

WANTED: The Future of Network Security for Service Providers – Now!

Sponsor: Juniper Networks

Author: Mark Bouchard

AimPoint Group
keeping IT on target

Introduction

The mandate for service providers is simple: transform or face irrelevancy, and eventually insolvency. The commoditization of transport services means that selling connectivity is not sufficient. To be competitive, service providers must differentiate by offering customers a diverse range of value-added and increasingly content-rich services. Moreover, to *remain* competitive, it is essential that transformation be embraced as a continuous process. There is no time to stand still. New and innovative offerings based on new and innovative technologies must be brought to market practically before customers even know they want them.

A significant obstacle to all of this transformation and differentiation, however, is the absence of a sufficiently capable network security solution. The vast majority of available security products are strikingly incapable of achieving the balance necessary to adequately address ongoing trends driving the need for substantially higher throughput, lower latency, and greater stopping power – not to mention other essential criteria, such as adaptability, reliability, unified management and cost effectiveness. The presence of capabilities and strengths in one or more of these areas is inevitably coupled with weaknesses in others.

What service providers need instead is a next-generation network security solution, one that has been architected from the outset to meet all of the applicable requirements simultaneously. Such a solution would allow them to pursue new technologies, service offerings, and ways of doing business without having to make risky compromises or incur extraordinary costs to remain competitive.

This paper explains why a next-generation network security solution is needed *now*, what such a solution entails, and the wide range of benefits that it yields.

The Service Provider Dilemma

There is little doubt that IP-based infrastructure is the way forward for telecommunications service providers. The capital and operational efficiencies of a highly flexible, converged network are simply too compelling to avoid. Add to this the customer demand for new services, ones that make their lives and businesses more enjoyable, more effective, and less costly – such as video on demand, VoIP, IPTV, software as a service, and a broad portfolio of other NGN applications – and one conclusion seems clear. Whether or not a given service provider actually “owns the WAN pipes,” going forward the greatest opportunities for growth and an ever-increasing slice of their revenue will involve IP-based services.

It follows that preserving this all-important revenue stream depends on efficiently providing an effective level of confidentiality, integrity, and availability – in other words, security – for these services. Failing to do so will inevitably lead to theft of services, disclosure of sensitive customer information, and degraded performance, all of which contribute to customer dissatisfaction and, ultimately, defection. But herein lies the dilemma, for it's the very growth, richness, and diversity of these IP services that is creating an increasingly challenging protection scenario and, in turn, eroding the effectiveness of conventional network security products.

The Growing Volume and Sophistication of Network Traffic

One of the major challenges impacting the effectiveness of network security solutions is the growing volume and sophistication of network traffic. On the supply side of the equation, this is due in no small part to the aforementioned pressure on service providers to offer their customers an increasing variety of ways to take advantage of networked solutions. But it's the demand side that's really driving the boat. This is particularly true of enterprise customers. Like the service providers themselves, these organizations also need to take better advantage of information technology, not only to improve their ability to capture and retain customers, but also to reduce costs through increased operational efficiency.

The rising “speed of business” is practically dictating an increased pace of development and deployment of new applications to support both back-office and customer-facing processes. The same reasons are driving greater “exposure” of applications that organizations already have, for example, by making internal systems and solutions externally accessible and/or by extending coverage to additional populations of users and devices. Economic and time-to-market demands are encouraging pursuit of non-traditional approaches for solution attainment, such as software-as-a-service and other forms of managed/cloud-based services. And, perhaps above all else, there is the changing nature of applications.

This too is being driven by the realities of business, but also by the natural evolution of technology. Applications are becoming far more complex, involving far more components and capabilities than was previously the norm. Take, for example, the concept of a Web 2.0 mashup, where numerous underlying connections are made to compute and compile a useful aggregation of multiple, networked data sources and services, each typically representing an application in its own right. Or consider a multi-media collaboration application, such as a telepresence solution. The latter serves to highlight yet another important characteristic, namely a growing sensitivity of many applications to latency – a product of the mounting interest to deliver ever-richer content in conjunction with real-time interactions.

The impact of these various trends is twofold. First, there has been, and in all likelihood will continue to be, a veritable explosion in the amount of traffic that service providers are required to process and protect. The second issue is that whatever processing and protecting is done, it must be accomplished with a minimum of introduced latency – an objective that is becoming even harder to achieve given the growing volume and sophistication of threats.

The Growing Volume and Sophistication of Threats

No longer satisfied with merely building their reputations, hackers are now intent on actually making money. Rather than creating threats that “rattle our cages,” they are now focused on designing exploits that successfully evade or overwhelm the majority of commonly installed countermeasures. This has led to a number of notable changes with regard to the nature of the threat landscape.

For starters, there has been a veritable explosion in terms of the sheer quantity of threats being released in recent years. The speed with which new exploits are developed and launched has also risen dramatically, particularly relative to when associated vulnerabilities are disclosed. Zero-day threats, once described in hypothetical terms, are now an all-too-common reality.

