

Network Security Perimeters

By Chris Ellis

www.EllisTalks.com

Sponsored By:



Produced By:



Table of Contents

Network Security Perimeters: How much is enough?3

Definition of an NSP3

NSP architecture6

Type of threats you are vulnerable to7

Market trends8

Beyond basic NSPs: Scalability & performance Issues9

Building an NSP for performance9

Enhancing NSP performance10

Conclusions11



Chris Ellis is an IP networking specialist who has spent most of his career as a consultant analyzing, designing, and deploying IP networks. His career of over twenty one years has seen a particular focus on the engineering of secure IP networks as well as next generation networks that offer quality of service, high performance and high availability.

Chris has developed content for, and delivered numerous courses and seminars for customers including Spirent communications and Alcatel.

Chris created EllisTalks by blending his technical knowledge with his highly acclaimed presentation skills. The EllisTalks brand of IT learning products is targeted at busy people and his innovative video streaming tutorials continue to break new grounds.

Information on Chris Ellis is available at www.EllisTalks.com and Chris' email address is chris@EllisTalks.com

Network Security Perimeters: How much is enough?

Network Security Perimeters (**NSPs**) have become necessary as a result of our increasing dependency on electronic communications via the Internet, characteristically for access to e-commerce services. Prior to the explosive growth of e-commerce, **firewalls** were sufficient for the most part. It was possible to configure a single firewall to prevent unauthorized traffic coming into your company because the rule was simple: No traffic is allowed into your corporate site, unless that traffic originated from within. Then the firewall could match the outbound request with the inbound response, granting it permission to enter the network. This is known as a stateful inspection firewall.

But that traffic model has changed. A single firewall no longer can protect today's corporate environment. An electronic presence is imperative in business today, and we now *encourage* external traffic to come into our networks. They come to purchase our products, to access our services; and to see our marketing messages. Today's network model is akin to Wal-Mart's: Welcome guests with a smile and encourage them to linger while enjoying the shopping experience. Unfortunately for the wired world, there are vogue aberrations — those that steal into an online site and try to "take the place down."

So the issues facing every organization with an online presence are: What is an appropriate NSP configuration given your corporation's security policy? Can your network presence — your Web infrastructure — meet the expected increase in traffic volumes that your organization is predicting? Can it do so while maintaining appropriate security levels, and still meet adequate performance levels?

Those are the challenges facing today's IT executive. In this special report the issues of NSP design, performance and scalability are addressed. In the next report the issues of NSP manageability and ease of use will be examined.

Definition of an NSP

- **NSP** usage
- NSP components

A Network Security Perimeter is analogous to the type of perimeter security you might see around an army compound: barbed wire, high fences, motion detectors and armed guards at gates ensure that no unauthorized visitors can gain entry.

In the wired world, **firewalls** replace the barbed-wire fences, Network Intrusion Detection Systems (NIDS) replace the motion detectors and IP Security (**IPSec**) style authentication mechanisms replace the interrogation at the gate. While it isn't a perfect analogy, it's best to think of an NSP this way: as a *collection* of devices whose group objective is to detect, alert and possibly eliminate traffic with malformed intentions.

NSP usage

It would be correct to assume that there is not a one-size-fits-all approach to building out an NSP. If you are a home user, your NSP might consist of an off-the-shelf security appliance. That market segment has low cost, minimal configuration firewalls that do a pretty good job of deflecting threats, right out of the box.


At the other extreme sits the large enterprise environment - one that has made a significant investment in Web technologies - many of which today are mission-critical. They probably support a variety of sales and marketing initiatives, or are in the business of disseminating information. These Web sites by their very nature invite your customers to come in and learn more about your company, your competitive differentiators and your value proposition; and to browse products, facilitate online purchases and generally have easy access to information.

Because of the open-door nature of these Web sites, a robust NSP is imperative. If you've outsourced your online activities to an application service provider or a hosting organization, then you want to know how their NSP has been deployed. Even so, you still might host Internet services for your employees, so your site is still vulnerable to an attack. In other words, you must have protection if you are connected, and the bottom line in building out an NSP is much more than just connecting a single firewall.

