



**Axel Buecker
Per Andreas
Scott Paisley**

Understanding IT Perimeter Security

This IBM® Redpaper publication takes a close look at the enterprise IT network perimeter, which has been diluted from a well-defined set of ingress and egress points to a mesh of undetectable flows from devices capable of accessing and penetrating corporate resources. The days of keeping the bad guys out by building a well-defined wall are definitely over. Businesses and organizations require collaboration with internal and external business partners, customers, and employees, which further removes walls and protective barriers.

In this paper, we discuss how the variety of endpoints that were once considered to be *inside* have now become the perimeter itself. With this idea in mind, we investigate how you can build a strong security solution to protect the valuable assets accessible through the IT infrastructure.

The target audience for this paper is IT architects, IT specialists, and security administrators.

A little history

IT perimeter security is a fairly broad term that has a diverse set of implications and meanings. It is quite common to misunderstand the nuances implied by the term.

In the beginning of the digital computer age, computer systems were single stand-alone entities, tipsily located in a physically secured room known as the *machine room*. Input and output media were hand-carried into the room or handled at the *remote job entry* (RJE) location. Both the machine room and the RJE were subject to the physical security and access controls in place at the given location. Because the computational center was located in well-defined locations, it was easy to identify sources of entry. Therefore:

The perimeter was very well defined, and security could be enforced on a physical level.

The next phase of computing introduced terminals, where a keyboard and monitor were wired directly to the central computer system. With this approach, input could be submitted from various locations, no longer restricted to the physical location of the central computer system. However, a certain proximity to the central computer system was required. Access controls were needed as an integrated part of the computer. Physical constraints dictated the distance between the computer and its terminal and keyboard.

The perimeter was still well defined. However, physical security was no longer sufficient.

The proximity clause changed once modems, or other means of single-point remote access, were introduced to enable computers or terminals to communicate directly with the central computer system. This new infrastructure layer also required an additional access control layer because access control enforced at the central system was no longer sufficient.

Even though these systems were “remote,” the perimeter was still defined. Security enforcement required additional access controls. Furthermore, because multiple users could access the same CPU, each user had to be monitored (authenticated) for this use.

The single-point remote access paradigm shifted dramatically with the dissemination of the Internet, which connected these large CPU systems, such as mainframes, *to each other.*

Authentication at the local system was lost as systems became *personal*. These personal computers (PCs) had no need for authentication; as the name implies, they were intended for personal use and home use. These PCs also began to be networked together through the use of modems and eventually connected to the Internet through technical means such as high-speed broadband or DSL access devices. The connections today provide access lines even into private homes, and the bandwidth provided to the user community is growing every year.

As the systems grew and became more powerful, authentication was reintroduced into the personal computers as the PCs themselves could be remotely accessed.

This development is used by companies and employees alike to promote home working environments and past-office-hours work through connections to the corporate network. These connections use high-speed lines as a carrier into and over the Internet through a corporate VPN (virtual private network) entry point into the corporate network infrastructure. For example, the corporate network extends seamlessly to the employee’s home.

The same technology is being used for customer and business partner access to the corporate application and data resources as well.

Devices that can use Internet technology are no longer bound to well-known personal computers. Many sorts of wireless devices allow people to transparently gain access to the Internet and, with this access, to gain access to a corporate IT environment. Because the systems and applications are interdependent and connected, it is often difficult to even know where the application is hosted or on what computer the application is executing. Devices such as vending machines, telephones, medical equipment, and manufacturing equipment all have the ability to access the Internet and can even be accessed remotely.

The perimeter is now becoming fuzzy. Any sort of computing device may become the perimeter itself, and these devices in many cases are mobile.

This introduces us to a new concept. If the network perimeter has eroded, then what is the perimeter? The network perimeter has become a dynamic changing barrier that you must redefine and protect. The problem arises when you think and view the network perimeter as a static barrier because *it is not!* The systems that interact with the network perimeter make this network dynamic, and thus you must protect it by defining a *system perimeter* that understands and is capable of being a part of the network perimeter. Another issue is that applications introduced by a Web browser and run on local machines are difficult to control with traditional network perimeter tools. The systems do not even have to move to introduce unwanted access on the system itself.

The winner and loser conundrum

Today's IT environment is the result of numerous battles between technologies. Technologies, which give the user equal or better possibilities, are always competing. While users decide and define the market, the better technology does not always win. Trade-offs for security, bandwidth, stability, and speed are always fighting for popularity.

Betamax/VHS/Video 2000 battle

We are all familiar with the old battle between the different video systems, such as Betamax, VHS, and Video 2000. A winner was declared by popular vote. It is fascinating that the format with the lower video quality signal won! However, VHS also had the ability to store more video on a single tape. Thus the *best technology* did in fact win.

“IP on everything”

This prominent quote originates from Vinton G. Cerf, VP and Chief Internet Evangelist at Google. He wore his famous T-shirt with this quote at an IETF conference in 1992.

The Internet Protocol suite won the battle (if there ever was one) and is now the dominant suite of Internet communication protocols. The victory has been so overwhelming that it seems as though almost everything is IP enabled.

The IPv4 suite, however, has inadvertently introduced several security concerns that must be addressed when IP enabling all kind of devices.

IP Version 6 (IPv6) attempts to address many of the concerns; however, IPv6 security, and specifically perimeter security, will always be a concern as devices become the new perimeter.

The IT perimeter: A definition

The IT perimeter - back then . . .

Defining the perimeter has usually been an easy task. There was the cave to be protected; there was the village; there was the castle with its wall. It was easy to define, visualize, and create a protection policy to enforce and protect the obvious boundaries of this perimeter.

History in some ways repeats itself. First there was the computer. Then a number of computers, followed by the first small network of computers. Then the network grew and was connected to other networks—today all IP devices are directly connected to networks and each other.

Until recently the IT perimeter definition was less complicated and better understood. A firewall was all that was necessary to define the network perimeter. Everything inside the firewall was considered a trusted insider; everything outside was not so good . . . from a network perspective at least.

