IDC
Analyze the Future

WHITE PAPER

Unique Security Challenges in the Datacenter Demand Innovative Solutions

Sponsored by: Juniper Networks

Christian A. Christiansen John Grady
Phil Hochmuth
October 2013

IDC OPINION

The enterprise network has become increasingly complex as mobile workers, cloud applications, ecommerce, and virtualization have become more prevalent. New security solutions have emerged to address the increased attack surface created by emerging technologies. However, there remains a lack of clarity around where these security solutions are best deployed. IDC sees three unique use cases for securing the enterprise network:

- **Campus edge:** At the campus edge, internal users access external applications. Organizations require granular control over applications based on each user, as well as the ability to detect malicious content. Next-generation firewalls (NGFWs) and unified threat management (UTM) appliances provide this functionality on a consolidated platform.

- **Datacenter edge:** At the datacenter edge, external users access internal applications. Users are both numerous and unknown, making threats harder to detect with a high degree of certainty. Additionally, with many applications hosted in this part of the network having direct revenue implications, availability is a key concern. Solutions that mitigate distributed denial-of-service (DDoS) attacks and block hackers from Web applications with specificity are necessary.

- **Datacenter core:** Internal resources and data are hosted in the core of the datacenter. Typically, at least a portion of the infrastructure in this part of the network is virtualized. Latency is not acceptable in the core datacenter, meaning that performance is of key importance. Solutions that provide consistent policy enforcement across both physical and virtual environments, high performance, and the scalability to protect dynamic environments are required.

While each of these use cases has distinct features requiring unique solutions, integration remains important. The ability to leverage attacker intelligence across the infrastructure can improve security and simplify enforcement. Juniper Networks offers a portfolio of products built to address each of the previously mentioned scenarios specifically while also functioning as a holistic solution.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

SITUATION OVERVIEW

The Unique Security Needs of the Datacenter and Enterprise Edge

The rise of mobile devices, cloud applications, social networks, and Internet use in general has completely changed the way people live and work. While this has been a mostly positive occurrence, these enabling technologies have proven to be an efficient avenue for bad actors to compromise networks and spread malicious code. At the same time, organizations face a growing diversity of bad actors and attack methods designed to steal information and disrupt business. In response, security solutions are being developed to provide greater protection against attacks that exploit Web 2.0 and other user-centric technologies. Application-aware next-generation firewalls and unified threat management products are the best examples of new popular solutions built with these previously mentioned threats in mind. However, while these products have garnered the lion's share of attention in the network security industry over the past few years, they are designed to secure only a specific portion of the network — the campus edge. Focusing too much attention on this approach leaves a gap in many networks when considering the threats, security requirements, and requisite performance requirements for modern datacenters (both the perimeter and the core), which remain significantly different from elsewhere on the network (see Figure 1).

