



Missing Link 
Security Services
Mark Bouchard, Founder

Securing Internal Networks: The Evolving Role and Requirements for Intrusion Prevention Systems for Public Sector Organizations

Sponsored by: Juniper Networks

Introduction	2
The Perimeter is Dead! Long Live the Perimeter!	2
The Increasing Elusiveness of Threats	2
The Proliferation of Points of Entry	3
The Compliance Mandate	3
A Concise Strategy for Internal Network Security	4
What "Ideal" Really Means	4
The Countermeasures that Matter Most	4
A Closer Look at IPS for Internal Networks	5
The Devil is in the Differences	5
Similarities Matter Too	7
Conclusion	7
About the Author.....	7
A Word From the Sponsor.....	8



Government and public sector agencies worldwide are under mounting pressure to better secure their internal networks. For many of these organizations, the continued occurrence of successful attacks is clear evidence that perimeter-centric defenses are not sufficient. Overall, this situation is not surprising given the assumptions that need to hold true for a perimeter-based security strategy to be effective, including that:

- Perimeter controls will reliably stop all threats they encounter;
- There are no alternate paths for external threats to take; and,
- The threat from internal sources is negligible.

None of these have ever been completely safe assumptions. More troubling, though, is that ongoing changes to the computing landscape continue to erode their validity.

This paper explores the prevailing conditions driving the need for enhanced, internal network security as a prerequisite for establishing a logical set of technical countermeasures that will help meet this increasingly common objective. The focus then shifts to one of the most promising of these countermeasures, the network intrusion prevention system (IPS). Details are provided on how this historically perimeter-centric control has evolved, and the key requirements that should be met for a solution that will be operating at the core of the network, as well as at other strategic locations internally.

To be perfectly clear, the perimeter is not really dead. There is no denying that a number of conditions have emerged in recent years that elevate the importance of establishing better defenses for internal networks – an aspect of their security architectures that most public organizations have spent relatively little effort on in the past. However, the resulting shift to address internal network security does not mean that the conventional Internet perimeter can now be ignored. Doing so would be contradictory to the principle of defense-in-depth. Instead, a more balanced approach is required. Ideally, multiple, complementary layers or “perimeters” should extend from well-defined points of entry and natural chokepoints all the way to the information resources that are ultimately being accessed.

In any event, the changes driving this new reality can be grouped into three categories: those that pertain to the threat landscape, the operational landscape, and the regulatory landscape.

Motivated less by gaining notoriety and more by making money, hackers have accelerated their efforts and are now placing greater emphasis on evading commonly deployed countermeasures in order to obtain valuable information. Threats are being generated more quickly than ever before, thereby limiting the effectiveness of patching and other reactive countermeasures. In addition, an increased level of creativity is leading to threats that are both more diverse and more elusive. The predominant concerns of the past such as file-level viruses, worms, and denial-of-service (DoS) attacks are being over-shadowed by an array of newer contenders that include spyware, spear phishing, keylogging trojans, and rootkits. Even more troubling are targeted attacks, where hackers build custom exploits to take advantage of characteristics and weaknesses unique to a specific organization’s computing environment. And then there is the issue of threats migrating up the computing stack. By focusing on system and application-layer weaknesses, malware can be designed to slip through the network-layer countermeasures that dominate most organizations’ defenses.

These developments are particularly troubling for government and other public sector organizations for many reasons, including the following.

- They exacerbate the mounting concern over the possibility and likelihood of electronic invasion and/or disruption by foreign entities.
- They substantially increase the risk levels applicable to the significant amounts of the nation’s critical infrastructure, classified information, and public safety services operated and/or overseen by federal, state, and local government agencies.
- They further complicate and increase the cost of ongoing continuity of operations (COOP) initiatives by driving up the frequency and diversity of incidents.
- They jeopardize the integrity and availability of large portions of public sector network infrastructure which, due to chronic underfunding/underspending, still depend on all sorts of legacy systems – the point being that such systems are highly vulnerable and, at the same time, relatively difficult to remediate due to patches not being available in a timely manner, if at all.

The result of the increasing elusiveness of threats, in any event, is that a security solution must incorporate an



increasingly wide range of protective mechanisms to be effective, including ones with greater visibility and control at the application layer.