The IT perimeter - today. . .

What does the revised definition of the perimeter look like today, or how should it look? The perimeter is becoming more and more defined by each node on your network and not the network itself. In addition, network protocols themselves have been enabled to allow applications to traverse through the firewall and run on local machines (for example, Java™ and JavaScript™).

Some of the devices that break traditional perimeter security are:

- ▶ Applications that traverse through firewall policies
- ▶ Mobile devices
- ▶ IP-enabled devices internal to the network
- ▶ External devices that are “allowed” on the internal network “temporarily”
- ▶ Wireless access points that are unknowingly deployed
- ▶ Direct Internet access from devices

Applications have to be accessed by users and other applications to fulfill their purpose. This access, however, can expose the application to unwanted access. In general, ease of use is a concern; for example, users in general have to authenticate themselves only under rare circumstances to gain access to an application. The application is left open; thus promoting the application server to double as a perimeter.

Mobile devices are, in fact, mobile; their nature is to be moved and connect to various networks at various locations. Some connection points can be within the organization’s perimeter; others are not. This requires the mobile device in actuality to act as a perimeter, thus being enabled and configured to that end.

IP-enabled devices internal to the network often require a number of open ports in the firewall. Sometimes they even must be contacted from the Internet in order for them to function properly. To keep up with technological advances, these devices are often IP enabled after their initial configuration, and thus they are required to act as a perimeter as well, sometimes to protect an IP implementation where corners were cut to enable the device’s functionality.

External devices allowed on to the internal network temporarily can be a major threat for internal IT security. These devices are typically not scanned for viruses; access is often granted to an unrestricted network segment, and thus all devices in the network must act as one perimeter against these external devices.

The introduction of wireless technology probably had the most impact on opening internal networks to external threats. Unprotected and unknowingly deployed wireless access points still represent major loopholes into the enterprise network, as shown by various drive-by attacks.

Direct Internet access from any device is one of the most difficult to control from an IT organization’s point of view. These can be personal devices, not owned by the IT department. By connecting directly to a computer, these devices sometimes can enable the host computer to bypass traditional perimeter security controls - for example, when a “smart cell-phone” is connected to a computer and the computer can then access the Internet using the cell phones’ internal modem capabilities. The user can then disconnect this device and reconnect the computer back on to the IT infrastructure. Thus you need to take into account these types of capabilities and assume that this type of activity can happen in the network. You have to find a way to ensure that you can address this issue, and you have to redefine your perimeter with these types of access methods in mind.

Define your perimeter

Any network owner is required to know the full layout of the enterprise network. But if every node is the perimeter itself, then the layout of the network is less of an issue with regard to the perimeter boundaries.

Because the network has become extremely dynamic, you must ensure a vigilant exploration of this ever-changing network. Scanning and assessment must be continuous and ensure that you can identify misuse and abuse of the network and its IT resources.

The key to successfully defining the network perimeter is a combination of automated network tools and the ability to globally enforce host-based security software deployed to the mobile systems that you know access the network. Scanning and the discovering of unknown devices also must be considered because by definition, these unknown entities may constitute a perimeter breach.

Analysis tools

There are two basic approaches to analyzing the perimeter and the traffic around and through it using automated tools. In this Redpaper publication, we refer to these two types as *passive* and *active monitoring tools*. However, both methods have one thing in common—they produce log files, which always must be evaluated.

Passive monitoring tools

A good vulnerability and network scanner (such as the IBM Internet Scanner® or IBM Proventia® Network Enterprise Scanner) can be an effective way of discovering devices connected to the network, and what the discovered devices are capable of doing in the network. In addition to device discovery, these scanners can report on the vulnerabilities of the network devices they scan as well as reporting the devices discovered in the network.

The vulnerability assessment application can scan the network for weaknesses and identify more than 1,300 types of networked devices, including desktops, servers, routers or switches, firewalls, security devices, and application routers. When these devices are identified, IBM Internet Scanner analyzes the configurations, patch levels, operating systems, and installed applications to find vulnerabilities that can be exploited by hackers trying to gain unauthorized access. These tools can assist you address the requirement of properly knowing your network layout.

These types of tools are considered *passive* because they do not scan all of the time. They scan the network only when they are invoked. Such passive tools either require setting a scheduled time to scan or are manually invoked.

Active network activity and monitoring software

Tools that scan the network 24x7 are considered to be actively scanning the network and its activity because they monitor traffic patterns, communications, and transmitted data. You can use the IBM Proventia Network Anomaly Detection System (ADS) tool to look for patterns and events, including unwanted IP structures and unknown communication patterns. The tool can help you draw a picture of the enterprise network and create an understanding of the communication patterns between participating devices, which in turn can provide a better understanding of the overall perimeter.

An overview of how ADS works is depicted in Figure 1.