FIGURE 1

Datacenter Versus Campus Edge Security: Different Attack Surfaces and Defenses

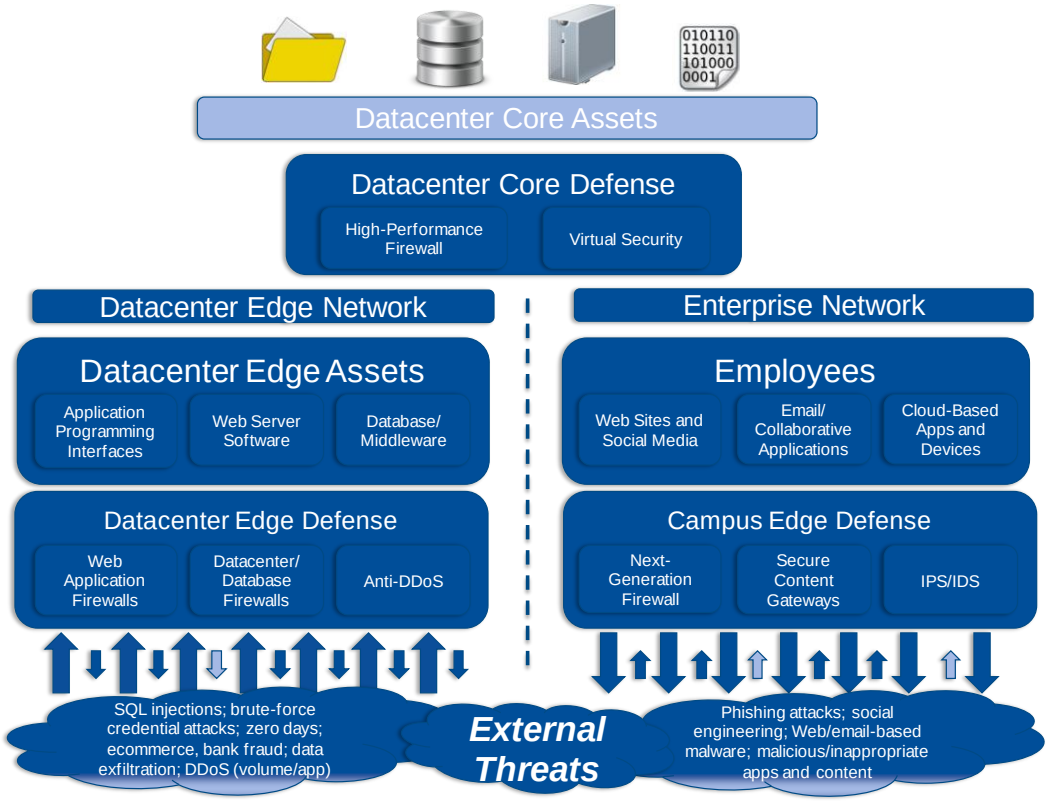

Source: IDC, 2013

Enterprise Edge

Securing the enterprise edge in the datacenter requires a different set of solutions than campus environments. Given that the majority of attacks targeting the datacenter and the externally facing apps that reside within them are "outside in" in nature, effective datacenter security solutions must successfully detect attacks earlier in the attack cycle, profile the actual attackers, and correlate a variety of information to prevent future attacks. When combined with the effective security technologies at the campus edge, this approach can provide significantly better protection for organizations.

Securing the Campus Edge: Addressing the Inside-Out Problem with NGFW/UTM

Campus security requirements that center on securing users and devices are well defined and addressed by most firewall vendors. Application visibility and control, IPS, and Active Directory (AD) integration capabilities are provided in almost every NGFW solution available today.

These multifunction security appliances first gained widespread adoption as the answer to the cost, deployment, and management complexities inherent in combining multiple point solutions in campus networks. Yet, over the past three years as Web application use (including social media) has proliferated in the workplace, the focus of many of these appliances has shifted toward providing more granular user controls via next-generation firewalls. In contrast to a traditional 5-tuple firewall, NGFWs deliver additional capabilities, including integrated IPS, application visibility and control, and user-based policy control.

The primary use case for deploying an NGFW at the campus edge is to gain visibility into and control of application traffic generated by employees, contractors, and other users on the internal network. In turn, the NGFW can protect users from infections caused by accessing malicious Web sites, downloading malware, and performing other actions that put organizations at risk.

Further, the ability to authenticate users based on Active Directory integration and subsequently enforce application-level policies based on their identity is a key requirement for campus networks. The near-ubiquitous use of social media and other data-intensive applications coupled with the increasing security threats targeting organizations through these platforms creates an urgent need to granularly manage access. Organizations must be able to give (some or all) users access to the productivity tools needed to do their jobs, such as Facebook, Skype, and LinkedIn, while limiting risky or otherwise unwanted applications such as peer-to-peer file sharing services. The application ID capability of NGFWs addresses this issue by identifying safe applications and enforcing application-level policies. As an added benefit, many of these solutions include integrated IPS to inspect return traffic for client-side exploits, antimalware, and URL filtering to prevent against device infections as well as user access management capabilities to aid with corporate regulations and compliance and so forth.

