KNOWLEDGE *is* power

Tactical Approaches to Using Network Security Technologies for Deeper Insights

BY JOEL SNYDER, OPUS ONE

What you don't know can hurt you, and is likely costing you money and increasing your security risks during an era of scarce resources. As security technologies evolve to work faster and more efficiently amidst Web 2.0, multicast video, streaming media, and peer-to-peer use on networks, they also work more deeply to reveal insights essential to maintain high security with limited staff. **Most enterprises have taken good control of their perimeter, but lack both visibility and control of security within their networks.**

This white paper proposes six key strategies that enterprise security managers can use to improve their network defense posture. These include both *visibility strategies*, based on intrusion detection, network flow analysis and system vulnerability analysis, as well as *control strategies*, based on enabling additional security on existing elements, adding elements to the network, and applying fine-grained access control to end user connections.

Incorporating these strategies into your tactical security plan will add security visibility and security control to existing networks, which will reduce overall risk and give the security team additional agility to securely support the business goals of the enterprise.

Overview

Enterprise security managers have successfully deployed firewalls to provide perimeter security for many years. While perimeter firewall technology continues to evolve, such as in the recent inclusion of UTM (unified threat management) features, most of us consider perimeter firewalls to be mature technologies, commonly accepted, and well understood.

However, the relative stability and immobility of firewalls presents a problem for the security manager, because the threat landscape continues to evolve, both rapidly and aggressively. In particular, this "threat evolution" has occurred in the presence of solid perimeter firewalls, which means that the newest and most dangerous threats are specifically designed to operate in an environment with strong firewalls in place.

How, then, can the security manager counter these new threats and reduce overall security risk? One easy answer is to proliferate more firewalls, pushing protection deeper into the network. While this is a *useful* and *familiar* solution, it is not a **complete** solution. To gain greater coverage against new threats, it is important to add new strategies and tactics to your network security architectures.

Today's market is flooded with commercial and open source solutions all purporting to help improve network security. Begin selecting the right strategy to navigate the offerings and advice by considering what perimeter firewalls *don't* do:

- *Firewalls are great at blocking traffic, but they don't tell you much about the traffic they did **not** block*
- *Firewalls cannot tell you how well your network is operating, or whether your network is internally secure*
- *Perimeter firewalls cannot control internal traffic*

These come down to two main deficits: a lack of visibility, and a lack of control. The key strategies to improve network security is to add both visibility into the network and control of traffic within the network.

In this context, **visibility** means knowing what is happening on the network from a *security* point of view. However, because network management and security management directly affect each other, visibility also means knowing what is happening on the network from a *network* point of view. These two are not far apart and directly affect each other. Security services can be a bottleneck on the network; network speeds and new applications can strain security services.

Control, in this context, means changing the nature of the internal network from a pure utility where anything goes, to a secure corporate asset. Adding control is not simply adding more firewalls; it is enabling control points throughout the network to direct and manage traffic. The result is a change in thinking about the network away from the content-neutral utility of the high-speed switched LAN to an intelligent, managed, predictable, and secured com-

Sponsored by:

NOKIA

NETWORKWORLD®

munications channel linking the people and systems in the enterprise.

In the rest of this white paper, we will discuss five specific tactics that network security managers can use to add both visibility and control to enterprise networks.

TACTIC: To increase your ability to see security issues within the network, add network Intrusion Detection Systems (IDSes) and Intrusion Prevention Systems (IPSes) to your core and DMZ networks.

The world of IDS has come a long way since early commentators, misunderstanding the technology, declared it "dead" back in 2003. Of course, some of the blame for this confusion comes from the technology proponents themselves, who over-advertised and over-hyped the capabilities and use cases for IDS.

The main misconception surrounding IDS, and its follow-on technology, IPS, is that it will somehow catch (and prevent in the case of IPS) an external attacker from breaking into a network. In most enterprises, a properly configured firewall, good patch discipline, and well-written Internet-facing applications will prevent the types of Internet-sourced attacks that an IPS is designed to block. This gives the IPS limited applicability in that environment.

However, IDS and IPS can provide enormous amounts of visibility into the security posture of a network. A good analogy is to compare a properly deployed IDS to a network protocol analyzer. A traditional protocol analyzer lets a network analyst understand what is happening from a data communications point of view. Similarly, the IDS is a "protocol analyzer" for the security analyst. When properly located and managed, an IDS gives the security analyst deeper understanding into the security of a network.

