they are categorized. If they're not part of the organization, they may have an entirely different set of rules and permissions. For instance, perhaps a guest user session should be quarantined completely if the individual is attempting to access a production server.

Reporting that is useful for security personnel and understandable by management personnel is needed.

Other write-in suggestions included "better reporting"—a need also identified in the SANS Security Analytics Survey[1]—"visualization of traffic patterns" and "SSL intercept," with other write-in answers indicating, again, a need for smarter network security through better context and visibility.

Finally, participants also indicated that a more usable interface would be a win for a NG-IPS. As mentioned, we expect everyday IT shops to maintain an IPS infrastructure, but in most cases IPS configurations and events are not understandable to most IT personnel. English-language help text, event text in simpler language and a simpler configuration process would go a long way toward helping in-house personnel, who may not be security experts, deliver a better overall security solution.

[1] www.sans.org/reading-room/analysts-program/security-analystics-survey-2013

As you can see from Figure 11 (below), multiple organizations are already integrating their IPS with firewall technologies, and there is some movement in routers and switches, which are often also used as firewalls.

Where they are dropping off is in the realm of cloud, virtual and outsourced monitoring tools.

Deployments in Line with Wishes

When asked what components they'd like to see as part of their IPS system, respondents' answers aligned with what they're already attempting to integrate through their IPS systems, as shown in Figure 12.

Of these, which components have you currently connected with your IPS, what components are still on your two-year roadmap, and which are not yet on your roadmap?

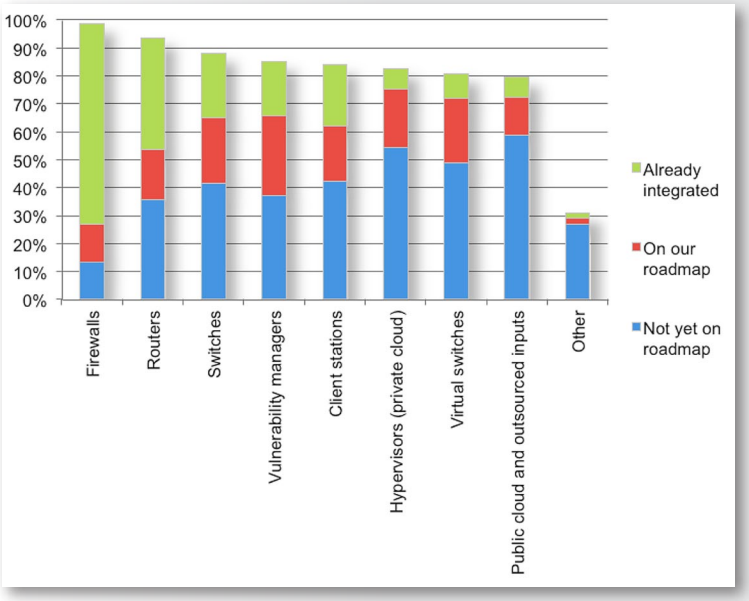

Figure 11. IPS Inputs Connected or Planned

If so, what infrastructure components would you like to see participating as agents in this IPS fabric for event data collection and/or enforcement? ***Select all that apply.***