Why NGFW Cannot Secure the Datacenter Edge

The NGFW architecture that solves the modern campus problem as previously mentioned conversely prohibits the NGFW from acting effectively as a security device at the datacenter edge where, paradoxically, the security problems are reversed. Protecting datacenters requires focus on securing applications and data versus users because the threats in this environment are not aimed at user devices or client-side exploits. Further, with millions of outside users accessing Web servers, high performance is a critical requirement for datacenters. This requirement prevents the use of integrated IPS and UTM solutions that make up an NGFW in many circumstances because often integrated IPS and layered security services degrade the performance of the firewall.

In addition, NGFWs are not well suited to the complex and large-scale DDoS attacks companies now face. While integrated IPS in NGFWs will detect and mitigate small-scale DoS attacks aimed at end-user devices, bandwidth exhaustion techniques and newer sophisticated techniques for server-side exploits, such as those described in the later sections of this document, cannot be detected or mitigated effectively. Application ID, which serves to identify HTTP/HTTPS-nested traffic on campus networks and is used to apply application-level policies, does not provide the deep inspection of application content necessary to detect abnormal behavior by unknown users accessing public-facing resources. Finally, the third architectural component of an NGFW, AD integration, is rendered useless in the datacenter where the users are unknown outside the network, which means they do not authenticate to an Active Directory.

Ultimately, the design and central use case for application-aware NGFWs (controlling the actions of known internal users and blocking access to external applications that live on the Internet) is best suited for scanning traffic at the campus network egress. In the datacenter where applications and valuable assets reside, the problem is exactly the opposite, with unknown external users accessing internal applications and creating an ingress (or outside-in) problem. In this scenario, an entirely different set of solutions is required.

Securing the Datacenter Edge: Addressing the Outside-In Problem with Web Application Security and DDoS Mitigation

Rather than controlling user actions, the main goal of datacenter security is protecting valuable internal resources and information from outside intrusion. The difficulty in protecting datacenters lies in recognizing the handful of bad actors among the thousands or even millions of unknown legitimate users accessing data — all while managing the demands of dynamic datacenters that require security solutions to provide scalability, dynamic configuration, and high performance.

To be considered an effective defense, a security solution must address five key requirements in the datacenter:

- **Detect and mitigate DDoS and application DDoS attacks (volumetric and emerging "low and slow" attacks).** In part because of the ease of access attackers have to botnets capable of sending significant amount of traffic, large and sustained DDoS attacks of over 1Gbps have become commonplace. Layers 3 and 4 ICMP, SYN, and UDP flood attacks are straightforward enough to launch. At the same time, organizations also now face more sophisticated and stealthy

DDoS attacks that target Layer 7 applications and often resemble legitimate traffic, making them more difficult to detect and prevent.

To defend against these attacks on the datacenter, dedicated DDoS solutions are necessary for a number of reasons. First, IPS capabilities in NGFW devices are not well suited to defend against large-scale DDoS offensives. Each new request that comes through the firewall or IPS is added to its state table and very quickly can result in the connection limit being exhausted and becoming a choke point on the network. Further, firewalls are not designed to detect legitimate traffic across multiple sessions working maliciously, making it difficult to detect "low and slow" application-level DDoS attacks. Finally, broad measures to mitigate DDoS such as blocking entire IP address groups or throttling traffic may stop the attack, but at the cost of affecting legitimate users.

Dedicated anti-DDoS solutions that are able to correlate information across sessions and granularly block malicious traffic without affecting legitimate users are the only way to adequately address the issue of DDoS attacks.

- **Deliver Web application security against unknown and zero-day attacks.** When attacking Web applications, adversaries use a mix of automated tools and human analysis to find easy routes to access sensitive data in datacenters hosting Web applications. Most of these efforts start with the massive machine-driven process of URL and Web application scanning to map the respective technical underpinnings of a potential target: server and Web server OS type, databases and middleware versions, and access control schemes. This effort filters potential targets down to smaller groups of more vulnerable systems. Individual attackers then begin the handiwork of inspecting and analyzing the Web infrastructure for potential exploits and weighing the amount of effort required versus the potential monetary or other value a successful break-in or breach would provide.