While an IDS or IPS does have some value on Internet-facing traffic, the real value comes from increased **internal** visibility of security problems. Different products act differently, but in the hands of a security analyst, the IDS or IPS acts as a window into the security posture of the network. The information provided by the IDS or IPS will help the security and network management teams uncover, as a start:

- *Security policy violations, such as systems or users who are running applications against policy*
- *Infections, such as viruses or Trojan horses that have partial or full control of internal systems, using them to spread infection and attack other systems*
- *Information leakage, such as running spyware and key loggers, as well as accidental information leakage by valid users*
- *Configuration errors, such as applications or systems with incorrect security settings or performance-killing network misconfiguration, as well as misconfigured firewalls where the rule set does not match policy*
- *Unauthorized clients and servers including network-threatening server applications such as DHCP or DNS service, along with unauthorized applications such as network scanning tools or unsecured remote desktop.*

If IDS is so useful at catching security problems, and IPS at blocking them, why haven't we seen enthusiastic acceptance? The answer lies in the management of data from IDS and IPS sensors. Until recently, it was virtually impossible to get good information out of IDS and IPS products because their management and analysis tools were so poorly written and mis-matched to the needs of the enterprise security analyst. IPS and IDS detection engines provide an enormous amount of information—some of which can be completely irrelevant. With proper management tools, this information can be filtered and qualified, turning raw data into actionable information.

In other words, writing an IDS or IPS sensor is easy (or at least relatively so; there are good and bad IDS/IPS sensors). Doing something useful with what you get from that IDS or IPS sensor is hard, and IDS/IPS and Security Information Management (SIM) vendors have only recently solved this hard problem. In the past, most enterprises focused their IDS acquisition process on small differentiations between the sensors, without really diving deeply into the management capabilities. For example, the topic of IPS evasion, a fascinating discussion item for researchers, received an immense amount of attention and focus. At the same time, the correlation between IPS alerts, security policy, and vulnerability information, has been virtually ignored—although this correlation is much more important to the enterprise looking for useful results from their IPS in real-world networks. The result was that the IDS or IPS failed to detect and block any significant number of intrusions (making the product not very cost-effective) and the analyst had poor visibility into the security posture of their network.

This tactic of increasing visibility using IDS and IPS technology, then, is not really about identifying and blocking intruders sneaking into your network from the Internet. This tactic is about using tools such as IDS and IPS for what they are best at: giving visibility into the security posture of your network identifying security, network, and performance threats, and minimizing the risk of those threats. A good IDS or IPS management system combines the raw data of the sensor with other content: network topology information, host application information, system criticality ratings, and vulnerability information. This is all correlated together to turn the raw data into actionable information, something the security and network manager can use to take intelligent action.

THE WORLD OF IDS HAS COME A LONG WAY SINCE EARLY COMMENTATORS, MISUNDERSTANDING THE TECHNOLOGY, DECLARED IT "DEAD" BACK IN 2003

TACTIC: To gain better insight into traffic and flows within the network, collect and analyze security and flow information from existing control points.

Most networks already have an abundance of instrumentation for collecting information about the traffic flowing through them. Devices such as firewalls, routers, switches, load balancers, and even the end systems themselves are capable of providing a great deal of information about what is happening inside of the network. Some of this information is available directly as a result of logs, such as from firewalls, IPSes and IDSes, or load balancers. In other cases, you may have to enable data collection by turning on features such as flow tracking with the Netflow or IPFIX protocols.

Traffic and flow information can help answer critical questions about the network: which systems are originating connections? Which systems are receiving connections? Who is talking to who? And how much traffic is flowing in each direction?

This strategy is a visibility strategy, but at the network layer more than the security layer. This strategy also represents an opportunity for the network and security teams to work more closely to understand what is happening on the network. Because each team has some of this data, and is interested in the results, a key value of collecting and analyzing flow information is bringing together data sources from different parts of the network to expose the traffic at as many layers as possible.