NSP components

Before we examine **NSP** performance and scalability issues, let's look at the type of devices that you will find within a large corporate NSP: routers, **firewalls**, **network intrusion-detection systems (NIDS)**, **security information management systems (SIMS)** and some **IPSec** VPNs.

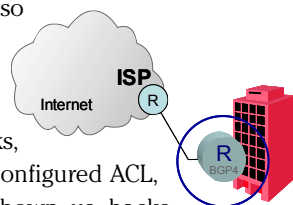
Terms in **green** can be seen in the Talking Glossary: www.EllisTalks.com/talking



Routers

The first router that your **ISP** connects to, is your border router (or access router) because it sits at the border of your security perimeter and provides access to the Internet. This router runs a protocol called Border Gateway Protocol version 4 (**BGP4**). BGP4 allows a border router to learn how to reach a destination address through the Internet. BGP4 routers face your ISP on the one side, and your corporate Network Security Perimeter on the other.

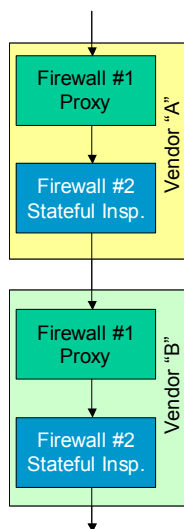
All routers use simple Access Control Lists (ACL), pronounced *ACK-ulls*, to determine which traffic - which packets have permission to enter your network. ACLs also control traffic flowing out of your network. ACLs work by examining source and destination IP address in each packet. In the early networks, this border router, with a properly configured ACL, was your firewall. As history has shown us, hacks quickly figured out how to beat this system using a technique known as address spoofing. Routers now are only part of the NSP solution and state-of-the-art firewalls are much more robust at guarding a network perimeter.



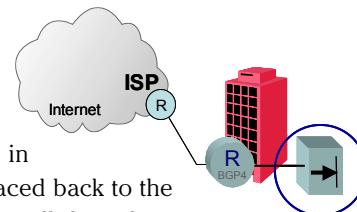
Firewalls

Behind the access routers are the **firewalls** - note the use of the plural. A single firewall is no longer considered sufficient. Multiple firewalls are required and the rationale is three-fold:

- Different *types* of firewalls (proxy vs. stateful inspection) can offer better protection. It becomes exponentially more difficult for a hacker to get through multiple types of firewalls.
- Using multiple firewalls from different vendor's offer enhanced protection. No two vendors could possibly design and implement their firewall code (software) exactly the same way. By taking advantage of that fact you are enhancing perimeter security - it's the old "two heads are better than one" axiom. Al Potter of ICSA labs endorses this model and he has a unique way



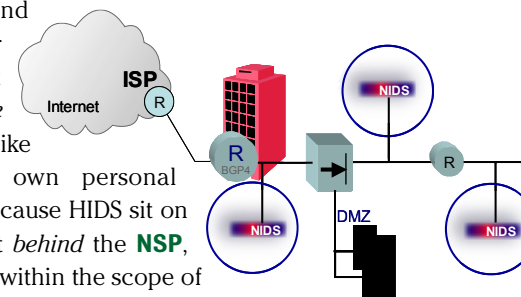
of explaining the strengths inherent in a multivendor, multifirewall approach: "About 90% of the bovine herds in North America can be traced back to the same parents. Because they all share the same genetic code and therefore the same genetic weaknesses, it would only take a single viral strain to wipe out almost the entire cow population". Not good if you're in the fast-food business; catastrophic if it's your **NSP** being eliminated.



- The third reason for using multiple firewalls addresses performance and scalability issues: More boxes translate into more horsepower. Not only use the firewalls in single file, but multiples of them in parallel to manage the workload. Then Firewall **load balancers** typically are used to evenly share the traffic amongst the firewalls.