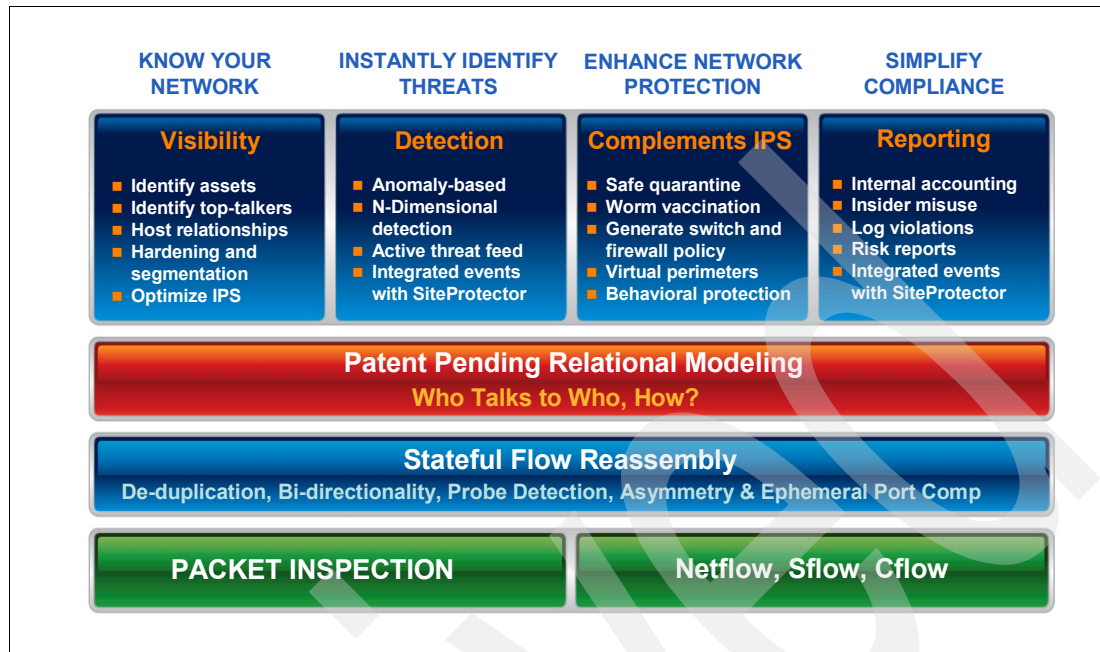


Figure 1 ADS outline

A sample deployment architecture for an ADS is depicted in Figure 2.

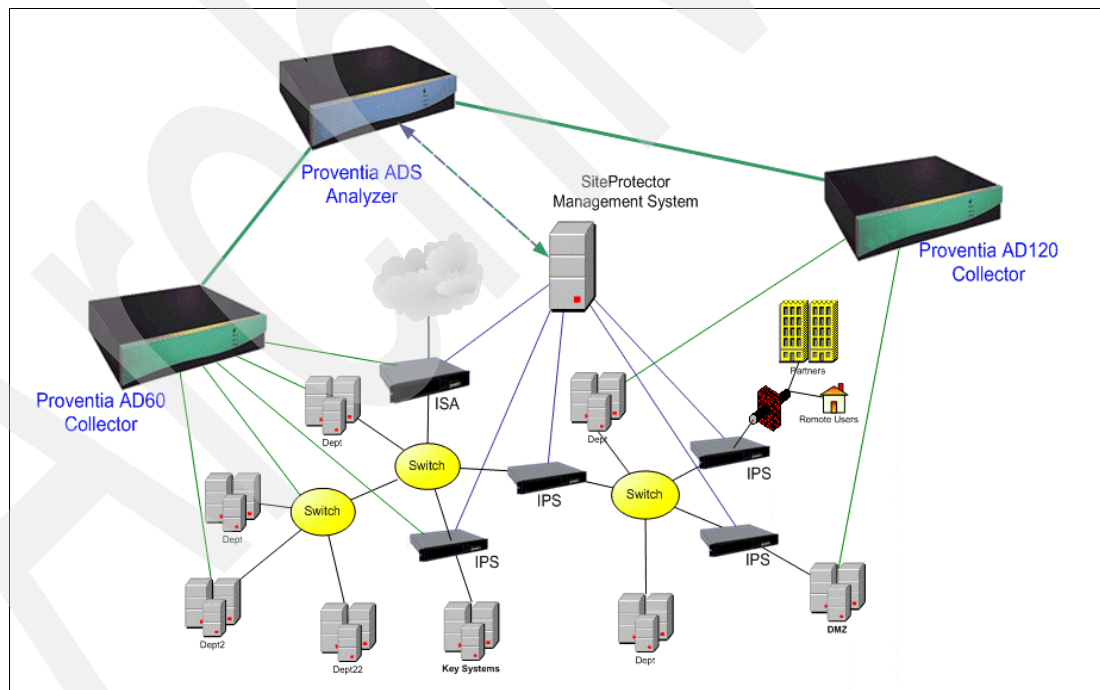


Figure 2 ADS sample deployment architecture

An ADS provides the ability to replay what happens in the network and thus can show you the network as it really *lives*.

The IBM Proventia Network Anomaly Detection System is a *network behavior analysis system* designed from the ground up as an internal network security system. By using network flow

data to determine which users and hosts communicate with each other, and how, Proventia Network ADS can deliver a continuous network inventory and a clear view of your network behavior. It can automatically detect unhealthy traffic, security threats, and noncompliant activities, such as abnormal network performance, worm propagation, and policy breaches.

Because the Proventia Network ADS has a network-wide perspective, it preserves business continuity by enabling you to track and harden threatened resources before vulnerabilities are exploited.

Proventia Network ADS can provide immediate value to your network as a standalone appliance but also integrates seamlessly with intrusion prevention and vulnerability management systems as a component of the IBM Internet Security Systems™ (ISS) protection platform. This integration provides additional value, helping you further develop and enforce security policies, demonstrate regulatory compliance, and harden your network against unauthorized applications and services, all while securing mission-critical data and resources.

Proventia Network ADS simplifies regulatory compliance by monitoring critical assets and applications and keeping track of change management. It identifies and takes action against malicious content, illegal access, insider misuse, and other security incidents, limiting the harmful effects of those incidents and providing critical information for incident response. This real-time security auditing and monitoring allows the creation of easy-to-read, in-depth reports to assist in meeting regulatory compliance objectives, especially the IT requirements set forth by SOX and CobiT.

Logs

All tools rely on analyzing lots and lots of events. The resulting logs must be evaluated constantly and consistently. This task can be partially automated; however, there will always be logs that need to be evaluated and cross-referenced by specialists.

Correlation between various logs should be investigated more rigorously and employed on an enterprise level.

Knowing your perimeter: What is next?

Once the network has been mapped thoroughly, it is time for you to consider revising the way the network is managed. You should segregate your networks into zones and define a data and asset classification. Furthermore, you should consider the fact that each host has become the perimeter.