Complicating matters further are the various forces driving practically every public sector organization to enable higher degrees of user mobility, remote-office interconnectivity, and external access to their networked systems. For instance:

- Increasing the ranks of teleworkers is a high-priority initiative. After all, enabling employees to work from home improves morale and productivity, reduces costs, and helps save the environment.
- Data center consolidation is another focal point for many agencies. The primary motivation is cost savings and improved operational efficiency, but an inherent consequence is that users in distributed offices must now access their applications remotely.
- Even without consolidation, however, there would still be a significant need for supporting remote connectivity to resources – in some cases from untrustworthy locations – due to the highly distributed nature of many government organizations (for example, the U.S. State Department) and necessary business travel.
- In an effort to raise the level of service afforded their constituents, many departments are seeking to improve interagency coordination via deeper degrees of network and systems integration.
- And then there's e-government, a vast effort to better and more efficiently serve the public by providing the masses with electronic access to a steadily growing set of government applications and systems such as online tax services and vehicle registration, to name a few.

Overall, the impact is the introduction of more points of entry for threats and a general erosion of the traditional perimeter. *This creates the need for more layers of protection, particularly controls that can be leveraged for all scenarios (both internal and external) by being located "within the boundaries" and in closer proximity to the resources being accessed.*

Depending on your perspective, the current batch of regulations pertinent to IT could be viewed either as a consequence of the previously discussed conditions, or as a completely separate driver. In either case, there are two important points that need to be acknowledged.

First, public sector institutions are not immune to the compliance mandate. If anything, their situation is even more complicated. Sorting through and rationalizing the various layers of legislation, oversight activities, and a plethora of disjointed policies can easily become overwhelming. For example, U.S. federal agencies must comply with a flood of statutes, executive orders, and official memoranda that pertain to their information systems. The Computer Security Act of 1987, the Federal Information Security Management Act of 2002 (FISMA), and a steady stream of policy statements being issued by the Government Accountability Office are just the tip of the iceberg. In addition, many departments must also comply with commercially applicable privacy and integrity legislation like the Health Insurance Portability and Accountability Act (HIPAA).

The second point to realize is that a thread common to many of the applicable regulations is the requirement that organizations (a) have a comprehensive information security program to ensure the integrity, privacy, and reliability of their IT operations, and (b) that this program should incorporate more than just perimeter-based controls to adequately prevent unauthorized access to critical systems and private information.

The net result of the above issues is that improving controls on internal networks is an important initiative for today's public sector organizations – or at least it should be!

But what does this mean in terms of specific security capabilities?

When it comes to improving controls on internal networks, numerous countermeasures ultimately deserve consideration. Yet given the size and scope of the internal computing environment, very few public sector organizations have the luxury of evaluating and deploying them all.

On the surface, an ideal solution would seem to be for public sector IT departments to replicate the full set of typical demilitarized zone (DMZ)-based controls at the "edges" of their internal networks, such as at or near all workgroup level switches. Maximum effectiveness could be achieved as threats would be cut off before they had any chance to propagate. The problem, at least with the current generation of security solutions, is that this approach is completely impractical. The cost and complexity in terms of security software, hardware, and



manpower would be overwhelming for all but the smallest of networks.

In contrast, a better balance of effectiveness, efficiency, and practicality can be realized by implementing a combination of controls at both the core of the network and at the interfaces to high risk/value segments, such as guest networks, wireless local area networks (WLANs), and any enclave processing classified information. In this case, threats could still gain entry into portions of the internal computing environment, but they would be prevented from spreading to the most critical areas and resources.

With a rough plan in hand, the next issue to address is what types of controls make the most sense and in what order. Regarding the type of controls to consider, there are three complementary countermeasures that form the foundation of an ideal solution:

- The ability to improve containment by isolating critical resources or otherwise minimizing access to them by using virtual LANs (VLANs), internal firewalls, Network Access Control (NAC), or other techniques to granularly enforce access control;
- The ability to proactively exclude sessions initiated from high-risk endpoints (by implementing host-integrity checking); and,
- The ability to filter allowed sessions for threats (by deploying intrusion prevention technology).