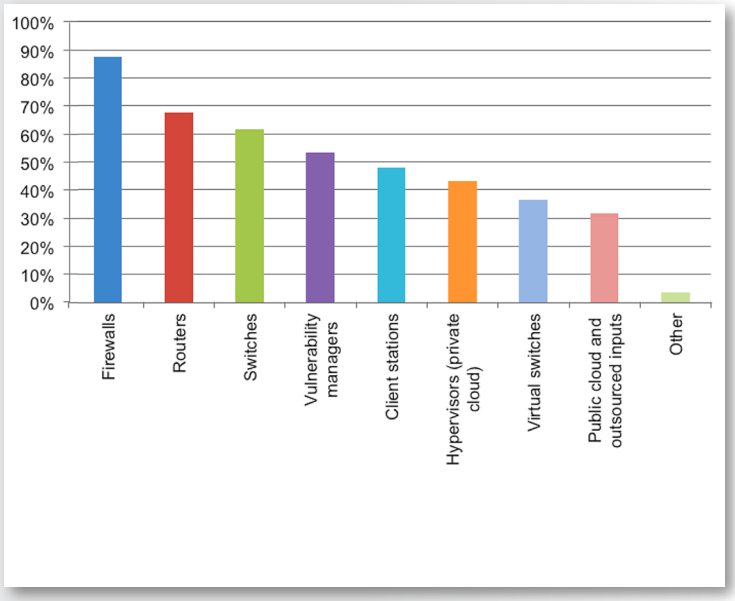

Figure 12. Desired IPS Inputs

Respondents are keen on IPS inputs residing on firewalls (87%), routers (68%) and switches (62%). However, a closer look shows almost half advocate for IPS input data being harvested from endpoint stations (47%), and 31% want public cloud inputs into corporate IPS systems.

This trend is reflected almost exactly in what our respondents already have deployed. Firewalls and, to a lesser extent, routers are prevalent platforms for IPS deploys and inputs.

Multiple Sources of Input

Of those organizations planning for workstation IPS inputs, 91% of respondents are looking for IPS inputs from traditional Windows, OSX or Linux clients. Almost half advocate for IPS inputs from corporate-owned phones (52%) and tablets (49%). Slightly fewer see a future where individually owned bring your own device (BYOD) tablets (39%) and phones (39%) will have apps to populate IPS data stores, as illustrated in Figure 13.

What client-side participation with an IPS for event data collection and/or IPS enforcement would be helpful in your environment?

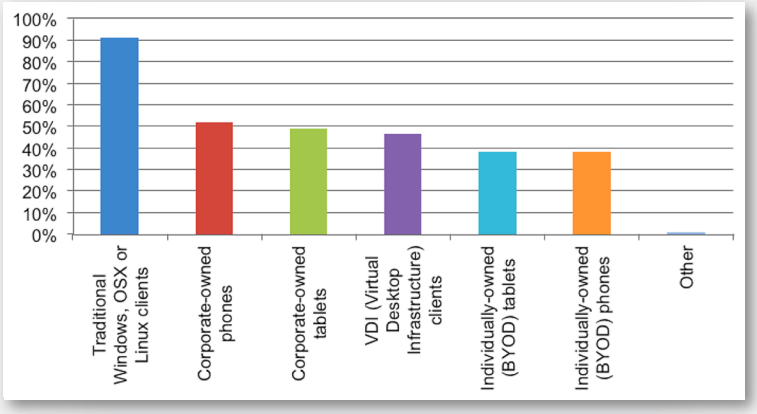

Figure 13. Future Participation in IPS Data Collection

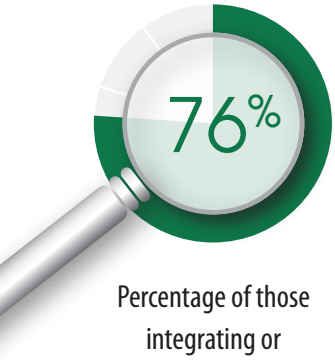

Percentage of those integrating or intending to integrate logging inputs

Roughly 76% of our respondents have either integrated logging inputs into their IPS systems or plan to within the next two years. This leaves a surprising 24% who do not have plans to use syslog, Windows logs or SIEMs as inputs for their IPS. This level of integration of varying log types is shown in Figure 14.

What client-side participation with an IPS for event data collection and/or IPS enforcement would be helpful in your environment?