NIDS

Intrusion-detection systems (**IDS**) come in two flavors: network-based (NIDS) and host-based (HIDS). They are similar in concept but play different roles. HIDS sit on a host — typically servers — and watch for attacks against a *single machine*. It's like having your own personal bodyguard. Because HIDS sit on servers that sit *behind* the **NSP**, they don't fall within the scope of this report.



Secure every client
in your enterprise.

[Click here for your free white paper.](#)

[Symantec](#)

NIDS, on the other hand, have been high flying in the press. A network-based IDS looks for abnormal traffic flows and flags them as being suspicious. Paramount to a successful deployment is fine-tuning the NIDS **sensors**. The sensors — which essentially are customized computers connected to your network via Ethernet — use two main techniques to effectively do their work.

Two main analysis methods

In the past, NIDS supported only one of the two methods, which include *pattern matching* and *protocol decoding* techniques. Most NIDS now use both methods to catch threatening traffic.

- **Pattern matching.** Pattern matching is based on comparing IP traffic to signature files which are libraries of known attack patterns. The signature files are updated regularly.

NIDS sensors are tuneable in a variety of ways — partly by limiting the number of signatures to look for. But at best it's a shell game. Setting sensor thresholds too low could let an attack to slip through; too high a threshold and you have the equivalent of a hyperactive airport security guard: The line of people requiring further interrogation just ballooned from 20 people to 3000. An impossible task for the network SWAT team. NIDS can generate so many alerts that they create a sort of human denial-of-service attack, known as **blinding-the-operator** within the industry.

Compounding this problem, or enhancing your security perimeter as the case may be, sensors are rarely deployed in ones or twos. They are strategically located in multiple locations, and it is not unheard of for a large organization to use more than 1,000 HIDS and more than 100 NIDS. Think “final game of the World Series” and count how many security guards are on duty and where they are stationed, and you get a feel for how your NIDS should be deployed.

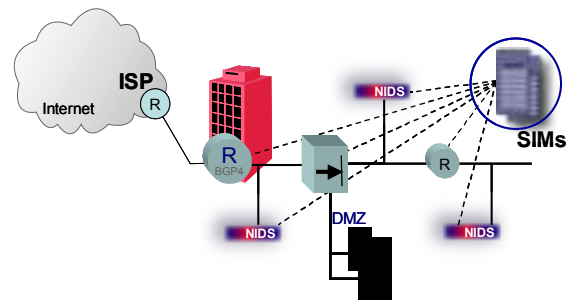
- **Protocol decoding.** Protocol decoding is based on the notion that any variation on the established protocol rules, points to a possible attack. All the IP protocols follow known, published rules in affecting communications.

Protocol decoding can catch attacks that pattern matching might miss. An analogy for protocol decoding would be an envelope sent to someone in Anytown, USA containing a letter dated September 32, 2002. Pattern matching applied against the envelope wouldn't catch the problem, but a protocol decoder would. The protocol decoder knows protocol *rules* and would quickly determine that such a date is illegal and identify that as a potential threat. In the context of IP protocols there are clear rules on what should appear in certain fields in each packet and there are boundaries

on the data found within the headers. For example The **DNS pointer-past-end** attack was based on invalid data in the header, and was caught with a protocol decoder; it was missed by a pattern-matching NIDS.

SIM

SIMS (also known as enterprise security management systems) have found an important niche in helping security analyst's deal with the enormous volumes of data that perimeter devices can produce. Even with appropriate **IDS** sensor tuning, multiple **sensors** can produce enormous amounts of information making it an impractical task to sort through and make any reasonable conclusions.



SIMS collect, aggregate and normalize vast amounts of data from a variety of networking components including:

- NIDS event data;
- HIDS event data;
- **Firewalls** logs and event data;
- Routers logs;



The Next Killer App Replication

It's a fact of business; secure, reliable backups are a necessity. ARE YOU UP TO SPEED? A new white paper will show you what data needs replicating and the reasons why companies must use replication software. Click here to receive your complimentary copy.