  With the Web development community constantly introducing new and untested code packages, tools, and languages that often contain vulnerabilities, combined with knowledgeable attackers actively searching for new targets, organizations with a strong Web presence face a significant challenge in securing resources. Web application security technology must provide protection in a dynamic and accelerating threat landscape.

  Yet traditional Web application firewalls (WAFs) have several weaknesses that make it hard for them to quickly defend against the constantly evolving attack methods used by criminals. Combinations of attacks, sometimes altered slightly, can evade the known-signature databases of WAF and application security defenses. The reliance of most WAFs on signatures to protect against known attacks renders them largely ineffective against unknown and zero-day attacks where no signatures have been written. Further, because rules and signatures must be written every time a new Web application threat is detected, there is also the associated cost of managing this time-consuming process. Solutions that leverage additional intelligence beyond signatures are required to protect applications against unknown vulnerabilities and custom-targeted attacks that are becoming more common.

- **Identify bad actors with a high degree of precision to make false positives statistically insignificant.** For many organizations, false positives are nearly as unwelcome as undetected attacks. This is especially true in the case of customer-facing applications, whose availability may directly impact revenue and brand equity. Because of the high potential of false positives, many organizations elect to deploy security solutions in detection-only mode to ensure uptime at all costs. For example, DDoS attacks are typically indicated by a sudden ramp in activity that overwhelms internal resources. However, because traffic spikes can happen organically based on a variety of factors such as sales or holidays, an online retailer may not set its security appliance to block this threat.

  For a solution to be deployed in block mode (which enables a much stronger security posture), it must be able to identify attackers by scanning for actions that can only be construed as malicious. This removes any doubt that the user could be legitimate. Instead of relying on signatures based on keywords and character strings often used in attacks, solutions should look directly at the action of a user, which largely removes the possibility of a false positive due to confusing users' intention.

- **Uniquely identify bad actors by going beyond IP addresses.** Identifying bad actors within a particular session is only the first step and should be supplemented by definitive ways to detect when they come back. Yet using IP addresses as the primary indicator typically does not work because attackers leverage proxies and other tools to mask their location to conduct subsequent attacks. Further, blocking IP addresses can lead to significant false positives that in some cases could mean blocking an entire country of users masked behind a single IP. Combining multiple factors beyond the IP address as a basis to identify attackers can create more meaningful and actionable data points. This approach means creating an attacker profile based on operating system, machine hardware specifications, and other components unique to an attacker's system that creates a fingerprint, which can help organizations definitively respond to their actions. Once attackers are specifically identified, they can be blocked with no impact to legitimate users. Additionally, moving beyond IP addresses allows hackers to be accurately detected in subsequent attempts to exploit the same or new Web applications.

  Attacker fingerprints can subsequently be shared across a common customer set or community and used to block or deny access of the identified bad actor across environments. In practice, this could create a "herd immunity" model similar to antimalware signature sharing and dissemination.

- **Propagate the enforcement policies to the network perimeter in real time so that the firewall can block an identified attacker from entering the network in the first place.** If an attacker is able to be accurately identified across multiple exploit sessions, the final question becomes how best to fully block the attacker from the network entirely. This is best accomplished at the perimeter through the firewall. Detailed information on known bad actors is largely useless unless it's actionable. In the case of shutting off attacks on Web apps or datacenter infrastructure, discovering an interaction with an identified hacker after the fact via a log file or SIEM analysis is too little too late. Fingerprint or signature information on identified attackers must reside at the first line of network defense — the datacenter firewall

infrastructure or enterprise perimeter. At this level, security infrastructure can automatically react when it recognizes a known malicious connection. Similar to how antivirus signatures became more prevalent in enterprise gateway technology in the mid- to late 2000s, automated attacker recognition will likely become a more standard technology in firewalls going forward.