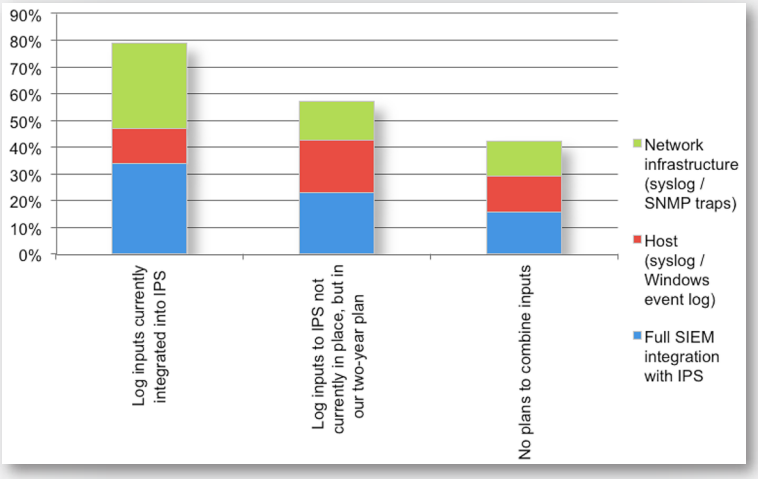

Figure 14. Log Input Integration with IPS

The inclusion of multiple sensory inputs working with the IPS for more accurate detection and blocking is introducing a new methodology in which IPS is no longer just a standalone product but part of a fabric of sensors, as discussed in the next section.

With a true "army of sensors" behind it, the IPS becomes part of what is coming to be known as a "fabric-oriented" approach to detection. In a traditional deployment, the IPS was a self-contained service, typically at some security perimeter, and often either adjacent to or part of a firewall. In this new fabric-based approach with the unified "sensors" mindset, the traditional IPS now becomes one more sensor among a larger ecosystem of monitoring and management tools. (IPS is perhaps one of the busier ones in that it's sniffing all traffic and re-assembling packet streams at a "choke" point.)

Wide Acceptance

This approach is not just a vendor trend. This is something survey respondents are actively seeking and working to deploy. The overwhelming majority of organizations responding to our survey (82%) say this trend toward distributed NG-IPS solutions is a key advantage in their strategies, as shown in Figure 15.

In this new "sensors" mindset, the traditional IPS now becomes just one more sensor.

Would a fabric-oriented approach to IPS involving other network operations and security infrastructure reporting sources be an advantage?

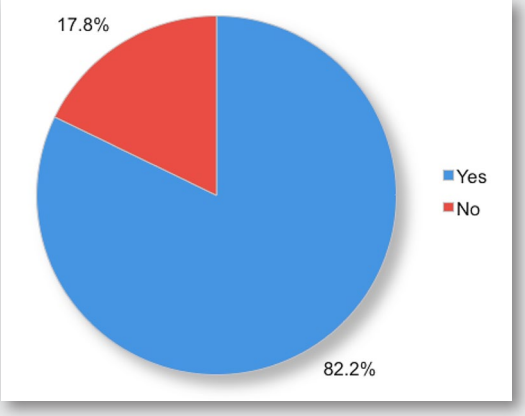

Figure 15. Acceptance of a Sensor Model of IPS

Many of the components to integrate an IPS fabric exist today; however, adding a security communications layer truly starts to describe an ecosystem of security sensors, with a central rule set, central data collection and reporting, as illustrated in Figure 16.