True, real-time data replication for Windows

[NSI Software](http://www.nsi-software.com)

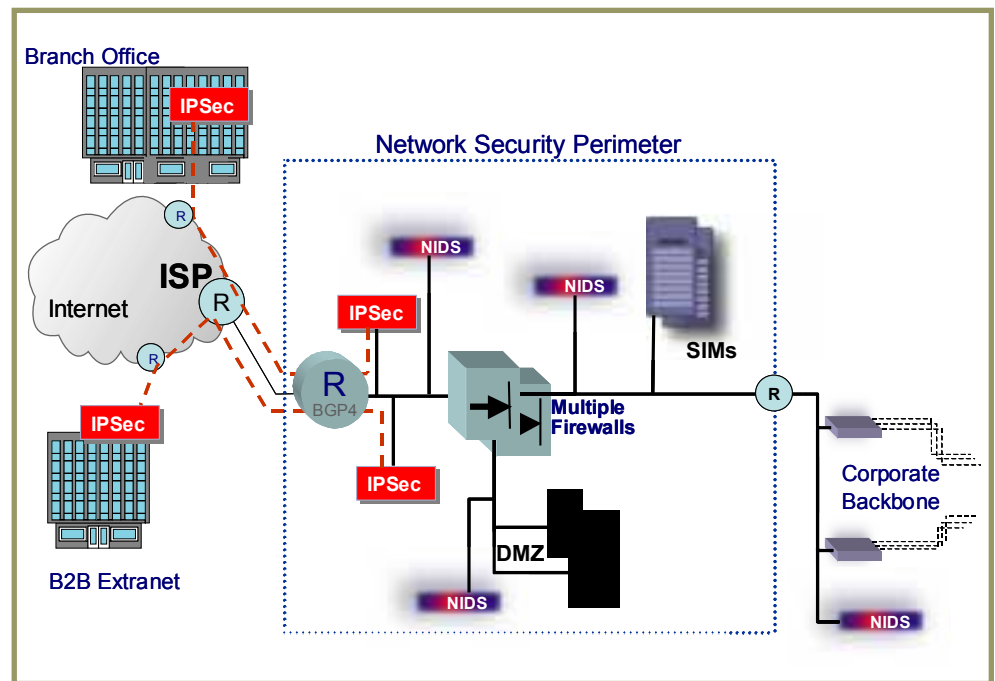
- Vulnerability scanners information;
- Manager of Managers data;
- Lightweight Directory Access Protocol directories;
- Public-key infrastructure system information;
- Operating system logs;
- Web logs;
- Telephone system logs;
- Biometric system data;

SIMS empower system administrators to manage this data better, and it lets them determine if two or three seemingly unrelated traffic abnormalities are related, ensnaring an otherwise impossible-to-detect attack. This would be similar to observing oddly behaving passengers (i.e. those with no luggage; using one-way tickets; same destination), arriving at different airline check-in counters at around the same time of day, and being able to pattern match their behavior enough to send up an alarm.

SIMS allow the security analyst to quickly drill up, down and across, and use various views of security to determine if the attack situation is real and understand the nature of the situation and respond effectively.

IPSec not only gives you encrypted tunnels for passing traffic across any IP network, it gives you strong authentication techniques. Authentication techniques give you high levels of assurance that the traffic has come from a trusted, known source. This assurance comes from special passwords that are secret to your organization alone, and these secrets are managed by the IPSec protocol suite.

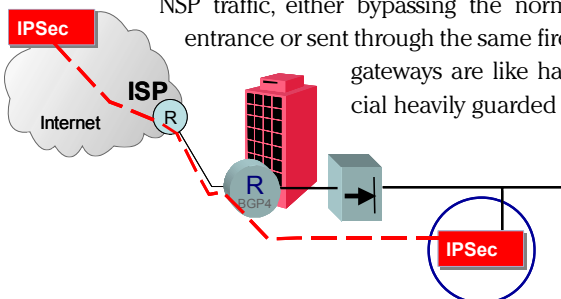
NSP architecture



IPSec VPNs

Some might argue that **IPSec** protocols don't belong in a discussion on **NSPs**. However, part of doing business in the wired world is allowing easy access into your protected environment by trusted individuals, for example your business-to-business partners, your road warriors, your branch offices and others.