Internal Datacenter

Securing the Core: Addressing Performance, Scalability, and Virtualization

While much of the focus of this document has been on securing the edge, another important aspect that should not be overlooked is protecting the internal datacenter. In the core datacenter, resources are internal (as in the datacenter edge scenario), and the users are internal (as with the campus edge use case). However, despite dealing with two known metrics (users and resources), distinct issues still need to be addressed at the core that differ from the issues at the campus edge and the external datacenter edge. The most important of these are performance and virtualization.

As datacenter fabrics have evolved and workloads and application traffic have increased, performance requirements have skyrocketed. Increasingly, 10Gbps and even 40Gbps connections are common as users demand instant access to applications and data. Viable core datacenter security solutions must support high levels of throughput (well over 10Gbps) and low latency, even when multiple security services are turned on. Additionally, these solutions must have the ability to scale seamlessly to support the needs of dynamic datacenter environments.

Perhaps even more importantly, security solutions must be able to secure virtual environments. Nearly 70% of enterprises are now deploying some form of virtualization or private cloud (according to IDC's *2012 Cloud Security Survey*). This means that in a majority of environments, workloads can move freely among physical servers, security zones, and physical locations. In virtualized environments, high volumes of "east west" traffic (i.e., server-to-server communications) can create complex security and compliance scenarios because data and workloads can traverse boundaries that were traditionally set via physical network topology or locations of machines. In other words, traffic flowing between virtual machines (VMs) never leaves the physical host and therefore is never seen by the physical firewall. Organizations require access control and visibility into inter-VM traffic both for compliance reasons and to protect against security violations. This means security policy and enforcement controls must now speak the language of virtualization and enforce policy based on logical network and application stack topologies. For securing east-west traffic, placement of the firewall in the virtual environment hypervisor itself has the dual benefit of visibility and granular access control over traffic flowing through VMs while maintaining high levels of performance.

As pervasive as virtualization is today, many environments still remain a mix of physical and virtual resources. So it is not enough to simply deploy a solution for a virtual environment. There must be integration between physical and virtual security solutions to ensure that resources are properly segmented and that policy is enforced consistently and continuously from the datacenter perimeter to the internal core, regardless of whether the environment is physical or virtual. Further, policy creation and management need to be centralized based on the resource, regardless of whether it is physical or virtual.

JUNIPER'S APPROACH TO DATACENTER SECURITY

To address the previously mentioned issues, Juniper has built a portfolio of solutions that moves well beyond the next-generation firewall. To ensure security and quality of service, security solutions must marry high performance with the ability to block advanced attacks with specificity while maintaining availability for legitimate users. Juniper delivers on these requirements through its SRX Series Services Gateways, WebApp Secure, DDoS Secure, and Spotlight Secure solutions (see Figure 2).

FIGURE 2

Securing the Datacenter: Utilizing Network Intelligence to Identify and Stop Attacks