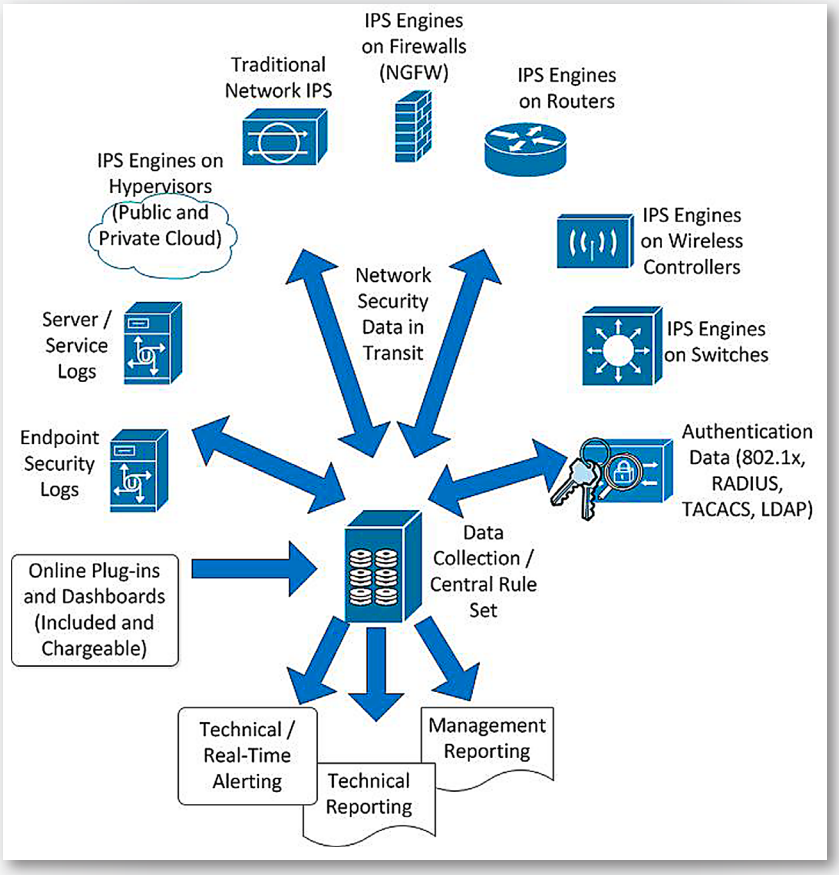

Figure 16. An Ecosystem of Security Sensors

As you can see, a lot of components and integration points are necessary for such a system to truly provide value.

Architecture Components

Rule sets, integration and standards will play a large part in how comprehensive this fabric-oriented approach will be. This includes:

- **The need for plug-ins.** It is likely that additional plug-ins will be necessary to import security data in much the same way as SIEM products operate today. Although many common plug-ins will probably be bundled into the data collection and reporting products, many plug-ins will likely be available as additional, chargeable items. In the near term, some plug-ins may have to be custom developed.
- **Correlation.** In a fabric-oriented IPS deployment, placement of the correlation engine that assembles all inputs now becomes critical. Bandwidth requirements and rate limiting of this data will become important to prevent security data from monopolizing bandwidth and affecting mission-critical production traffic.

- **Security and accessibility of the IPS data.** Protecting sensor data in transit will also become important. Because the quick transit of security data is important, the use of compression and encryption will become more prevalent in this space to help resolve these competing requirements.
- **A common language.** This communication also adds the requirement for a common language among all the new components in this remote sensor and reporting engine ecosystem. Although several such languages exist today, expect each vendor to try to standardize on a different one, limiting interoperability.

Most of the security vendors have a major push to develop a more "business language" reporting product that can be used with both IPS and SIEM data.

If a lot of what we've discussed seems like technology you can purchase today, that's because it is.

Security companies have been consolidating the product landscape, and many of the major IPS vendors have purchased or merged with a SIEM company. Most of the security vendors have a major push to develop a more "business language" reporting product that can be used with both IPS and SIEM data. Today's IPS, in its typical edge network deployment, is being rolled into the current crop of Unified Threat Management (UTM) firewalls. We see entirely new players who were not on the scene three to four years ago. Now, with significant portions of the firewall market, many IPS vendors have firewall functions and are growing into that space from the other side.

How Will This All Happen?

First, expect that the basic IPS engine will get broken up a bit. Expect to see purchasing an IPS become a more complex procedure in which organizations will purchase one or more IPSs for use on one or more platforms, with separate databases for data collection and reporting engines.