This is where IPSec fits in. It works alongside your NSP and lets authorized users send sensitive data across the Internet via encrypted "tunnels". The packets arrive via the IPSec gateway as part of your NSP traffic, either bypassing the normal firewall entrance or sent through the same firewall. IPSec gateways are like having a special heavily guarded entrance.



"70% of unauthorized access to information systems is committed by employees."*

Secure the network from the inside out with identity management.

Learn how »

waveset
Secure Identity Management

How sure are you about your users?

*Source: Gartner: Enterprises and Employees: The Growth of Distrust

Waveset

Type of threats you are vulnerable to

There are a variety of network-related threats that ISP and corporate users can fall victim to:

WHAT is the attacker trying to achieve?

- Undermine network availability.
- Gain access to sensitive corporate data.

Much of the attention in the press focuses on hackers trying to take corporate sites out of commission, and that is a serious threat. This can translate into a significant loss of revenue if you're a financial institution or e-commerce site. There is also the perception issue: Visitors unable to reach your corporate Web site may question your level of professionalism and competency.

However, gaining access to sensitive corporate information is also the subject of attacks. Getting into a military installation's database might be more rewarding, so to speak, than just messing up access on eBay for a day.

WHO is attacking your network?

- hobby hacker
- corporate espionage
- internal spy/mole
- disgruntled employee

The typical hacker profile ranges from the hobbyist and the on-a-mission misfits, to the corporate spies. The first group is usually content to hurt your business by affecting your network availability — through denial-of service (**DoS**) type attacks. But we're talking about a whole different ballgame with the latter group. Well-placed spies or moles make every effort *not* to disrupt network traffic flows, or to bring attention to their work as they slowly, methodically try to gain access to databases and information stored on sensitive servers.

Paul Sop, CTO for Intellitactics, has seen this dark side of corporate espionage: "I can count on my fingers and toes how many attacks that I've seen initiated from the inside, by corporate spies. Our SIM helped identify and alert the network security team about this one employee's usage patterns. He was later seen being handcuffed and escorted from the building. You'd be surprised to know how often this sort of thing is happening out there."

WHERE is the abuse/attack taking place?

- External attacks trying to get into your network.
- Abuse initiated from within your organization.
- Attacks from within.

Attacks can originate from the outside - in which case they are mostly **DoS** attacks, or attempts to penetrate your infrastructure. Inside, however your network security team can be kept busy on a number of fronts. There are the aforementioned moles and spies that might have to be dealt with. There might be dissatisfied employees that trip up systems or vent their frustration on the network infrastructure. These employees, when caught, make great examples of how your organization will not tolerate malicious use of corporate resources. And finally there is another class to deal with: the abusers. Not hackers or attackers per se but nonetheless using the network in ways not intended. These are the chronic online gamblers, MP3, and porno downloaders that waste valuable bandwidth and disk space not to mention their time at work. They potentially are putting the corporation at risk for improper Internet usage.

HOW the attack is accomplished?

- Executable based (virus).
- Network based.

Executable-based attacks include viruses, Trojan horses and worms because they work at the executable level - meaning they are based on software programs and as such run as executables. This malware can get into and run on anything that uses an unprotected operating system (servers, desktop computers and laptops), and unfortunately there are many such "Swiss cheese"



[NetSolve](#)

operating systems on many corporate desktops and server farms. Antivirus software is used to help protect against these threats. Virus software and the like often are not included in a **NSP** design because technically speaking they aren't part of the perimeter. They sit on the inside of your network - on the servers.

On the other hand, Network-based attacks are attacks that try to penetrate network devices. DoS attacks exploit known vulnerabilities in a system rendering it unusable; reconnaissance attacks scan your network and map network details for use in later attacks.