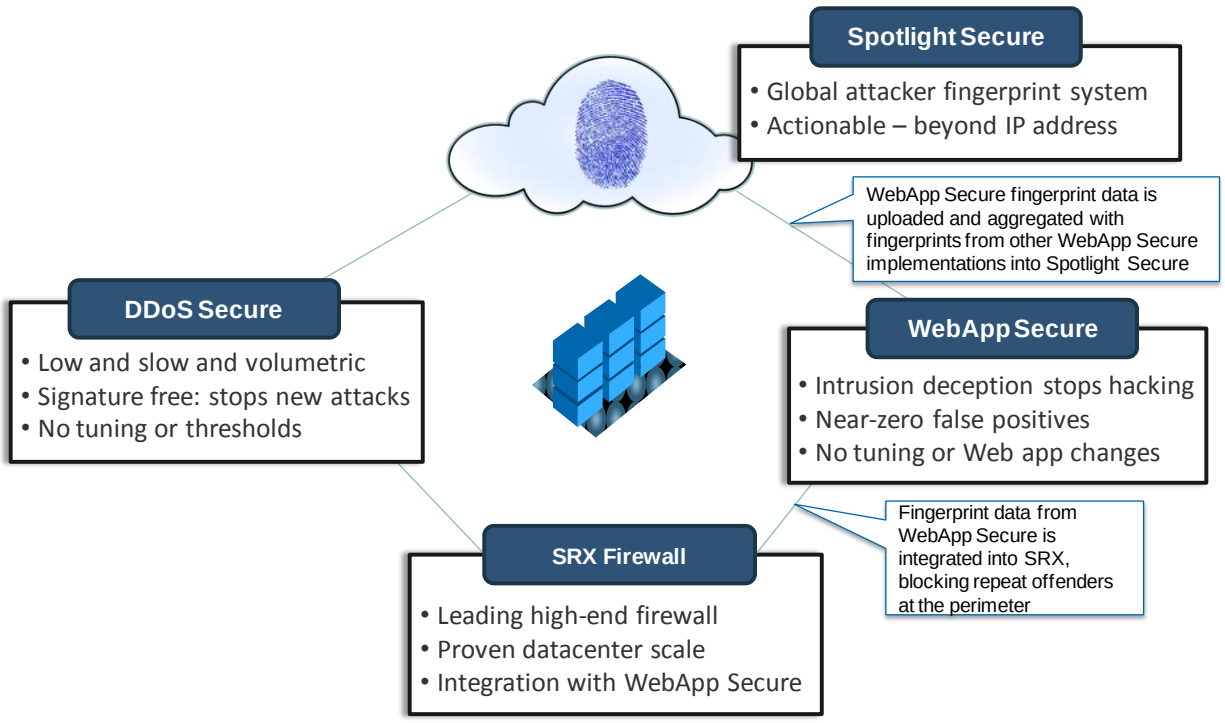

Source: Juniper Networks, 2013

SRX Series Services Gateways

The SRX Series of appliances is Juniper's flagship product in the datacenter security space. In total, 12 models serve the needs of small branch offices up to the largest datacenter. All run the Junos operating system and are capable of supporting a variety of security features. Five models are purpose built for the datacenter: SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800. These appliances integrate physical and virtual security through Firefly Host and feature separate data and control planes, which help ensure availability under heavy traffic loads. Additionally, network address translation (NAT), IPv6, high availability, Layer 7 security, prioritization, and quality of service are supported.

The SRX1400, SRX3400, and SRX3600 share a common modular architecture. Each appliance is powered by services processing cards (SPCs), network processing cards (NPCs), and I/O cards (IOCs). The SPC powers the security functionality on the appliance. Security services include firewall, IPS, AppSecure, and IPSec VPN. Compute resources from all SPCs are allocated across all security functions and not tied to a specific piece of hardware, providing for near-linear scalability. This resource management is powered by the NPC, which directs traffic to the various SPCs to leverage processing cores in the most efficient manner possible. The IOC provides for greater port density above and beyond what is standard on the appliance. Maximum throughput scales from 10Gbps on the SRX1400 to 55Gbps on the SRX3600, while concurrent sessions rise from 1.5 million on the SRX1400 to 6 million on the SRX3600.

The SRX5600 and SRX5800 are modular, chassis-based appliances for very large datacenters. Similar in architecture to the 3000 series, the 5000 series also leverages SPCs and IOCs, but without the need for NPCs. The five-slot SRX5600 is capable of 100Gbps throughput through next-generation SPCs (NG-SPCs) released early in 2013, while the SRX5800 supports throughput up to 200Gbps. Both appliances can support up to 6 million concurrent sessions with the NG-SPCs. To maximize uptime and minimize interruption, in-service software and hardware upgrades are supported on the platform, further supporting the scalability requirements inherent in datacenter environments.