When flow and traffic data are aggregated together, the network and security teams will be able to accelerate both their optimization tasks—refactoring or extending the network as needed to reduce delays and speed performance—and their debugging tasks. Questions such as "what application is using all the bandwidth," "who is running a rogue mail server," "what's that guy doing over there," and "what services are being used most" all become easy to answer when the data are available.

This all seems so obvious that network and security managers who are already taking advantage of this information may wonder why it's worth saying. The reason is simple: few enterprises are actually looking at this data in any depth. Simple open source products, such as the MRTG traffic grapher, are often the only traffic analysis tools being used. While MRTG gives a great overview of absolute bandwidth, it simply looks at counters—not traffic flows.

Asking for the data from existing devices is easy; analyzing the data may require additional tools.

In many cases, existing security tools such as firewall and IPS management systems not only handle security monitoring but can also provide traffic flow information useful to both the security and network teams. When these tools aren't available, more specialized products that aggregate traffic log information (such as Security Information Managers, SIMs) or that report on flow information (such as Netflow analyzers) can be added to the network. The key feature to look for in implementing this tactic is the ability to aggregate, summarize, and drill-down into the data. It's certainly useful to know what a single server is doing, but it's also just as important to know what all the servers in a rack are doing, or what all systems running a particular application are doing.

Examining traffic flow data may pose a different challenge: the traffic will generally be expressed in terms of IP addresses or even MAC addresses if it comes from switches and routers. While that's useful, people think in terms of applications and end users. To help put these traffic and flow data into the most useful context, look for network analysis tools that can integrate with other systems, such as the enterprise's directory or asset management database.

Flow data and traffic data in the large can also be complemented by traffic information in a more tightly defined context. For example, a firewall may be able to tell you how many flows would benefit from encryption acceleration.

A key part of implementing this tactic is focusing on the tools and devices you already have. While you may need to buy or add something to help with analysis, you should be able to make better use of the resources you've already paid for, and deployed deep in your network, to get greater visibility. If you're thinking of buying new network or security devices, consider their ability to export, display, or analyze traffic flow information as a product differentiator.

IN MANY CASES, EXISTING SECURITY TOOLS SUCH AS FIREWALL AND IPS MANAGEMENT SYSTEMS NOT ONLY HANDLE SECURITY MONITORING BUT CAN ALSO PROVIDE TRAFFIC FLOW INFORMATION USEFUL TO BOTH THE SECURITY AND NETWORK TEAMS.

TACTIC: Gain greater and more granular control over all traffic by enabling security on devices you already own, such as firewalls, switches, and routers

Enterprises have spent considerable time and effort buying high-end networking and security devices—and often under-utilize the security and control features of these devices. This can occur for organizational reasons, for example, if the network team and security team are not well integrated. But it also occurs simply because network and security managers have shied away from enabling security features beyond the bare minimum required to operate the network securely.

While vendors are happy to push new hardware and software to solve any new security problem, many enterprises can increase total security control using equipment they already have, that is already paid for, that they're already trained on, and that is already integrated into

their network. Existing equipment stretching from the perimeter and outer edge of the network all the way to the core can be turned into additional security control points at almost no expense.

For example, while an external Internet router is not a firewall, it can be a good "first cleaner" of incoming traffic, blocking obvious services that will never transit the external firewall, such as traditional NETBIOS traffic or SNMP traffic. External routers can also protect networks against spoofing. And, the external control plane of the network, between the external router and the firewall, can only be protected by the external router.

Even security appliances can be better utilized. All enterprise-class firewalls have a combination of access controls and some intrusion prevention features, such as denial of service protection, built-in. While an enterprise with a traditional IPS will definitely want the IPS to sit inside that external firewall, there's no reason why some basic IPS or denial-of-service protections cannot---or should not---be enabled in the external firewall. These can both complement the existing IPS, since many IPSes do not have denial-of-service protections, and can backstop the IPS.

Of course, this doesn't mean that every security feature in every device should be turned on willy-nilly. When a feature is enabled, the network manager has to have the confidence that the feature will not cause false positives, and that the alerts for that feature will properly propagate to the management console. However, areas such as control-plane management on the firewall itself and any connection to the external network are so obviously useful that they should be enabled.