Market trends

While most of the products mentioned (access routers, **firewalls**, **IDS**, **SIMS**) continue to expand their feature sets, the most noticeable trends are related to enhancing performance:

Horsepower.

First-generation **NSP** products like NIDS focused on detection; second-generation products are cranking up their horsepower recognizing that high volume traffic flows can lead to missed packets, which possibly can lead to a missed attack. Most devices support one Gigabit per second Ethernet links.

Parallel devices vs. inline.

NIDS are in-parallel devices meaning they observe traffic as it passes by, but traffic does not flow directly through the NIDS. In contrast, routers and firewalls are inline devices - traffic has to flow through them to get anywhere. There is a design trade-off between the two methods:

- In-parallel devices cannot bottleneck traffic. The trade-off, however, could be an over-stressed NIDS. When overtaxed they can miss traffic and therefore an attack. Worse, some NIDS fail to log that they were missing packets. These problems can be identified and verified with testing.
- Some vendors are making plans to offer inline products. The trade-off is a much higher detection rate (assuming they can keep up with the traffic flow. The trade-off, however, is the potential for the device to bottleneck traffic and degrade performance.
 - Noteworthy is that these products are mostly targeted at small-to-midsize enterprises, which typically do not have the traffic volumes of a larger corporate site.

Integration of firewalls with NIDS functionality.

Integration means fewer devices, which should translate into faster processing. However, the trade-off is clear: Can one box, *should one box*, be trying to do two functions? Do we need a

microwave oven that also can make coffee? If it excels at both, why not? On the other hand, if the oven undercooks when the coffee's being made, then we can live without it. Integrated components might be better suited to small-to-midsize enterprise sites, however their suitability in your NSP can be verified with testing.

Pro-active systems.

Another trend is in the migration from a "detect and alert" mindset, to a more proactive one: "detect and eliminate" the threat while the attack is in progress.

- Some **IDS** essentially disconnect the TCP session between the hacker and the end-system being hacked. - "*act now; ask questions later*", appears to be this year's security credo.
- Another method is for the NIDS to be able to communicate alerts back to the routers and **firewalls** that immediately shutdown access on those connections. This is done through shunning APIs.
 - Unfortunately this opens the door for an attack based on tricking your system into shutting down important connections, so it's a heated debate in the industry right now as to whether this approach should be used.

Powerful SIM data-analysis tools.

Not surprisingly, the SIM products are making themselves indispensable by understanding the difficulty faced by the network security team in dealing with excess numbers of alerts.

The approach is better, smarter, intelligent software that can really find the needle in the haystack. If a critical attack has been discovered, the SIM console alerts an operator with the facts and

FREE Security White Paper

INTERNET SECURITY SYSTEMS™

Are there unauthorized, vulnerable access points in your wireless network?

Wireless access points are designed for ease of use -- not for security.

Download this FREE white paper to learn how Internet Security Systems can help organizations realize the productivity gains of wireless (WLAN) without sacrificing the security of their data infrastructure.

Click to DOWNLOAD FREE White Paper >>

Wireless LAN Security
802.11b and Corporate Networks

www.iss.net

[Internet Security Systems](http://www.iss.net)

makes suggestions on how to rectify the problem. This goes a long way to removing the need for pro-active systems, and it correctly places the emphasis on assisting the security analyst to make the right decision.

More scalable products.

Businesses continue to deal with the need to handle thousands of customers, while offering good performance. A fortunate offshoot of scalability is one of **business continuity**. Inherent in a well-designed, scalable **NSP** architecture, is one that also features **high availability** with no single point of failure. That is because scalability often is achieved by using multiple devices used in parallel.