Regarding priority, there are two basic scenarios that require attention. First, if the location under consideration is at the core of the network, or elsewhere on the internal network for that matter, then emphasis should be placed primarily on threat filtering and secondarily on the other security mechanisms. This arrangement provides the greatest effectiveness with the least amount of negative impact. It stops threats reasonably well, yet is transparent to users and minimizes the likelihood that important sessions will be unnecessarily impeded.

The second scenario is when the location under consideration is closer to the edge of the network or at the interface for a sensitive zone, and the population of associated users and applications is smaller and more cleanly defined. In this case, an argument can be made to focus first on improving access control and containment. The administrative burden associated with setting and maintaining rules and policies should be more manageable and the likelihood of incorrectly blocking important traffic will not be as great. As a result, IT will be well served by taking greater advantage of the considerable stopping power of positive-model (deny by default) tools that inherently block threats by dropping all sessions that are not explicitly allowed by policy.

To be clear, it's not really a matter of choosing one countermeasure and not another in these cases. Ideally, all three countermeasures should be used. It's just a matter of moderating the extent they are relied upon in different scenarios to achieve an optimal balance between security effectiveness, administrative efficiency, and the continuity and productivity of critical communications.

Regardless of the approach that is ultimately pursued, it should be clear that threat filtering – typically in the form of network intrusion prevention technology – is a key component when it comes to securing internal networks.

When it comes to addressing internal security, the temptation is to treat it identically to perimeter security. To some extent, this does make sense. Many principles used to construct and operate perimeter security designs apply equally on the inside, such as the practices of hardening critical systems and establishing defense in depth. However, other principles do not work quite as well. For example, having an out-of-band management network for the internal environment is simply impractical from the perspective of cost and complexity.

The point is that to be suitable for internal network implementations, intrusion prevention solutions must evolve to account for several differences between the perimeter and internal computing environments, while still addressing core capabilities as well.

It is considerably more demanding for an intrusion prevention solution to operate internally – especially at the core of public sector networks – than it is to operate at external access points (see Figure 1). Specific considerations and how they can be accounted for are identified below.

Internet boundaries and perimeter connections rarely involve link sizes and traffic rates that exceed 100 Mbps. In contrast, internal networks are routinely sized with 100 Mbps, 1 Gbps, and 10 Gbps connections, and



typically experience traffic rates ranging from a few hundred Mbps at the edges to several Gbps at core and data center locations. Latency has also become a more significant issue with the rise in usage and growing dependency on time-sensitive applications such as voice over IP (VoIP), real-time collaboration, and those relying on rich media or video. In addition, there is growing contention between legitimate, mission-oriented applications and traffic that is far less “essential” such as personal Web surfing, Internet chat, and YouTube videos. Consequently, to be suitable for internal networks, an intrusion prevention solution should include:

- (a) A portfolio of appliances with models that support *real-world* throughput levels of a few hundred Mbps, a few Gbps, and 10+ Gbps; and,
- (b) Quality of service (QoS) features that help control latency for critical business and time-sensitive applications.

Special care should be taken when evaluating solutions in this area, however. Rated throughput levels are often based on ideal conditions and configurations, and they fail to adequately account for a mix of traffic types, threat loads, a full complement of policies and responses, or the resulting latency that is incurred. Results of tests performed or sponsored by vendors should be closely examined. Better still, it is recommended that public sector IT departments conduct their own tests.