Internally, security features on switches and routers should also be used where appropriate. *Table 1* lists some typical features on network infrastructure and security devices that are often ignored:

Extending existing infrastructure to gain security is economical in terms of capital and operational expenses. The goal of this tactic is to leverage existing control points within and surrounding the enterprise network. Buying new hardware and software is fun and exciting, but network managers can gain a significant bump-up in security by working closely with their existing LAN vendors to enable features that they have already bought and paid for. Similarly, security managers can often take their existing firewalls, devices they're already happy with, and do more with them.

TACTIC: Better understand the security posture of your network by making use of active or passive vulnerability analysis and network discovery tools

In the typical enterprise network, the provision of network services is distinct from the management of applications and application servers. The "application team" tells the network team what they need and when they need it, and the "network team" builds the network and services to support those requirements. This can extend to the security team as well: the application team will identify application ports and protocols, and the security team is responsible for managing traffic to those ports and protocols using the normal combination of firewalls, IPSes, and other threat mitigation tools.

Unfortunately, what you don't know *can* hurt you. As the application environments get more and more complex, and as tools such as virtualization cause easy proliferation of servers and services, the network and security teams may end up supporting a network with an unknown configuration: application communications that are not documented, application servers that are not properly managed, and test and lab systems that are scattered on production networks.

The network and security teams are obligated to know about these servers, these services, and these applications. They cannot build a production network service without this knowledge, particularly as enterprise network security grows more and more granular and pushes in towards the network core.

Gaining greater visibility into the network and these application servers and systems can be difficult if the application team doesn't have the information or doesn't see it as a priority. For that reason, a desirable strategy is to use active and passive network discovery tools to gather the information. Of course, no one wants to step on the toes of the applications and server teams, but this doesn't mean that the network and security teams should operate in the dark just because they can't get the information out of the application team in an accurate and timely way.

Having an accurate application layer map of the network is especially important when advanced threat-mitigation technologies, such as IPSes, are in place. While it is tempting to believe that an IPS can run successfully without any real

Table 1

External Router	External Firewall	Internal Switches and Routers
Anti-spoofing ACL	Basic DoS protection; anti-spoofing protections	Control plane management; separation of control from data plane; encrypted management
Control plane protection; encrypted management	IPS protections (complementary to true IPS)	Traffic management (anti-spoofing in both directions; worm detection; blocking unwanted/impossible traffic)
Blocking obviously unwanted traffic (SNMP, etc.)	Control plane protection; encrypted management	DHCP snooping; Dynamic ARP-detection
	Outbound traffic blocking (no pass-all rule!)	Link logging to existing security tools

knowledge of the underlying network, experience suggests otherwise. A well-tuned IPS is matched to the applications and systems running on the enterprise network to provide the highest catch rate for both security and policy violations, while minimizing false positives.

A combination of active scanning (with a light touch), useful for basic network discovery, and passive scanning, used to identify and map applications and systems that are otherwise invisible to active scanners, is the preferred approach. Active scanning has a high technical and political cost (see sidebar on Active versus Passive scanning technologies), but also brings benefits that passive scanners cannot, such as vulnerability analysis. Passive scanning avoids some of the pitfalls of active scanning, but is less precise in most cases. For example, vulnerabilities can be inferred, but it is difficult to determine if they have been patched. However, passive scanning also has technical advantages, such as the identification of applications that are actually being used and traffic/flow analysis for those applications.

In some cases, especially when vulnerability scanning is desired, the network, security, and application teams can jointly run a full-fledged scanner and patch management system. This brings benefits to each group, and helps to tie them more closely together, a desirable outcome in most enterprises.

TACTIC: Ensure only authorized and "safe" users connect to the network by using Network Access Control (NAC) to authenticate, validate, and control all network usage.

While perimeter security tools such as firewalls and VPNs have been strong guardians at keeping unwanted visitors from enterprise networks, the soft, chewy center of the corporate network often has few protections from unauthorized connections. Network Access Control, NAC, has emerged in the last two years as a set of technologies that can help extend strong access controls to network interiors.

NAC is often oversold as a product: you should "buy a box of NAC" or something like that. In reality, NAC is a new philosophy of network management and security. Prior to NAC, the general thinking about access controls in enterprise networks was to focus on perimeters and obvious choke points, such as DMZ networks or server farms. With NAC thinking, the network is turned inside out. Instead of having access controls at the perimeter or at major entry points, a "NACified" network is one where every point of connection becomes an access control point, hence the name: Network Access Control.