WebApp Secure

WebApp Secure is an intrusion deception solution designed to defend Web sites and Web applications from sophisticated attacks. Rather than rely on attack signatures that can result in false positives, WebApp Secure injects fake code into the Web site or applications, creating honeypot-like traps that mirror common preattack actions such as probing for hidden directories, links to confidential resources, or authentication vulnerabilities. As attackers work through their toolkit, they fall into one or more of these traps, alerting administrators to their presence. Because the traps only simulate the vulnerabilities attackers look for without tying into the application code itself, there is no threat of infiltration. Further, because the detection points are designed around actions that are unquestionably malicious, false positives are limited.

Once an attacker has been identified, the attacker then must be blocked. Since IP addresses can change and in many cases represent multiple users, relying solely on IP as an identification factor is an ineffective way to block bad actors. WebApp Secure solves this by placing an encrypted persistent token on the machine of attackers using Web browsers to launch attacks. When attackers use software and scripts rather than browsers, the offending machine is fingerprinted. In both examples, once an attacker has been identified, a smart profile is created with the IP address, fingerprint or token information, and types of attacks attempted. Various responses can be initiated based on the severity of the threat. Low-threat attackers may simply be warned with a message that they are being watched. For more advanced attackers, the site can be displayed as unavailable while remaining unchanged for legitimate users. To gain as much information as possible, the attackers can be given responses that make it seem as if the probing is successful. This not only serves to keep attackers busy on a fruitless endeavor but also helps WebApp Secure collect as much information as possible by keeping the attackers engaged.

The solution is deployed as a reverse proxy with load balancing as either a hardware-based appliance or a virtual appliance. Clustering is supported for high availability and increased throughput. Additionally, WebApp Secure is PCI compliant.

DDoS Secure

Designed to combat denial-of-service attacks, Juniper's DDoS Secure leverages signatureless technology to detect and mitigate "low and slow" application-level attacks. Rather than leveraging thresholds and tuning to mitigate attacks, the solution relies primarily on an algorithm-based risk scoring system dubbed "CHARM." CHARM provides a real-time risk score for each IP connecting to the datacenter. Incoming traffic is analyzed and differentiated between legitimate human traffic (which is typically irregular in nature) and machine- or botnet-generated traffic (which is typically more regular and logical). A higher CHARM score indicates lower risk. In addition to inbound risk scoring via CHARM, the solution analyzes resource responses to the external request. Resource health is measured at Layers 3, 4, and 7 for URL response time, HTTP server codes, and backlog queue, among other metrics. By correlating inbound risk and outbound response, DDoS Secure is able to detect stealthy attacks that typically evade signature-based defense.

DDoS Secure can be deployed as a 1-RU physical appliance or virtualized machine. Because the solution correlates inbound and outbound information for risk assessment as opposed to using predetermined threshold blocking, false positives are limited, which makes the solution easy to manage and deploy. Additionally, BGP integration enables the solution to work with cloud-based mitigation solutions to combat large-scale volumetric attacks.

Spotlight Secure

Spotlight Secure aggregates the attacker identification information gathered by WebApp Secure and feeds intelligence to supported SRX, DDoS Secure, and WebApp Secure gateways. Customers subscribed to the service benefit from additional intelligence collected from Juniper's installed base. Compared with an IP reputation feed, Spotlight Secure provides specific attacker identifications and their threat profile, enabling protection regardless of IP spoofing. Known attackers can be blocked before they're able to launch an offensive attack. The correlation across multiple technologies allows customers to have a better view of their threat landscape and make the connection of seemingly piecemeal attacks.

Integration of SRX Gateways, WebApp Secure, and Spotlight Secure

The SRX Series integration with WebApp Secure is straightforward, yet effective. The SRX acts as the enforcement point at the perimeter that stops attacks before they can begin. The process works as follows: An attacker is detected at WebApp Secure using intrusion deception. A device-level tag is pushed to SRX in real time so that any future connections from that attacker can be enforced at the perimeter. This gives the benefit of enforcing further out at the perimeter while still maintaining device-level accuracy.