Some innovative niche products also have been produced that allow for a highly scalable NSP. These components are examined in the next section and include:

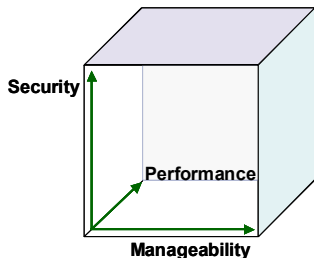
- **BGP4** route controllers that intelligently handle multiple **ISP** access links;
- firewall clustering techniques;
- Firewall load balancers;

Beyond basic NSPs: Scalability & performance Issues

NSPs have spawned a new breed of appliances and network devices. Some do not perform security functions per se, however they are adjuncts to an NSP and have an important role in assuring that the NSP can operate in a high-availability, scalable high-performance fashion.

Building an NSP for performance

It is one thing to build an **NSP**; it is quite another thing to build one that your customers will like to use. Security comes with overhead. It takes time and effort for the police to search for and pick off the bad guys, likewise for network security. Your NSP design and implementation may negatively affect your network performance if not done correctly.

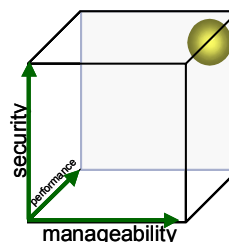


A little degradation in performance might be acceptable if it means a robust, hardened security perimeter has been estab-

lished. Too much degradation, however, and you might lose customers because of click-aways, which translate into lost business.

Why would they click away? "Because they got tired of waiting for the HTML page to appear," says Philip Joung of Caw Networks. "Typically users are willing to wait up to 8 seconds. After that, they get impatient and move on." This is especially hard on e-commerce sites because it directly translates into lost revenue.

The challenge, therefore, is to build not only an NSP that has robust security, but to build one that offers high performance, plus the ability to scale - to handle an ever-growing user community base. Scalability is equally important because for your business to flourish, you need an ever growing number of visitors and customers to come to your site.



So the ideal location to be, in terms of creating an NSP, is in the top-right corner of the SPM (Security, Performance and Manageability) box. There you have achieved robust security levels that offer excellent performance and scalability, with a minimum of management headaches.

Getting the right fit

However, getting into the top-right corner of the SPM box might not be economically feasible. If budget concerns were set aside for a moment - and we could dream in color - we could design a network of highly scalable proportions - one that offered extraordinary performance levels.

But for those who need to build out effective **NSPs** with good performance, they must find a balance between budget constraints and network requirements. The answer might be in finding the right performance levels for your network, for your needs. ICSA Labs' Al Potter agrees with this notion and compares it this way: "When you buy a coat, you don't take the biggest one off the rack - you get one that fits you. Same with performance: People should be more worried about - does it fit?"

Performance tuning

Establishing and validating performance objectives is best done in a controlled environment. This necessitates an environment that lets you test your NSP deployment before going live - a highly recommended practice. CAW's Philip Joung says, "a well-designed NSP will have been tested beforehand, in order to determine its stability, performance characteristics, and ultimately its ability to keep a network secure."

A controlled environment lets you tweak parameters with NSP devices - be it your firewall configuration or your **IDS sensors** -

and compare your test results to previous tests in the exact same configuration. This is enjoyable work because you can baseline a network, run the tests and get black-and-white feedback on your efforts. No vague guesstimations here. Being able to measure browser to Web server response time, as you incrementally increase the number of sessions and packets flows, is what allows one to pinpoint where the architectural weaknesses are.

Testing the validity of your security tuning requires enormous amounts of, and a wide variety of traffic. Products like CAW's WebAvalanche - which can simulate over a million Internet users all sending different types of traffic streams, can let you analyze the NSP at full-throttle conditions. Then you manually introduce an attack. Using your NIDS or **SIMS** console, you can scrutinize the NSPs behaviour and determine if the attack slipped through undetected. If your NSP missed the attack, it might be due to improper perimeter configuration or any one of the NSP device's inability to sense certain attacks under high traffic volume conditions.

By stress testing your NSP in a controlled lab environment, you accomplish four things:

1. You can determine what your NSP capacity is in terms of numbers of users (millions of connections);
2. You can determine client-to-server performance metrics over a range of traffic volumes (millions of packets per second);
3. You can determine client-to-server performance characteristics over a range of IP services (TCP,UDP, **HTTP**, **HTTPS**, **SMTP**,**DNS**,**TELNET** etc).
4. You can test your **NSP** functionality and therefore your network's vulnerability under high stress conditions as well as under normal load conditions.