One of the most important things that NAC does is focus on the user. Rather than trying to apply controls based on IP addresses and subnets, NAC moves the security decision to the actual user making the connection. It is this change in focus, from the subnet towards the user, that represents what is important about NAC.

While there are products that can help in implementing NAC, the main point of this strategy is not to suggest running out and implementing some vendor's idea of NAC. Instead, the tactical approach calls for taking the key ideas of NAC: user authentication, access control, and end-point security, and moving those ideas out to the wiring closet and all points where end-users connect to the network. When NAC thinking is applied to enterprise networks, the network changes from an "anything goes" utility to a secured access point for enterprise applications. If the control goal represents half of this white paper, then NAC is the idea of control taken to its logical conclusion.

Because NAC is a change in philosophy, and a set of technologies, rather than a product, each enterprise will choose different parts of NAC to use in their network. Some will focus on end-point security, others on fine-grained access controls, and others on authentication---or a combination of all three. For greatest success, network and security teams will work together to determine what the real goals for NAC are, then deploy the technology elements that support those goals.

One of the challenges that accompanies NAC is the inter-disciplinary nature of the technology. End-point security within NAC is rightfully within the area of expertise of the desktop or Windows team. Access controls, whether expressed as VLAN selection, Access Control Lists (ACLs), or simple Go/No-Go decisions, have to be managed by the network team. And the actual NAC policies and access control rules will have to be defined by the security team. Bringing these three teams together is a critical step in any successful NAC deployment.

DESPITE THESE CHALLENGES, IT IS IMPORTANT TO REALIZE THAT ALMOST ANY AMOUNT OF NAC THINKING INJECTED INTO THE NETWORK WILL BE A HUGE INCREASE IN SECURITY.

Despite these challenges, it is important to realize that almost any amount of NAC thinking injected into the network will be a huge increase in security. For example, simply dividing a network between guest and non-guest users (or even blocking all guests from attaching to the same wired network as staff) will dramatically reduce the window of opportunity for malware to accidentally spread from an unprotected system. While NAC products and services inhabit a spectrum from "a little bit of control" to "a huge amount of control," almost any first steps towards NAC adoption are a great leap forward in security controls.

TACTIC: Look ahead to new applications and new SLAs by instrumenting your network to measure performance and plan for growth in security.

Much as we'd like to stop growing and building our networks, it's not going to happen anytime soon. New applica-

tions (with tighter service level agreements, SLAs) and greater use of existing applications will drive higher levels of traffic throughout the network. This means that planning for growth is an integral part of day-to-day operations of each network. As networks and security features merge over time, this means that planning for security growth is also a daily task.

Two techniques can help ensure that security is not a bottleneck to new applications. First, properly instrument your network to collect statistics and performance measurements so that you know the true performance of the network and the underlying security technologies. Secondly, identify points where security features are a potential bottleneck, and design for the growth and scalability of those points.

When gathering statistics, you'll want to go beyond simple measures such as throughput (although throughput is still a very, very important measure) and reliability and include other metrics, including latency, jitter, and even connection rate in your reporting. While the traditional applications have long-lived connections for which throughput and latency are most important, newer applications, especially multimedia ones, are more sensitive to other network performance metrics such as jitter and connection rate limits.

Since most enterprise networks will have some voice applications running on them, QoS has become a common configuration option. If you have QoS, you also want to collect statistics on how traffic is being put into different QoS queues, and what the performance of each of these queues is: how many streams are being rate-limited, what is the utilization of each QoS queue, and, if you can get it, how often do congestion events occur and how long do they last?

This visibility goal of gathering all these statistics and performance metrics is to provide a planned and gentle upgrade path for security and network equipment, avoiding application bottlenecks, rip-and-replace and forklift upgrades. This means that security points within the network should be identified, and then designed to scale in performance as the application requirements increase.

Traditional security products, such as firewalls, are generally tested and specified to work with "typical" traffic loads, either large-size UDP packets (if the vendor wants to have a very high performance number) or with HTTP traffic (if the vendor wants to have a worst-case performance estimate). New uses of the network, such as VoIP, multimedia, and mobility applications, may fit neither of these two profiles. While it's tempting to guess that a new application performance will fit somewhere between the best case (large UDP packets) and worst case (many HTTP sessions), that's not always a good assumption. It is always possible that a new application resets the bar for "worst case performance" in a product. This is why instrumentation, data collection, and monitoring of network performance are critical visibility strategies that pay off when application mixes change.