CHALLENGES AND OPPORTUNITIES

While Juniper's approach to securing Internet-facing applications and resources in the datacenter goes well beyond the next-generation firewall model that many other vendors currently espouse, there remain challenges that Juniper must overcome. From a technology perspective, while WebApp Secure is an innovative product, improvements must be considered to gain widespread adoption. Specifically, an improved approach should be considered to combat attackers using virtual machines. Part of the value proposition of WebApp Secure and Spotlight Secure is the fingerprinting capability, which allows returning attackers to be quickly identified and information to be shared across all customers of the service in part via a persistent token. However, attackers using a virtual instance to analyze a Web site may spin up a new virtual machine on a regular basis, removing the persistent token as a point of identification. While the solution is designed to detect and block VM-based attacks using other components of the fingerprint, the loss of the persistent token in attacks from virtual machines is notable.

From a go-to-market perspective, Juniper's brand and products are well known and respected in the network and network security areas and are growing in visibility and reputation in the enterprise datacenter. However, the buying center for Web application firewalls and Web application security tools often resides with the Web application development and deployment teams; the specialists fine-tuning the WAFs are sometimes separate from the network security teams doing the same activities on network firewalls. This could be a new sales cycle experience for many Juniper account representatives and partners. Web application DevOps teams often have very different needs and benchmarking requirements for performance when evaluating technology compared with network infrastructure or IT security. The core background of these teams is outside the familiar network/security operational roles with which Juniper is most familiar.

Finally, Juniper will have to quickly establish partnerships to help build momentum in these new markets. Having exited the load-balancing/application delivery market years ago, Juniper will likely be selling into existing installed bases of application controller solutions, the providers of which may also offer WAF and Web application protection capabilities on top of Layer 4/7 load balancing and acceleration. Partnering with providers that do not offer security solutions may help Juniper quickly capitalize in accounts where another provider is delivering ADCs, but not security. In addition to these connections, vendors that offer database security, as well as Web infrastructure stack providers (Oracle, IBM, Microsoft, Red Hat, etc.), will be key technology and integration opportunities for Juniper as it gets deeper into the DevOps realm.

CONCLUSION

The volume and sophistication of threats targeting both critical enterprise infrastructure and individual employees continue to rise on a seemingly asymptotic curve. The goals of both types of attacks are similar, if not exactly the same: steal the target organization's most sensitive, valuable data and digital assets, whether by compromising an individual employee or by attacking the Web and datacenter infrastructure directly to gain access to valuable digital information. It is critical that

enterprises understand the fundamental differences in attack methods and security tools associated with each scenario.

While securing end users' Web activities and systems access is important, IT security professionals should also take a hard look at the security technologies directly fronting the corporate crown jewels — applications, servers, and data that reside in the enterprise datacenter. To that end, the following approaches are recommended to security and risk professionals:

- Instead of a "shields up/heads down" approach to security, enterprises must gain more visibility into attackers' methods, and even identities, or attackers' identifiable characteristics. Blocking broad swaths of IP addresses, based on reputation lists or as a result of specific incidents, is ineffective from a security standpoint and inefficient from a business IT operations perspective.

- Invest strategically in datacenter security technologies built to address future threat trends and attack methods. Relying on mature attack detection methods, such as gateways, based on attack signature and virus definition, and firewall technologies, won't be enough to keep up with sophisticated, evolving threats on Web apps and server infrastructure.

- While addressing evolving threats at the datacenter edge, enterprises must also consider how internal datacenter network architectures have changed and rethink security approaches with virtualization — both server and network — in mind. The inability to deploy and manage security policy across the virtualized datacenter can lead to compliance and regulatory scrutiny as well as audit failures — or worse, make critical virtual workloads and data more vulnerable to external and internal attacks.

Addressing both external and internal threats to the datacenter requires a diverse set of security technologies that can correlate and react to threats across various defense layers while being managed and configured by a common policy control framework. Juniper's inside-and-out approach to datacenter security can enable enterprises to build out applications and services that drive productivity improvements and revenue in the datacenter while minimizing exposure to and risk of attacks and breaches.