On top of everything else, increased creativity on the part of hackers has led to a threat population characterized by significantly greater diversity and, on average, significantly greater elusiveness. The predominate concerns of the past such as file-level viruses and worms have been over-shadowed by an array of new contenders, including spyware, spear phishing, keylogging trojans, rootkits, and targeted attacks. Even more troubling is the trend of threats “migrating up the stack” to take advantage of much harder to protect application-layer weaknesses.

Once again, the impact these changes have for service providers is essentially twofold. First, in case it wasn’t already clear, the prevailing conditions re-confirm the need to actively provide protection for the network traffic and associated services that represent the lifeblood of their companies. The second issue has to do with the scope of this protection. Multiple types of countermeasures – some involving intensive, in-depth inspection techniques – will need to be applied in order to adequately account for the diversity and elusiveness of today’s threats. This latter item, in particular, has several important implications for what constitutes a suitable network security solution for service providers.

The Impact on Network Security Infrastructure – Key Requirements Going Forward

The challenges outlined in the preceding sections are instructive in that they reveal the minimum set of requirements that now define an appropriate network security solution for service providers. Specifically, to be considered effective, a solution must fully address the following set of essential criteria.

- **Security** – To meet the need for greater stopping power a solution must incorporate multiple countermeasures, featuring a blend of positive- and negative –model techniques that provide protection not only at the network layer, but at the application layer and for individual elements of data as well.
- **Scalability** – System capacity must be readily scalable from relatively modest traffic rates of a few Gbps to an aggregate throughput of greater than 100 Gbps. This is not enough however. Session handling capabilities must also scale to match these throughput levels. For NGN applications, this translates to being able to support several million simultaneous sessions.
- **Latency** – Given the performance characteristics of today’s applications and the far-flung nature of their users (due to overlapping globalization and mobility trends), this criterion can no longer afford to be treated as secondary to having generous amounts of throughput. The two are definitely inter-twined, but latency should also be considered in its own right. Solutions must be architected to minimize the amount they introduce and should also incorporate capabilities to prioritize the processing of designated, time-sensitive traffic streams.
- **Unified management** – Administration of the solution’s various capabilities should not require the use of multiple management tools or consoles. Ideally, it should be possible to configure all of the different inspections required for a given traffic stream within a single rule or policy statement. In addition, the solution should feature extensive role-based administration capabilities to enable granular separation of duties.
- **Reliability** – Highly proven software, redundant components, and support for high-availability configurations are absolutely imperative given the potential impact of failures on service providers and their customers alike.
- **Adaptability** – The solution should be able to accommodate additional functionality – especially new security capabilities – as it becomes available and/or warranted without the need for a major “rip and replace” exercise.
- **Networking/compatibility** – To ensure applicability in the broadest set of use cases, the solution should include at least basic support for a wide range of networking technologies (NAT, address assignment, VLANs, and security zones). Further infrastructure consolidation and simplification can be realized, however, if full-featured instances of certain technologies are fully incorporated, such as routing and switching capabilities.
- **Cost effectiveness** – This is already accounted for in part by each of the preceding criteria, but the general idea is that the solution should be designed to reduce infrastructure complexity and total cost of ownership relative to available alternatives.

Conventional Approaches Come Up Short

The different types of network security products currently in use by most service providers address these requirements to some extent, but they typically have significant limitations as well.

Best-of-breed appliances. Single-service products featuring best-in-class security software continue to receive a lot of attention, particularly for countering new classes of threats or implementing newly emerged security technologies. Even with specialized hardware, however, multiple instances are often needed to achieve sufficient scalability. And this gets multiplied, of course, by the need to have a similar set of devices for each and every countermeasure a service provider decides to deploy. Security and throughput objectives can generally be achieved, but not without substantial cost – not to mention latency, complexity, and rigidity.

Blade systems. Chassis-based systems that accommodate multiple server blades are definitely a step in the right direction but, unfortunately, not a very big one. New functionality or greater throughput is supported by simply adding more blades. As such, these systems definitely deliver a measure of consolidation and reduced complexity, at least relative to best-of-breed appliances. However, they fail to address the need for lower latency. To be fully “processed and protected” packets need to transit the system’s backplane multiple times and are re-processed by each of the different security “modules.” In addition, little, if anything, has been done to unify management. And service providers can still run into throughput constraints due to all of the redundant processing that is required and the fact that most of these systems rely on generic hardware (and operating systems) for the individual blades.

Multi-service gateways. Combining multiple countermeasures with a purpose-built appliance platform has a few interesting advantages. Products with true service integration not only provide a complementary set of security capabilities, but do so with unified management and a fairly efficient processing model that introduces minimal latency. On the downside, throughputs above 1 Gbps are practically unattainable, at least not without adversely impacting latency or cutting back on the variety of inspections that are conducted. Numerous units will be needed in most use cases, leading to commensurate increases in cost and complexity. Having a fixed form factor also limits the adaptability of such products.

The net result is that service providers are caught between a rock and a hard place. They can choose: not to pursue new technologies, services, and ways of doing business because there isn’t really an effective way to economically secure them; to pursue these business-critical items but not bother with securing them at all; or, to navigate various tradeoffs and compromises to pursue and protect them as best they can with the less-than-ideal security solutions that are available. In fact, at any point in time, chances are that most service providers are employing all three of these approaches to some extent.