Final Thoughts

Throughout this white paper, we've emphasized the twin goals of *greater visibility* and *greater control* within enterprise networks. It's important not to confuse either one with greater security. Both are a means to provide greater security, and both may be essential to pushing the security bar further up. But simply having a lot of visibility, if you don't do anything with the information, isn't going to improve security. Similarly, putting a lot of control points into a network won't help security, without a security policy that says what you're controlling and why. ■

Goal	Strategy	Tactic
Visibility	Increase your ability to see security issues within the network	Add network IDS and IPS to your core and DMZ networks
Visibility	Gain better insight into traffic and flows within the network	Collect and analyze security and flow information from existing control points
Control	Gain greater and more granular control over all traffic	Enable security on devices you already own, such as firewalls, switches, and routers
Visibility	Better understand the security posture of your network	Make use of active and passive vulnerability analysis and network discovery tools
Control	Ensure only authorized and "safe" users connect to the network	Use Network Access Control (NAC) to authenticate, validate, and control all network usage
Visibility	Look forward to new applications and new SLAs	Instrument your network to measure performance and plan for security growth

gettingreal

Applying Intelligent Threat Mitigation in Enterprise Networks

Enterprise network managers are bombarded with advice on how to do their jobs. The Internet has unearthed an unlimited supply of experts ready to blog on the "right approach" to network security, teamed with an equally unlimited number of PowerPoint presentations on which products they should be buying to protect themselves from this afternoon's new threats. Meanwhile, in uncertain economic times, every security decision and acquisition has to be accompanied by an explicit question: "is this the best way to be spending our money, ***now?***"

Unfortunately, standing still is not an option. New threats are appearing on a daily basis. Work styles are changing, as more and more people blur their work and personal lives, driven by the proliferation of mobile devices and anytime/anywhere connectivity. And applications themselves are changing, with Web 2.0 and multimedia putting a very different load on enterprise networks.

The solution to this problem is **Intelligent Threat Mitigation**. Deploying security technologies within the enterprise network that help you work smarter, by giving you visibility and control. Intelligent threat mitigation helps you know your network, unburdens your security staff, and extends the life of your security infrastructure.

The first step towards intelligent threat mitigation is knowledge: know your network, because what you can't see *can* hurt you. As Joel Snyder points out, Intrusion Detection and Prevention (IPS) technology is a key part of network knowledge. Nokia's partnership with Sourcefire brings the most visionary IPS developer together with the best-known security appliance manufacturer. With 90% of the Fortune 500 depending on Nokia's security products, the combination of Sourcefire's technology and Nokia's uncompromising quality, global support, logistical expertise, and proven track record offers best-in-breed software combined with best-in-breed hardware, operating systems, and support.

Nokia IPS is a critical tool to gain knowledge and insight into enterprise networks. While most first-generation IPS products have focused on stopping unlikely attacks, Nokia IPS is different. Nokia IPS combines the power of Snort—the de-facto standard for intrusion prevention with over 150,000 active users—with real-time endpoint intelligence, user identity tracking, and automated IPS tuning into a scalable and manageable network knowledge tool. Nokia IPS certainly detects and blocks intrusions, but it also provides a key additional benefit: network knowledge.

NOKIA IPS IS A CRITICAL TOOL TO GAIN KNOWLEDGE AND INSIGHT INTO ENTERPRISE NETWORKS. WHILE MOST FIRST-GENERATION IPS PRODUCTS HAVE FOCUSED ON STOPPING UNLIKELY ATTACKS, NOKIA IPS IS DIFFERENT.

With Nokia IPS in place, the network manager knows what systems are on the network, what applications they are running, who is using them, and how much traffic is flowing between different devices. All this critical contextual information supplements the prioritized list of security problems, updated constantly, with the highest priority problems on the highest priority systems at the top of the list. By combining context with security information, Nokia IPS turns data into knowledge.