Enhancing NSP performance

Several new approaches to solving the performance conundrum have emerged on the market over the past year. These devices are adjuncts to an NSP and as such they are not security devices. Nonetheless they offer extremely high benefits in terms of increasing the scalability, the redundancy and the performance of a network security perimeter. These approaches include:

- **Multihomed ISPs** using **BGP4**
- Hi-capacity access links
- Intelligent BGP4 route controllers
- Firewall **load balancers**
- Firewall clustering techniques

Multihomed ISPs and high-capacity access links

Using multiple **ISP** connections solves several design issues:

1. It lets a corporation use more than one service provider. If problems arise with one SP, it doesn't spell disaster for your online presence because you have another service provider connection to rely on.
2. It allows for redundancy, which means access to the Internet has no single point of failure. This is good from a business-continuity perspective.
3. It allows for much higher overall bandwidth - two fat pipes can carry twice as much information as one. Having two access connections means having two **BGP4** routers. To maximize performance on these dual links, a relatively new breed of products has emerged, known collectively as "Intelligent route controllers".

Intelligent route controllers

These products talk directly to the BGP4 border routers and juggle their routing tables in such a way that the routers achieve highly efficient load balancing over the two access links. This is much harder to do than it sounds, especially when you consider that two different service providers might be connected to the two routers.

The bottom line is that a corporation can achieve scalability, high performance, **high availability** and have no single point of failure on the access side, through the use of these products.

Vendors in this space include netVmg and RouteScience

Firewall Load balancers

Firewall **load balancers** have been around for some time and are well understood and work in a stable and effective way. A firewall load balancer keeps track of how busy each of the **firewalls** are, then dynamically directs traffic to the least-busy one. This lets you to build a multivendor, multitype firewall configuration as discussed earlier within this report. The bottom line is firewall load balancers help you achieve high performance.

Vendors in this space include Alteon, Foundry Networks, Nortel and Radware.

Firewall clustering techniques

This is a relatively new space. As the name implies, firewall clustering, is a technique that allows up to 16 firewalls be interconnected over One-Gigabit-per-second Ethernet links. This allows the firewall cluster to scale up to millions of concurrent users. This clustering method generally supports industry standard hardware (Sun, Compaq) and common firewall platforms

(Check Point Software and Symantec). This method also supports a redundant model that is resilient to failure offering no single point of failure

Vendors in this space: RainFinity and Stonesoft



[Request Reprint](#)

[Request a reprint of this report.](#)

Conclusions

- **NSPs** have become necessary as a result of our increasing dependency on electronic communications via the Internet.
- NSPs are a collection of devices including **BGP4** routers, **firewalls**, **IDS**, **SIMS** and **IPSec** gateways.
- A best practice technique is to use multiple types of firewalls within the NSP, as well as to use firewalls from different vendors.
- Many NSPs use firewall **load balancers** to evenly distribute the traffic loads between their multiple firewalls.
- NIDS use **sensors** to detect attacks, and many sensors (100+) are used in large enterprise deployments.
- NIDS require appropriate tuning, otherwise they can blind the NIDS console operator with too much event data, putting the security of the network at risk.
- SIMS are key to analyzing collected data. They aggregate and normalize vast amounts of event data from NSP components.
- Threats can originate from the inside as well as the outside of your organization.
- There are two main types of attacks that can confront an NSP:
 - DoS attacks attempt to reduce network availability;
 - Theft of information: Corporate trade secrets on one end of the spectrum and simply network port information on the other; both can be detrimental in the wrong hands.
- NSPs can achieve high performance levels through the use of Multihomed ISPs with Intelligent BGP4 route controllers and high-capacity access links. Firewall load balancers and firewall-clustering techniques offer high performance as well as reducing single points of failure in the NSP.