Figure 1 – Why High Performance is Needed at the Network Core

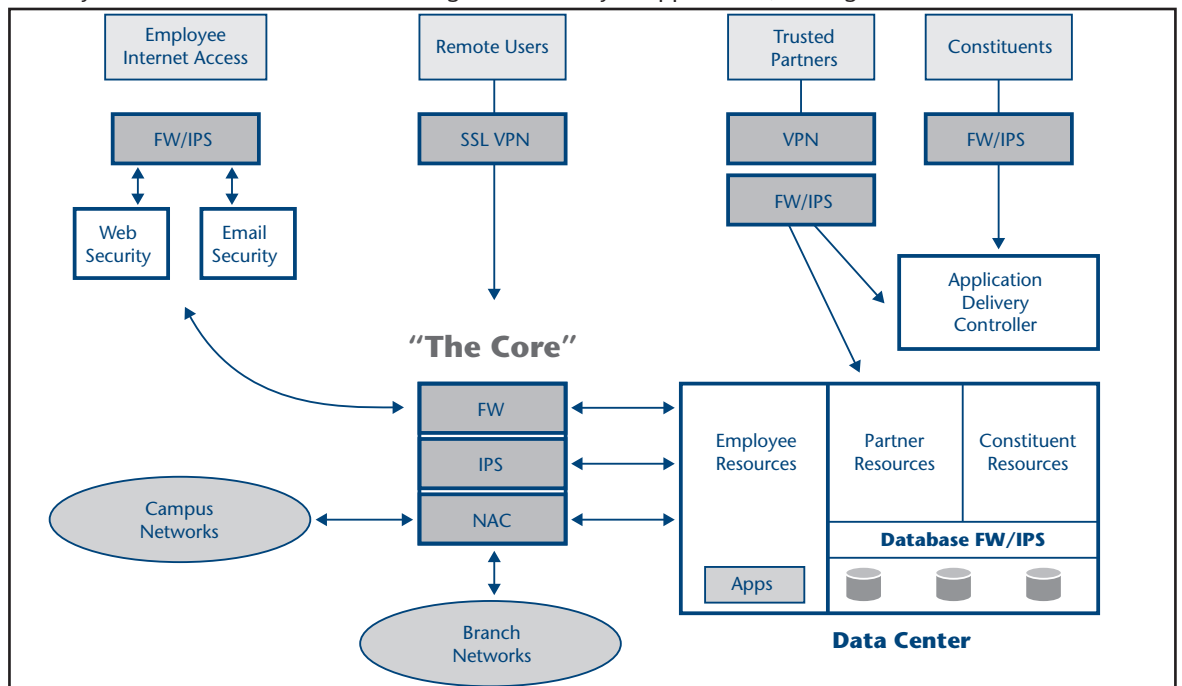
In general, the traffic encountered at perimeter locations is dominated by a relatively small number of Internet protocols and applications like HTTP(s), Simple Mail Transfer Protocol (SMTP), FTP, and Domain Name System (DNS). In contrast, internal networks are home to hundreds and potentially even thousands of protocols and applications, including many that are homegrown. The implication is that to be suitable for internal networks, an intrusion prevention solution should include:

- (a) Protocol and application awareness that is not only broad and deep, but also accurate (in that it does not rely solely on port/protocol designations); and,
- (b) The ability for IT to add its own detection signatures to better account for homegrown and/or highly customized applications.

These capabilities are particularly important given the reliance today’s solutions have on mechanisms such as protocol anomaly detection to thwart the steadily growing menace of zero-day threats.

It is also worth noting that the ability to prevent application-layer attacks is inherently compute-intensive. Digging deeper into packet payloads and reconstructing entire sessions consumes significantly more system resources (memory and CPU cycles) than when decisions are based primarily on information contained in packet headers. This further reinforces the need for systems with high-performance designs and architectures.

Not only does the internal network have a greater diversity of applications, but a greater number of them



are critical to the smooth and continued operation of the government. As a result, reliability of internal



countermeasures is a much greater concern, and not just for devices that will be operated at the network core. Accordingly, an internal IPS solution should include appliance models that incorporate the following optional features:

- (a) Native failover and clustering capabilities, as well as a dedicated HA port;
- (b) Separate control and data planes, so that sensors remain accessible and manageable even under duress;
- (c) A native bypass capability, so that a sensor can pass traffic even if it has a critical hardware failure; and
- (d) Redundant elements, such as power connections and hard drives.

Instead of having a dozen or so classifications to identify external users, internal role definitions can easily number in the hundreds. And instead of having a well-defined DMZ hosting a few tens of systems, the internal network can easily involve hundreds of segments, zones or enclaves, and thousands of devices that need to be protected. When this is added to the aforementioned diversity of protocols and applications, the relative complexity of various policy management, sensor configuration, and event management tasks should be readily apparent. This is why it is imperative for intrusion prevention solutions that will be operating internally to also include capabilities such as:

- (a) Centralized, multi-device management with flexible grouping;
- (b) Delegated and role-based management, so that local administrators can set custom policies and inspections according to local conditions;
- (c) One or more predefined or “recommended” policies for both traffic inspection and event responses; and
- (d) Flexible yet straightforward ways to establish and enforce different policies for different users, enclaves, and resources, such as having VLAN-aware rule sets.