A Next-Generation Architecture that Delivers

What today’s service providers need is another alternative: a network security solution architected to maximize attainment of the previously identified requirements and that is capable, therefore, of fully enabling the various initiatives service providers must pursue to remain competitive. Based on their respective strengths and weaknesses, a combination of the traditional chassis and multi-service gateway approaches would certainly be a logical foundation from which to build. Indeed, a very attractive option would be a chassis-based design that features:

- Interface flexibility, ideally in the form of modular cards/blades
- Dedicated, distributed hardware and software-based intelligence for ingress processing (e.g., DoS screening, session lookup), internal distribution and balancing of session traffic, and egress processing (e.g., traffic management)
- A high-speed, non-blocking switching fabric that provides any-to-any connectivity between all slots/blades
- Scalable processing capacity, ideally in the form of blades that are dynamically programmable and that automatically inherit the configuration of all other processing cards (thereby avoiding the management overhead and inevitable utilization inefficiencies associated with having each blade configured to perform a specific subset of tasks)
- A dedicated control plane to enable full-time accessibility to management functions
- Redundant hardware components and support for high availability configurations
- A modular operating system capable of being incrementally upgraded (for adaptability purposes) and serving as the “central store” for a robust set of security services that are tightly integrated (to avoid redundant packet processing)
- Fully unified policy and configuration management
- A robust networking feature set, including routing, switching, and virtualization capabilities

A network security solution designed to these specifications is attractive not only due to the tremendous capabilities it provides – see Table 1 – but also because it is based on a highly proven architecture. After all, the architecture described herein is basically the same as that used in today’s large enterprise and carrier-grade routing and switching platforms.

Table 1: Comparison of Different Network Security Product Architectures

	Best-of-Breed Appliances	Traditional Chassis	Multi-Service Gateways	Next-Generation Architecture
Security	☆	☆☆☆	☆☆☆	☆☆☆
Scalability	☆	☆☆	—	☆☆☆
Latency	—	☆	☆☆☆	☆☆☆
Management	—	☆	☆☆☆	☆☆☆
Adaptability	—	☆☆☆	—	☆☆☆
Cost/Complexity	—	☆☆	☆☆	☆☆☆
Networking	☆	☆	☆	☆☆☆

The Benefits of a Next-Generation Network Security Solution

The technical advantages and capabilities that can be attributed to a next-generation network security solution – one that is consistent with the previously described architecture – have already been fairly well illuminated. They include substantially greater scalability and security coverage with considerably less latency and infrastructure complexity. However, there are numerous business-oriented benefits that should also be acknowledged. For service providers, a next-generation network security solution:

- **Enhances competitiveness** – rapid deployment of new and innovative value-added services is enabled by removal of the obstacle of always having to first acquire and deploy more security infrastructure to ensure service integrity.
- **Improves responsiveness** – the ability to rapidly provision services and offer fully customizable security capabilities helps deliver on the need of enterprise customer's to execute changes quickly to keep up with the ever-increasing "speed of business."
- **Reduces total cost of ownership** – network and security management are greatly simplified; the same, highly economical solution can be used to protect virtually all of a provider's core infrastructure, including both external and internal –facing data centers; and, cloud-based/managed service offerings can efficiently be secured by accommodating tens-to-hundreds of customers per platform and by enabling an optimum balance to be attained between I/O interface and throughput requirements.
- **Facilitates future growth** – a high degree of adaptability ensures applicability even as the threat, application, technology, and business landscapes continue to change, along with the portfolio of services offered in response to new trends.

The bottom line is that a next-generation network security solution enables service providers to accelerate the transformation, innovation, and differentiation required to drive continued growth while still containing costs. Service offerings and the infrastructure that supports them can easily be scaled without the usual delays and capital expenditures required for new hardware installations.

Summary

Today's service providers have reached a critical juncture. Their ability to continuously offer new and innovative IP-based services and support the burgeoning population of NGN applications – an absolute necessity when it comes to remaining competitive – is being impeded by the lack of a suitable product for adequately securing these offerings. What they require is a next-generation network security solution, one that takes a page from the book used to design carrier-grade routing and switching platforms. By employing a similar architecture and taking advantage of similar design principles, such a solution is capable of delivering superior protection with minimal added latency and greater scalability, adaptability, and cost-effectiveness than would otherwise be possible. The net result: service providers can pursue new technologies, service offerings, and ways of doing business without having to make risky compromises or incur extraordinary costs.

About the Author

Mark Bouchard, CISSP, is the founder of AimPoint Group, an IT research and advisory services company specializing in information security, compliance management, application delivery, and infrastructure optimization strategies. A former META Group analyst, Mark has assessed and projected the business and technology trends pertaining to a wide range of information security and networking topics for more than 12 years. During this time, he has assisted hundreds of organizations worldwide with strategic and tactical initiatives alike, from the development of multi-year strategies and high-level architectures to the justification, selection, and deployment of their security and networking solutions. A veteran of the U.S. Navy, Mark is passionate about helping enterprises address their IT challenges.

A Word From the Sponsor

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.