Network knowledge is not restricted to IPS platforms. Network managers who have chosen Nokia's IPSO-based security appliances for Check Point firewall deployment have tools to improve their visibility into their own networks. An advantage Nokia brings to the table is that Nokia IPSO is not just a bare-bones operating system—it's a full-featured platform, with a hardened operating system, integrated routing protocols, load-sharing active/active clustering, and a secure remote management toolkit to maximize network availability.

As a visibility tool, Nokia IPSO lets the network manager drill down into traffic flows through the security appliance to understand traffic patterns, connection lifetimes and sizes, session and transaction rates, packet sizes, and load on the security appliance. These visibility features in Nokia IPSO provide actionable intelligence to the network manager, enabling them to understand and optimize traffic, troubleshoot problems, and do capacity planning. Nokia IPSO can also export traffic information using Cisco's popular NetFlow™ data format for analysis by other tools—including Nokia IPS.

The second step towards intelligent threat mitigation is unburdening your staff. A network that is unmanaged, is a network that is unprotected. The key is to get valuable staff

out of repetitive and low-value tasks so that they can focus on what's important: managing the network. Using tools such as Nokia Horizon Manager, Nokia's centralized security appliance management tool, along with a consistent set of powerful platforms, like Nokia's IPSO security appliances, frees up staff time to let them concentrate on areas where critical thinking and human insight are important. For example, when an educational government agency saw the task of rolling out over a dozen firewalls to support 70,000 users, Nokia Horizon Manager quickly paid off by reducing installation time and management effort, keeping focus on critical issues of 10GbE performance and latency management in an environment of quickly changing traffic patterns.

Nokia Horizon Manager helps to unburden staff by automating time consuming tasks, such as running backups, upgrading appliances, making global changes, and provisioning critical applications. Because Nokia Horizon Manager looks across all Nokia security appliances, it also provides real-time monitoring of appliance information as well as traffic data, along with alerting. And, with Nokia Horizon Manager, it's easy to roll out new devices error-free, because Horizon Manager can ensure consistent appliance and application settings. Your "best practice" rules for deploying firewalls are enforced in the management console, which helps to manage complexity and prevent errors.

Another key towards unburdening your staff is using smarter tools in the first place. For example, with Nokia IPS, the network manager has a "security dashboard" available which gives an at-a-glance view into the most critical events and the most important metrics. Rather than wasting valuable time trying to decide "is there something I need to know about," the network manager can glance at the dashboard and immediately see whether any security events need their attention.

The third step towards intelligent threat mitigation is extending the life of existing infrastructure. Installing new hardware and configuring new software is expensive, time consuming, and can interrupt network services. A smarter strategy is to select security platforms that will be long-lived, reusable, and scalable on everything from a DSL line to a 10 Gbit/second connection.

The Nokia IP Security appliance line is a very cost-effective way to standardize on a single platform that can support firewall, VPN, and IPS applications. With a single operating system, a single management interface (and central management through Nokia Horizon Manager), and interchangeable modules, IPSO security appliances can scale up to the largest enterprise networks—and represent a minimum learning curve for network and security engineers. Unlike other vendors who change architecture, hardware, operating system, and management system from year to year, Nokia IPSO provides a single foundation for the entire product line that retains knowledge and experience, minimizing your costs when it comes time to grow or upgrade.

Although Nokia IPSO offers a long life and platform stability, this does not mean that new technologies are not available in the Nokia IPSO product line. In addition to high-speed networking, including both 1GbE and 10 GbEthernet, Nokia IP Security Appliances take full advantage of Intel's new multi-core architectures, with both dual-core and quad-core models available—which lets you take full advantage of Check Point's CoreXL multi-core acceleration technologies for additional performance in the same chassis.

Nokia IP Security Appliances also make use of Nokia's Accelerated Data Path Services Modules, which provide a speed boost to existing hardware by accelerating Check Point firewall performance up to 10Gbit/second per module—without a new license or new application. For example, a large international bank was able to replace Check Point Secure Platform hardware based on commodity servers by standardizing on Nokia IP2450 security appliances and ADP services modules. This was both a business decision, based on Nokia's exceptional support track record and solid reliability, along with a technical decision, because the new hardware platform reduced latency and increased performance, without having to change firewall platforms. Alternative chassis-based solutions from other vendors were dismissed as not cost effective for the same levels of performance. ■