Network definition

Networks are the mechanism for electronic communication between IT systems. Views of the network and security have changed over time. Network security used to be focused on hard boundaries, with limited access to and from the Internet. Now networks must provide a variety of communications in and out of an organization in a carefully controlled manner.

Network zoning

A key concept in defining a modern perimeter is to create security zones of the network infrastructure as depicted in Figure 3. It is no longer sufficient to use perimeter firewalls to segment important areas. All areas of the network must be part of a security zone, and all nodes must be able to act as the perimeter.

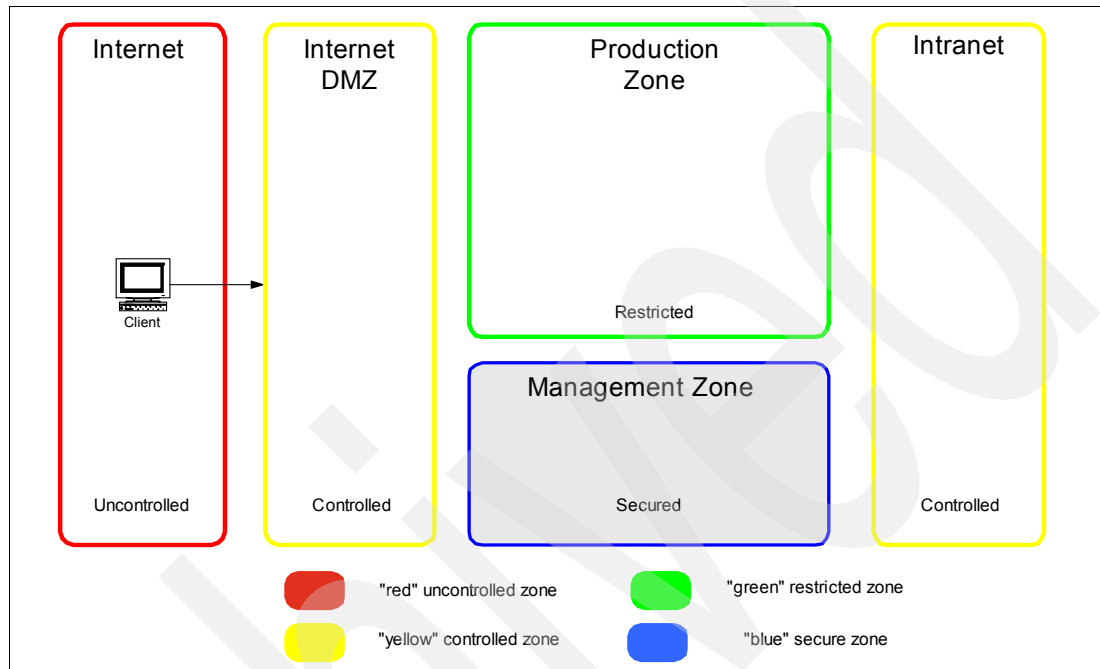


Figure 3 Network zones

Security zoning requires initial classification, and it requires that zone definitions include the various types of mobility and enforcement. A key benefit of a security zone is that in the event of a security breach or incident, the breach will be limited to the zone itself.

For example, if the organization only authenticates users traversing the central VPN solution, the enterprise network is at risk, because most VPN clients are freely downloadable and configurable. It is not sufficient to require user authentication; the workstations should be authenticated as well.

Network boundaries or perimeters are used to isolate networking zones with differing security policies. These boundaries are created to implement restrictions on the type of traffic that is allowed in a zone - for example, restricting access to only HTTP traffic on port 80 and HTTPS traffic on port 443 inbound from the outside to a zone of Web servers. You use a firewall to allow this traffic and block all other. In its simplest case, a firewall is a device that implements a policy regarding network traffic. It creates boundaries between two or more networks and stands as a shield against unwanted penetrations into your environment. But it is not meant to be your only line of defense; rather, it is a mechanism that slows the progress of an intrusion.

One method of shielding information about the network the firewall protects is by re-addressing the packets so that outbound traffic appears to have originated from an address associated with the firewall itself. This re-addressing is called *network address translation* (NAT), and its primary function is to hide the trusted network from untrusted networks.

Classification

Today the classification effort is mainly based upon *data classification* and, to a certain degree, *user classification*. In the modern enterprise, the actual hardware or communication infrastructure is rarely classified. Quality of service (QoS) efforts in the network side have led to a minor classification of communications floating in the network. However, this effort has more to do with upgrading and downgrading traffic types than with a full-effort classification.

An up-to-date classification effort should include the following:

- ▶ User
- ▶ Data
- ▶ Hardware
- ▶ Communication

It is desirable that all four mentioned classes undergo a classification effort, which determines when, where, and how the protection effort should be increased or decreased.

User classification

User classification depends on the role association of actual users of the network and the resources that interact with the network. You must take the discipline of identity management seriously and extend it into the future. You must define users not only by assessing the access they need, but also by their location within the network topology. Finally, you must include security zoning in the user definitions.

Data classification

Data classification is a mature area that is concentrated on file content more than general data classification. It is possible for most users to store documents on the local hard drive, regardless of company policy.

Today data classification in the enterprise is partially implemented through the backup mechanisms in place. However, you should examine whether the data classification is meets quality standards and actually usable.

Hardware classification

Hardware classification today is mostly focused on the tangible assets within an organization. Desktop computers or mobile computers are locked to a immobile item when left in a room with general access. You classify mobile computers for the mobile workforce separately because they are exposed to physical access by outsiders if not properly secured while off premises. You have to make the same considerations with mobile phones, such as PDAs or BlackBerry devices, which are capable of storing classified data.

Printers are typically located in locked printing areas, with only authorized personnel having access to that area. Servers are located in data centers, where access is limited and very controlled.

In addition to the regular computing devices, other IT-related *gadgets* are commonly being deployed today. These include, but are not limited to USB devices such as thumb drives or external hard disks, MP3 players, video cameras, and smart card readers. These devices can also be regarded as the new perimeter and have to be included in hardware as well as data classifications.

Communication classification

Security zoning enables you to classify all communication in the network. Communication must be classified for network management to know which traffic is legal and which is not. You make decisions based on whether traffic must traverse a firewall, and you have to decide where to place intrusion detection systems (IDS) and intrusion prevention systems (IPS), where and when to require user authentication, and so on.

One of your classification efforts should be to avoid or deny encrypted traffic within the enterprise network. All traffic should be available for immediate inspection, through IDS or IPS devices. For example, you can use communication classification to require that IDS must prevent the access of unclassified traffic in the network. This traffic can then be denied access as part of the general policy.

Mobility and connectivity

You must examine the mobility of the nodes on the enterprise network as well, including the network connections the nodes engage in. You must investigate the wireless LAN infrastructure, especially with security in mind. Further, you need to revise how mobile users can connect when not on the enterprise network.

LAN connectivity

LAN connectivity is assumed to be acceptable at any location within in the organization. However, the question is whether this connectivity is still acceptable when security zoning is in place. This consideration may imply that a user is granted physical access to network ports only in specific locations.

The improvement of wireless technology can assist in turning the wireless disadvantage into an advantage by introducing and deploying a full wireless network infrastructure for internal enterprise users as well as visitors such as customers and contractors. You may want to leave the traditional LAN port connections to devices requiring (still) higher network speeds.

Other technologies are available that meet the challenge of specifically authenticating nodes that access the network through a physical LAN port. These solutions can grant or prevent network access based on several criteria, such as the compliance posture of the node or the combined authentication of node and user. In any case, you must weigh the pros and cons, such as costs and risk mitigation aspects, that these solutions can offer.

Wireless connectivity

Most new portable computer devices (and many other mobile devices that are connectivity related) are now delivered equipped with a wireless adapter, enabling wireless connectivity to networks at locations where this service is available. Such locations are numerous; they are found within most organizations and also at public locations, such as airports, cafés, and hotels. This increased availability enables connectivity to (mostly) the Internet through an Internet Service Provider (ISP). Once connected to the Internet, users can connect to the internal enterprise network utilizing a virtual private network (VPN) connection, as discussed in the following section.

Wireless connectivity at non-enterprise locations thus enables users to be available or online within an enterprise network at times and places never before possible. To utilize this advantage, the user must be allowed to connect to unknown wireless networks - for example, it is up to the user's discretion to decide whether or not a network is trustworthy, thus extending the enterprise perimeter and its security to any particular user.

In most interposes, the wireless network infrastructure traditionally traverses the building perimeter, making internal network infrastructure visible to outsiders. Today, however, the enterprise network perimeter no longer follows the building perimeter.

VPN connectivity

VPN connectivity to an enterprise network is an integrated part of every network infrastructure today. VPN connectivity can be divided into two main categories: remote user VPN and site-to-site VPN connectivity between networks.

The latter can be compared to leased-line connectivity between two networks, utilizing the Internet as a carrier. The perimeter of the enterprise is thus not changed because the same rules that have been in place for leased-line connectivity usually apply. However, we strongly recommend an IDS/IPS device in the receiving DMZ, where the VPN connection is terminated.

User VPN connectivity to the enterprise network is almost a given in organizations, enabling users to access internal resources as though they are locally connected to the enterprise network. With user VPN connectivity, the Internet is also used as a carrier medium for the connections, which are encrypted. However, this new wireless infrastructure approach moves the enterprise perimeter out to each remote user because the mobile computer becomes the boundary for the internal enterprise network - wherever it is located.

The mobile computer normally receives an internal enterprise network IP address (depending on the software used), and as long as user authentication is passed successfully, users are treated with the same trust as though they were within the physical building boundaries.

Within VPN software today, the device connecting through the VPN can either be authenticated or required to run certain applications or executables when connectivity is approved. However, it is important to note that most VPN connectivity authentication is still solely based on user credentials, leaving the device out of the authentication process.

Device ports

Other dangers may include Bluetooth connections, USB ports on the devices, and of course non-enterprise network connections.

USB ports on devices in particular are possible targets for attacks on an organization. It is necessary to employ software-monitoring connections to USB ports on devices such as mobile computers, desktops, and servers. Many hazards can originate with USB ports as was the case with floppy drives and CD-ROM drives.

Only a few technologies are available to log connectivity to USB ports on devices. You should examine and evaluate these technologies because they fill a loophole in the enterprise network that is often overlooked.

Enforcement

Referring back to the network zoning diagram in Figure 3 on page 8, the main task in IP perimeter enforcement is to map network zoning and thereafter implement it. This task can be demanding, but it offers benefits, most notably a clean network infrastructure.

All regulations must be enforced; otherwise, the regulation is meaningless. And all regulations must be available. To create a regulation, it is necessary, as a first step, to classify your data and your assets.

The next step is to ensure your regulations make sense, not only to the regulators, but also to the users. This includes reporting and dealing with breaches. *Reporting* is the important word here. Reporting must be an integrated part of the enforcement and must be available to the users.

For example, if the organization decides to enforce device authentication for user VPN connections on top of user authentication, the authentication scheme not only must be understandable to the user, but it must also be easy to use. Otherwise, users will resist authentication and over time apply pressure to change it (or even worse, get it dissolved).

To enforce regulations, it is necessary to have the means to enforce, for example, management of devices as a central point and to enforce device authentication policies. With an infrastructure capable of handling all tasks at hand and configured to suit users' needs and support regulations, the task becomes easier.

The infrastructure is comprised of hardware, software, and management components, but it relies on correct classification and transparent regulations. Part of this must be change-control processes, auditing features, and accepted procedures.

Perimeter security

As discussed previously, every computer system with potential networking capabilities can be considered a perimeter host device. In this light, it is fair to assume that host-based security is very often neglected.

Today organizations would rather protect the network instead of protecting individual host systems. The one exception is anti-virus products installed on most hosts, but almost no other host-based security systems are deployed.

Most IT professionals still rely in the network protecting the specific host systems from malicious content. This assumption, however, is no longer valid and sufficient. Thus you must look for solutions that regard the host as the new perimeter and look at each, network and host, individually as well as in a combined network topology context.

Host definition

In general you have to distinguish between user-oriented workstations and service-oriented servers. Today numerous user workstations are deployed as mobile computers to enable a mobile and flexible workforce. Desktop systems are used for service-oriented tasks with more than one individual user accessing the system.

Servers are usually deployed in central locations, and user access to these machines is granted only through special services like such as file sharing, printer sharing, or application sharing.

You have to consider, however, that every host system uses the same networking infrastructure and the same protocol stack (IP). Thus every single host must be protected in a similar manner from malicious code and attacks. Because you have to handle malicious attacks differently on servers than on user workstations, different products are available for workstations and servers. These products utilize similar sets of technology to recognize malicious code or attacks, but they have to behave differently on the different host systems.

For example, you cannot just simply shut down a service on a production server. Figure 4 illustrates where the relevant Proventia Desktop and Proventia Server products should be deployed within an overall sample IT deployment.

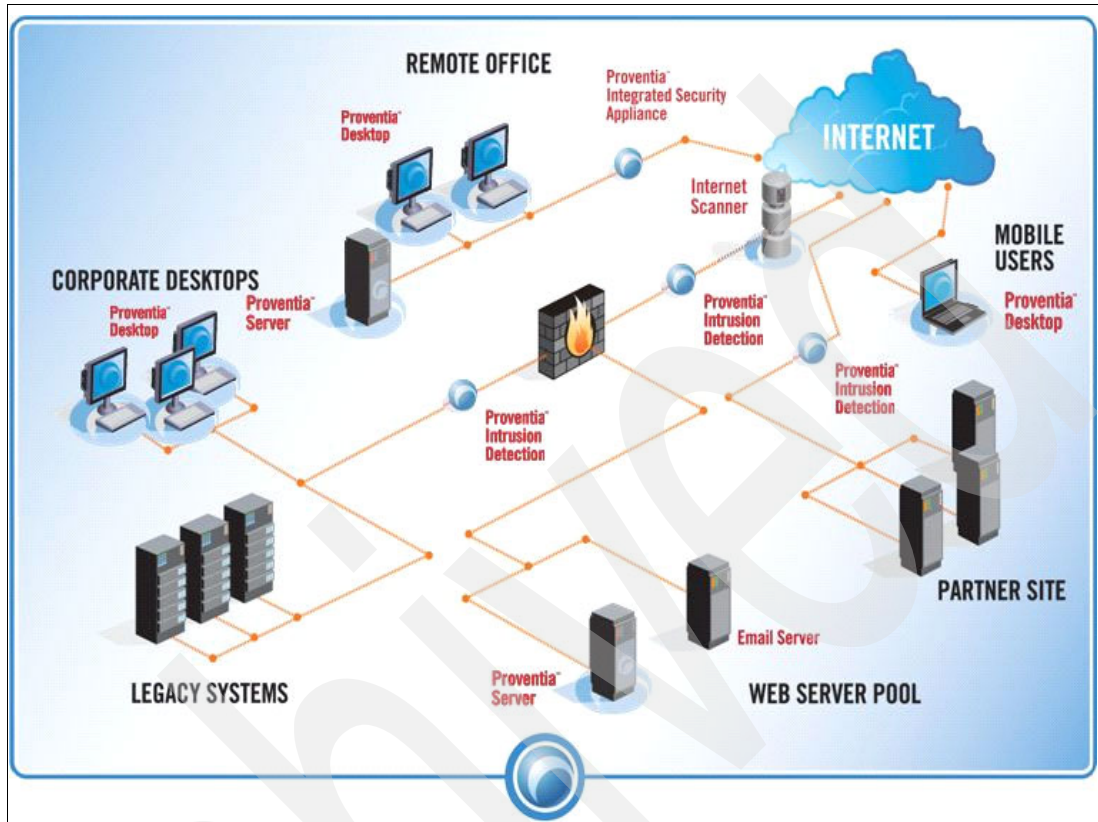


Figure 4 Placement of IBM Proventia components

In addition to desktop and server protection, we also show Figure 4 where to best place the aforementioned intrusion detection and scanner services. A cost-effective integrated security appliance is placed within the remote office perimeter to act as a combined IDS/IPS device. These individual solutions can be centrally managed using the IBM Proventia Management SiteProtector™ product.

Let us now take a closer look at both the desktop-oriented and server-oriented products.

Desktop products

When deployed on all corporate desktops and mobile computers, IBM Proventia Desktop Endpoint Security provides a multi-layered approach for protecting your systems from malicious intruders. This multi-layered design is depicted in Figure 5.

This does not mean that firewalls and access control lists suddenly become obsolete. It simply emphasizes that protecting the perimeter requires more technology than ever before.

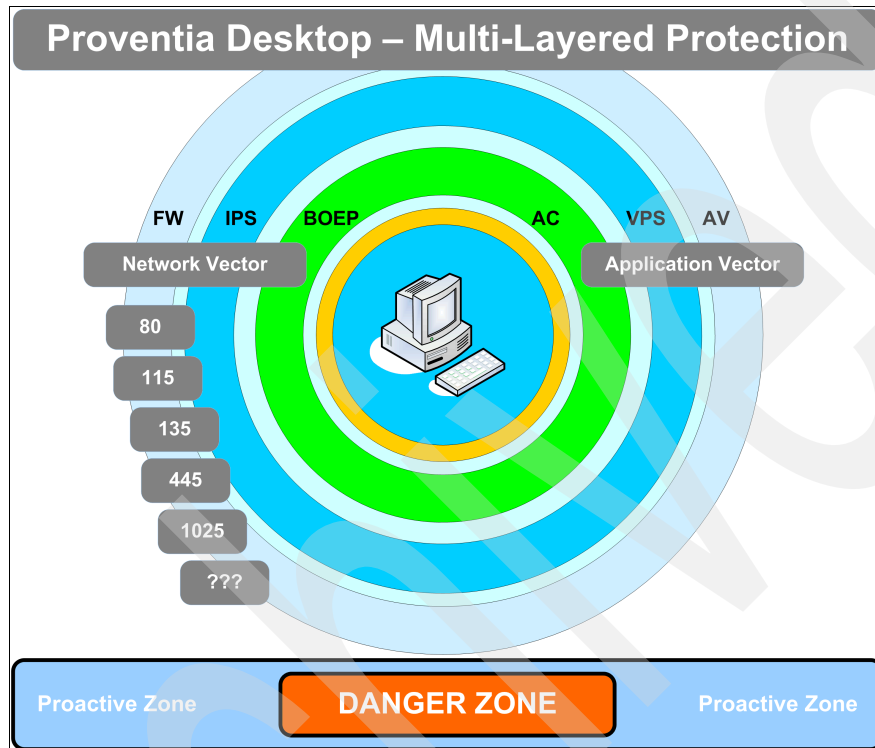


Figure 5 Multi-layered protection

The various components (refer to Figure 5) utilized in IBM Proventia Desktop Endpoint Security are:

- ▶ FW (Firewall)
 - The firewall module operates reactively.
 - User impact can be high if this component is not centrally managed.
- ▶ IPS (Intrusion prevention system)
 - The IPS system protects against all known vulnerabilities and exploits.
 - The user impact is low.
 - It stops the attack.
- ▶ BOEP (Buffer Overflow Exploit Prevention)
 - BOEP protects against *known and unknown* buffer overflow exploits.
 - These exploits represent the majority of attacks.
- ▶ AC (Application control)
 - AC is based on configuration.
 - AC potentially protects against all *known and unknown* attacks.
 - User impact can be high if this component is not centrally managed.

- ▶ VPS (Virus Prevention System)
 - VPS uses behavioral patterns that can detect and block a large number of viruses.
 - It does not require any signature updates.
 - VPS is explained in more detail in the section that follows.
- ▶ AV (Signature anti-virus)
 - AV protects against most known file attacks.
 - AV is almost always reactive.
 - Constant updates are required.

These functions are considered part of the Virtual Patch® initiative, and some are discussed in the sections that follow.

Server product

Both the IBM Proventia Server Intrusion Prevention System and IBM RealSecure® Server Sensor products, together referred to as the *IBM server protection suite*, provide powerful protection technologies in a single multi-layered agent to guard business-critical systems and data from any attack, outside or inside the enterprise. They proactively protect servers from malicious attacks while supporting your compliance needs. To combat threats, the server protection suite combines several protection technologies into a single, multi-layered agent similar to the one shown in the desktop product in Figure 5 on page 14. Offering broad operating system and platform support, the server protection suite also guards business-critical systems and data, helping you to meet stringent audit and compliance standards.

In addition to protection technologies, the server protection suite also helps solve the following key business problems:

- ▶ Data security
 - Provides historical data that enables an organization to find the origin of a change, breach, or string of behavior.
- ▶ Insider threats
 - Tracks the who, what, when, where of user and administrator behavior.
- ▶ Compliance
 - Provides reporting necessary to prove the security of sensitive information.

Protection technologies

We now take a brief look at the following technologies:

- ▶ Protection Analysis Module
- ▶ Virtual Patch
- ▶ Virus Prevention System (VPS)
- ▶ Buffer Overflow Exploit Protection (BOEP)
- ▶ Application control

Protocol Analysis Module

The innovative, patent-pending *Protocol Analysis Module* (PAM) is the underlying detection engine that serves as the foundation for network, server, and desktop IDS/IPS and endpoint security solutions.

PAM separates itself from the competition by:

- ▶ Vulnerability modeling
- ▶ Port variability (port-independent protocol decoding)
- ▶ Number of protocols decoded
- ▶ Number of methods and algorithms
- ▶ Employment of these methods, either singularly or in combination based on the type of attack
- ▶ Order in which these methods are applied
- ▶ Quality of the algorithms themselves

Virtual Patch

Virtual Patch is a method for blocking the exploitation of a vulnerability without applying a vendor patch. Virtual patching involves automatic monitoring, attack recognition, and blocking of specific communications and resource usage based on several factors, including the vulnerability posture, patch level, and OS of the target host.

IBM ISS employs the Virtual Patch *protect-while-you-patch* principle through vulnerability-based intrusion prevention and risk assessment. The Virtual Patch *protect-while-you-patch* principle uses a combination of the IBM ISS patented vulnerability and threat-detection and threat-prevention algorithms and methods along with a module for impact analysis and attack pattern recognition. Using the Virtual Patch principle enables you to easily employ protection policies to protect critical assets from attack and misuse until a physical vendor patch or manual corrective action can be taken. The IBM Virtual Patch enables full protection from potential compromise when combinations of IBM ISS network, server, and desktop IPS solutions are used.

Virus Prevention System (VPS)

The breakthrough Virus Prevention System (VPS) is available in the IBM Proventia Desktop Endpoint Security products. Considered the next generation anti-virus, VPS uses behavioral patterns that can detect and block up to 93% of *new* viruses, without requiring an update. VPS uses a virtual system to detect, analyze, and stop entire families of viruses. It is truly preemptive, not reactive, technology.

Buffer Overflow Exploit Protection (BOEP)

BOEP identifies attempts to execute code (system calls) on writable memory regions. This distinction is important: BOEP does not prevent all buffer overflows, but only those that overrun their bounds and attempt to execute in writable regions of memory.

Application control

Application control, a technology that enables the specification of trusted versus untrusted applications, is also available. If used correctly, application white lists and black lists are an effective method of policy enforcement (through file name and MD5 checksum).

Conclusion: Where the heck is my perimeter?

In short, your perimeter is right in front of you: All computers and IP-enabled devices in the organization can be the perimeter. And every one of them should be treated as such. To accept this statement as valid, we invite you to consider the following brief example.

Where are you reading this document? If it is a hardcopy or print out, you are probably safe in assuming that data cannot leak from your perimeter. What is the classification of this document you are reading? This specific document is public, but it was not public while it was being written. It is important to know the limits of the content of the files and data you have access to.

However, if you are reading this document electronically, consider the connection of the electronic device you are using. Are you connected via an 802.11x wireless access point? Are you connected through a network-enabled device such as a cellular modem? Do you have a VPN back to your home office? Do you know if you have a connection to the Internet? You might be viewing the document on a small hand-held device that has multiple methods of connecting to the Internet. Unfortunately, convenience tends to be the priority in regard to data access, and providing convenient access to data often conflicts with providing secure access.

The point here is that *any* available path for data to migrate should be the first issue taken into account when considering the perimeter. Data can move through a network connection or a simple data connection such as a USB stick. These are the very pathways—even this document you are reading at this very moment—that provide outsiders access to your system. The perimeter begins where the data moves. As long as data movement is intentional and within the confines of your data classification and movement policy, you are in good shape.

Unintentional data movement or data leaks can be better controlled when you consider that the perimeter is literally right in front of you. A combination of network controls on the *home network* with the addition of host protection is critical to knowing where the dynamic perimeter is and how to protect the data that migrates through it.

In conclusion, some accepted methods and definitions are no longer as valid as they used to be. Thus, you must re-evaluate and reconsider them.

- ▶ Encrypted traffic on the internal IT network should be denied.
- ▶ Firewalls regulate traffic but may not prevent an attack through an open, known port.
- ▶ Anti-virus software is reactive.
- ▶ Application control is necessary.
- ▶ Mobile computers, desktops, and servers must be treated equally when it comes to malicious content exposure.
- ▶ User authentication might no longer be enough but nevertheless should be more widely used.
- ▶ The internal network is a mesh of ADSL, wireless, Internet, and leased-line networks, where the user often determines the carrier and freely roams between them.

In this paper, we have intentionally expanded the accepted perimeter definition a little. We have further built a case for deploying detection, protection, and analysis systems that can assist you in the task of redefining your perimeter.

Redefining the perimeter is a lengthy task, and you will encounter several blind alleys. To assist you in avoiding the blind alleys, the IBM Method for Architecting Secure Solutions (MASS) can be a cornerstone of the considerations for redesigning and rethinking the perimeter and its definition. See “Additional reading” on page 18 for more information.

Perimeter protection starts on the host and is determined to a wide degree by the user. The user must be enabled, educated, and assisted to become responsible by making the

necessary tools, controls, and regulations available. This effort should play an integral part in the future of every organization in order to operate a secure and compliant IT environment.

Additional reading

If you want to read more about perimeter security, we suggest the IBM Method for Architecting Secure Solutions (MASS) that is being used by IBM Global Service employees in security architecture engagements. It can help you understand and categorize security-related problems and discussions in today's e-business-driven enterprise IT infrastructures. An introduction to MASS was originally posted in a special edition of the *IBM Systems Journal on End-to-End Security*, vol. 40, no. 3. You can also find this article in *Enterprise Security Architecture using IBM ISS Security Solutions*, SG24-7581. In addition to being a source of further technical details, this particular book is a comprehensive overview of all IBM ISS solutions.

The task of developing IT solutions that consistently and effectively apply security principles has many challenges, including the complexity of integrating the specified security functions within the several underlying component architectures found in computing systems, the difficulty of developing a comprehensive set of baseline requirements for security, and a lack of widely accepted security design methods. With the formalization of security evaluation criteria into an international standard known as *Common Criteria*, one of the barriers to a common approach for developing extensible IT security architectures has been lowered; however, more work remains. The MASS methodology uses a systematic approach for defining, modeling, and documenting security functions within a structured design process to facilitate greater trust in the operation of resulting IT solutions.

The team who wrote this IBM Redpaper publication

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Austin Center.



Axel Buecker is a Certified Consulting Software IT Specialist at the International Technical Support Organization, Austin Center. He writes extensively and teaches IBM classes worldwide on areas of Software Security Architecture and Network Computing Technologies. He holds a degree in computer science from the University of Bremen, Germany. He has 21 years of experience in a variety of areas related to Workstation and Systems Management, Network Computing, and e-business Solutions.

Before joining the ITSO in March 2000, Axel worked for IBM in Germany as a Senior IT Specialist in Software Security Architecture.



Per Andreas is an IT Architect in Network and Security Services working for IBM Denmark. He has 9 years of experience in the network and security field. He holds a bilingual bachelor degree from the Aarhus School of Business and an IT Degree from ITU in Copenhagen. His areas of expertise include network transformation and transition projects and firewall management for global clients.



Scott Paisley is a Principle Security Architect at IBM Internet Security Systems in Fairfax, Virginia. He has 21 years of experience in systems integration, computer networking, and computer security. He has worked at IBM Internet Security Systems for 9 years. Prior to joining IBM Internet Security Systems, he worked at the National Institute of Standards and Technology. There he focused on systems integration products, Web design, and systems administration design, and wrote programs for Internet technologies. He is a frequent speaker at leading industry events, such as Forbes CIO Forum, Forbes Risk Management, Interop New York, and InfoSecurity New York. He holds a bachelor's degree in computer science from the University of Maryland in Baltimore.

Thanks to the following people for their contributions to this project:

Nancy Crumpton
Editor
International Technical Support Organization, Austin Center

Archived

Archived

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-4397-00 was created or updated on November 2, 2009.



Send us your comments in one of the following ways:


- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo) ®
IBM®
Internet Scanner®

Internet Security Systems™
Proventia®
RealSecure®

Redbooks®
SiteProtector™
Virtual Patch®

The following terms are trademarks of other companies:

Java, JavaScript, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.