It is also desirable that the intrusion prevention devices be part of a broader family of solutions that incorporate complementary capabilities – especially host-integrity checking and granular access control/containment – to enable better coordination and possibly even integrated policy and event handling.

The growing need to operate at the core of public sector networks and other internal locations certainly adds some new criteria into the mix. However, fundamental functionality and features still deserve close attention as well. In particular, when selecting an IPS, public sector organizations should not lose sight of the need for a leading solution to also incorporate:


- A comprehensive set of threat detection mechanisms that include stateful signature, protocol anomaly, traffic anomaly, and heuristic-based methods, to help reduce false negatives;
- Increasingly rich contextual awareness of the resources that are being protected to help reduce false positives and improve administrator efficiency;
- Robust event correlation, workflow, and reporting capabilities to improve operator efficiency;
- A comprehensive set of automated response mechanisms to support near real-time interdiction of detected threats; and,
- Frequent updates for content, such as signatures, which are generated in a timely manner by a top-notch threat and vulnerability research team.

There are many secondary criteria that matter as well. But this list and the items identified in the preceding sections are a great place to start.

The steady erosion of the factors that underlie the effectiveness of a perimeter-based security strategy is forcing ALL organizations, including those in the public sector, to more directly address the security of their internal networks and systems. In this regard, a strong foundation can be established by combining three highly complementary countermeasures:

- Host-integrity checking, for its ability to keep risky endpoints from accessing the network in the first place;
- Granular access control, for its ability to inherently preclude and contain threats; and,
- Network intrusion prevention, for its ability to scrub allowed sessions of any threats that might otherwise get through.

Intrusion prevention systems, in particular, are well aligned with a key objective for internal networks and, therefore, should not be viewed solely as a perimeter-based security solution. When properly tuned, IPS can



surgically prevent “bad” traffic while simultaneously minimizing the potential of needlessly impeding traffic that is benign, or worse, instrumental to smooth IT operations. However, the suitability of intrusion prevention also depends on the evolution of associated products to account for significant differences between perimeter locations and the internal computing environment. Among other considerations, operating at all internal locations – but especially the core of enterprise networks – requires substantially greater system capacity and performance, coverage for a broader array of protocols and applications, uncompromising reliability, and granular yet highly efficient management capabilities.

Mark Bouchard, CISSP, is the founder of Missing Link Security Services LLC, a consulting firm specializing in information security and risk management strategies. A former META Group analyst, Bouchard has assessed and projected the business and technology trends pertaining to a wide range of networking and information security topics for over 10 years. He is passionate about helping enterprises address their information security challenges. During his career, he has assisted hundreds of organizations worldwide with strategic and tactical initiatives alike, from the development of multi-year strategies and overall architectures to the justification, selection, acquisition, implementation and operation of security and networking solutions.

Juniper Networks Intrusion Detection and Prevention (IDP) products offer the latest capabilities for in-line network IPS functionality that protects the network from a wide range of attacks. Using stateful detection and prevention techniques such as robust protocol anomaly detection and support for over 5500 signatures, Juniper Networks IDP provides zero-day protection against worms, trojans, spyware, keyloggers, and other malware from penetrating the network or spreading from infected users.

Designed to provide the highest reliability, Juniper Networks IDP appliances are equipped with several reliability features. The built-in bypass capability not only offers reliability but a significant cost savings since no external bypass units are needed. The separation of control and data plane ensures continued communication with the appliance, especially when the network is under heavy strain from traffic or attacks. The redundant design and performance scale of up to 10 Gbps real-world throughput ensures the highest level of reliability and performance needed in the network core.

Juniper Networks IDP appliances are designed to interoperate with Juniper Networks SA SSL VPN and UAC Infranet Controllers. This interoperability offers granular access control based on specific users and real-time behaviors, rather than predefined static access roles. The ability to react to threats and compromised credentials in real-time ensures comprehensive security coverage for the entire network.

Juniper Networks IDP products are managed by Juniper Networks Network and Security Manager (NSM), a centralized, rule-based management solution offering granular control over the system’s behavior. NSM also provides access to extensive logging, fully customizable reporting, and management of all Juniper Networks firewall/VPN and IDP appliances from a single user interface.



2000285-001 Aug 2008