The basic IPS engine will likely become a commodity available on multiple platforms, including:

- IPS systems on dedicated hardware for high performance situations and "traditional" IPS placement will continue.
- We'll see more IPSs delivered as images for operation inside of a hypervisor. In hypervisor situations, the IPS might interoperate with one of the more advanced virtual switches. Over time we can expect to see more virtual IPS solutions with hooks directly into the hypervisor, using interfaces such as VMCI to collect IPS information directly, bypassing the virtual network.
- We'll see IPS systems deployed either in the firmware of a router or switch, as a virtual machine (VM) hosted within the router or switch, or as a card plugged into the router or switch backplane.
- As the real estate covered by these "remote instance" or "sensor" IPS systems grows, expect the price of this component to drop.

With the basic IPS now available on so many platforms, a central server for collection and correlation of IPS events will become a large part of the solution.

Next, we'll see the passive vulnerability scanner function broken out; in fact, we're seeing that right now. Think of this as the "low-hanging fruit" detector for internal systems. They operate mostly as "reconnaissance engines" within the IPS and are generally smart enough to identify hosts, operating systems and running applications. They are well placed to detect traffic differences that might identify operating system versions, application versions and any patches or updates that affect network traffic. So, if you have an out-of-date browser or application, or a host still running XP, a strategically placed passive scanner is a nice way to collect that information. Passive scanners are also nice in that they can detect simple version and patch issues without the intrusive and time-consuming scan that a traditional scanner or enumeration tool might require.

Event collection will become an integral part of this. With the basic IPS now available on so many platforms, a central server for collection and correlation of IPS events will become a large part of the solution.

Dedicated reporting products are being sold by many IPS vendors. A fundamental level of reporting will remain on the basic IPS engine, but a more fully featured reporting engine with multiple inputs, more dashboards and a much larger database for a larger history will be able to handle the data from all those other "remote instance" IPSs. This allows customers to get basic IPS function with an entry-level budget and allows them to grow into a more complete solution as the need arises and the budget allows.

With all this in place, it's not a large stretch at all to start adding other inputs, such as traditional syslog feeds from other hosts or network devices and Windows event logs. Just as this discussion is starting to stray from network-based IPS in this area and into SIEM territory, the IPS vendors are moving in that direction. If taken a few steps more, with inputs from endpoint protection, authentication logs and the like, what used to be the flagship IPS product quickly becomes one input of many into a much larger SIEM product.

Conclusion

As we can see from our survey results, the lines in network security are blurring. Integration and overlap between SIEM and NG-IPS products, overlap between management and technical reports, and attempts to make advanced training less of a requirement to configure and operate NG-IPS and SIEM products are all making for an interesting Venn diagram.

A trend we're seeing in security products is more data and better ways to represent it. On the input side, IPS products are becoming much more diversified, with inputs from more network devices to cover ever-larger portions of the network and with inputs from logs, endpoint protection services and other sensors and agents. As the volume of input data grows, IPS products are becoming more capable of gleaning more usable intelligence from this data, becoming much more application aware, user aware and context aware.

On the output side of IPS products, we're seeing similar growth, with enhanced reporting and easier report generation. Reports are becoming both more usable for technical personnel and more understandable for the managers, CSOs and directors that are ultimately responsible for the security of the organization. Creating reports and representing data in different ways is becoming easier and will require less specialized training as NG-IPS products use more dashboard and wizard-based approaches in creating outputs.

As NG-IPSs become more application aware and the number of inputs grows, we can look forward a shift in IPS function. The traditional IPS approach has always been to define "known bad" activity and alert and report on that. The mythical goal of IPS systems has always been to define "known normal" and alert and report on abnormal activity. With all of these changes in hand, we might just get there.

About the Author

Rob VandenBrink is a consultant with Metafore in Canada, specializing in security, networking and virtualization. He holds several industry certifications, as well as a master's degree with the SANS Technology Institute. He is a co-author of SANS SEC579—Virtualization and Private Cloud Security. Rob's current research projects include fiber channel security and antiforensics using FPGAs and GPUs. Rob is also an incident handler with the Internet Storm Center. Look for his posts at **http://isc.sans.edu**.

Sponsor

SANS would like to thank this paper's sponsor: