


Addressing Emerging Threats and Targeted Attacks with IBM Security Network Protection



Redguides
for Business Leaders

Paul Ashley
Chenta Lee
Craig Stabler



- Examination of evolving security threats including malicious applications and network activity
- Introduction of the IBM Security Network Protection system components and deployment architecture
- Presentation of a typical client scenario with a solution that meets the challenge



Executive overview

In networks today, organizations are faced with hundreds of new web and non-web applications that are available to their users. Social media applications, peer-to-peer file transfer applications, Voice over Internet Protocol (VoIP), web-based email, cloud data storage, and many others are all readily available. The ease and speed at which these new applications can be installed or simply accessed reduces the effectiveness of a perimeter-based security architecture and provides many new types of risks. These applications can be used by an attacker to obtain initial access into the organization and bypass any perimeter-based security.

This IBM® Redguide™ publication introduces the *IBM Security Network Protection* solution, which is a *next generation intrusion prevention system* (IPS) that extends the capabilities of traditional protocol-based IPSes by providing application visibility and control. By using IBM X-Force® Research And Development, this solution provides critical insight and control of all user activities by analyzing each connection to identify the web or non-web application in use and the action being taken. The IBM Security Network Protection solution can then decide to allow or block the connection, and can inspect even those connections that are encrypted by SSL. Additionally, the X-Force IP Reputation information can be used to understand whether sites that are accessed are hosting malware, are BotNet Command and Control servers (C&C servers), or are phishing sites, and other important information.

The IBM Security Network Protection can record connection information, including user and application context, and can use this information for local policy refinement, including bandwidth management. Alternatively, the connection information can be sent to a *Security Information and Event Manager* (SIEM) for security analysis and longer term storage.

The IBM Security Network Protection consolidation of the traditional IPS function, in combination with sophisticated user-based application control and IP Reputation, can provide an integrated security solution. This approach allows for faster deployment and simplification of the administration that is associated with the deployment of multiple products, reduces the cost of ownership and complexity, and provides for better return on investment (ROI).

The target audience for this publication is business leaders, decision makers, network managers, IT security managers, and IT and business consultants.

Introducing the current threat landscape

This section introduces some of the key security challenges that organizations are facing today. The threat landscape has changed dramatically in recent times. New and sophisticated targeted attacks that are designed to gain continuous access to critical information are increasing in severity and occurrence. Some examples are *advanced persistent threats* (APTs), *stealth bots*, *targeted application attacks*, and *designer malware*.

You can obtain an up-to-date understanding of the threat landscape by referring to the IBM X-Force Trend and Risk Report¹, which is provided four times a year. The X-Force Research and Development teams monitor the latest security threats, including software vulnerabilities and public exploitation, malware, spam, phishing, web-based threats, and cyber criminal activity. An updated summary of this research is included in each report.

The 2013 mid-year Trend and Risk Report showed the growing trend in cyber incidents. Numerous and documented attacks to organizations both small and large occurred throughout 2013, with the scope and frequency of data breaches continuing in an upward trajectory. Many of these attacks cost the organization millions of dollars in lost business and reputation.

The 2013 mid-year report grouped the incidents based on attack behaviors. Social media sites were used by attackers who exploited a user's trust in the site. From legitimate looking spam containing malicious links to fraudulent emails appearing to be from friends, social media sites were used to both obtain information about users and as an initial access point into organizations. Attacks against mobile devices increased dramatically, in particular against mobile devices running the Android operating system. Although fixes were released to fix these mobile vulnerabilities, only a few users actively patched their mobile devices to prevent attacks. These compromised mobile devices provide an access path into unprepared organizations. Distributed Denial of Service (DDoS) attacks also became more prevalent, but were used mostly to distract organizations from another attack. Another form of attack was against watering holes, which are those trusted sites from which users download applications or updates. This type of attack, which is known as *poisoning the watering hole*, is used to deliver malware to even the most technically savvy organizations. Finally, attackers reused tried and tested techniques, such as exploiting web application vulnerabilities, poor password management, and vulnerabilities in common tools (such as Java vulnerabilities in browsers).

One of the key changes is the loss of the security perimeter in organizations. We have moved to a world of interconnected devices and services. A lapse in policy enforcement at any point in the network can undermine the whole system. Organizations today are faced with new risks through the uncontrolled use of web and non-web applications. Employees everywhere want to use all of the same applications at work as they do in their private lives. This phenomenon, often called *consumerization of the enterprise*, creates risks. With relative ease, a user can access or download and install these applications for non-business purposes. The applications have the potential to expose the organization to new threats, consume bandwidth for non-business reasons, and reduce worker productivity. Additionally, the growth of mobile, and cloud-based computing increases the risk of security breaches that originate inside the organization. In many cases, the organization does not have visibility or control of application usage.

Firewalls cannot provide protection to an organization against this rogue application use. These applications typically communicate through fewer ports (often HTTP/HTTPS) and are actively deceptive (they use nonstandard ports and port and protocol hopping). Only through deep packet and session inspection can these applications be identified and controlled.

¹ <http://www.ibm.com/security/xforce/downloads.html>

Because of these new application-based threats, organizations are now asking themselves the following questions:

- ▶ How do I identify which applications my enterprise users are using?
- ▶ How much bandwidth is being used by each of these applications?
- ▶ Can I limit bandwidth consumption by application and user?
- ▶ Can I provide different users different application access?
- ▶ Can I control specific application features?
- ▶ How do I control and know which websites that users should be accessing?
- ▶ How do I stop data from being leaked from my organization?

Many organizations approach the problem by using a basic “block all” approach, that is, they deny all access to these applications. They quickly found that this approach is not a suitable one because many of these applications have a business purpose for some parts of the organization. For example, marketing is much more effective when it uses social media applications.

Other organizations used simple web-based URL filtering to solve the problem, but quickly found that many of the applications that are being used (for example, VoIP, peer to peer file sharing, and cloud data storage) are not web-based applications and cannot be controlled in this way. Additionally, many of the URL filtering solutions are “block all” and do not provide granularity to application actions. An example is web-based email. Control must be at a granular level, such as read, write, append file, and chat, and must be individually controlled.

Another challenge for organizations is the deployment of firewalls, intrusion prevention, URL filtering, vulnerability management, and other technologies that are not integrated. The usage of these poorly integrated solutions is creating “security sprawl”, lower overall levels of security, and raising cost and complexity for the organization.

The IBM Security Network Protection solution is designed to meet these new challenges. By incorporating best-of-breed intrusion prevention, application visibility and control, IP reputation, and SSL inspection into one solution, and integrating this solution with security intelligence, any organization can gain an increased level of security, visibility, and control.

Protecting your network against threats with protocol-aware detection

The core function of next generation Intrusion Prevention Systems is to perform deep packet inspection to detect threats and attacks from both internal and external sources. The IBM Security Network Protection solution functions by using the protocol-aware *Protocol Analysis Module* (PAM) engine. The PAM engine is developed by the IBM X-Force Research and Development team.

IBM X-Force Research and Development

The IBM X-Force team receives and analyzes billions of security events each day, which are sourced from the IBM Managed Security Services teams. The IBM X-Force team uses this information to develop new security protections for vulnerabilities, even in the case of no known use. In addition to this function, the IBM X-Force has close research ties with leading software vendors, which allow for early awareness of new vulnerabilities.

The security protections that are developed by the IBM X-Force team are then incorporated back into the PAM engine, which is the core component of all IBM network protection solutions. The fact that PAM is continuously updated by the IBM X-Force team as new vulnerabilities and uses are discovered provides unparalleled network security protection for IBM clients.

The various focus areas of the IBM X-Force Research and Development team are shown in Figure 1.




Figure 1 Focus areas of IBM X-Force

The Protocol Analysis Module

PAM is a modular framework that is designed by the IBM X-Force Research and Development team. The PAM software provides in-depth security and protection by analyzing every packet that traverses the IBM Security Network Protection solution.

Pattern matching versus protocol analysis

The PAM engine is a custom-built software engine that provides full protocol, content, and application awareness that goes beyond signature-based detection mechanisms. Many other IPS solutions use *pattern matching* to detect the signatures of specific techniques. This approach does not keep up to date with the evolution of current techniques and can also lead to a high rate of incorrect detection of attacks. Additionally, pattern matching engines require a unique signature for each usage technique. The PAM engine, with its *protocol analysis* approach, is vulnerability-based, not signature-based, meaning that it can evolve as threats and new attacks are discovered. Using advanced heuristics techniques, the PAM engine can detect evolving threats that are modified to avoid IPS signature-based products. These detection techniques enable the PAM engine to detect many zero-day attacks.

Because of its modular architecture, the PAM engine can inspect, maintain state, and, if necessary, block protocols throughout the different layers of network communication. PAM can understand and examine over 300 network and application layer protocols, 150 data file formats, 800 web applications, and 1600 actions. The engine can process low-level protocols, such as the Internet Protocol (IP), to detect and block attacks at this level (such as denial of service attacks). PAM can perform a deep analysis of data that is transferred by high-level protocols, such as Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), and Remote Procedure Calls (RPCs).

In addition to providing deep packet analysis at many different levels of network communication, PAM can inspect at high speed and with low latency. It is fully multi-threaded to enable it to monitor multiple connections in parallel. This configuration ensures that your network is secure while still maintaining the high levels of performance that are required in today's business environment.

Key components of the Protocol Analysis Module

The PAM engine combines many key features to provide a complete network security solution. The following core components are features of the PAM:

- ▶ IBM Virtual Patch® technology: Shields vulnerabilities from exploitation, independent of a software patch.
- ▶ Client-side application protection: Protects users against attacks that target applications that are used everyday, such as attacks that are embedded in Microsoft Office files, Adobe PDF files, and multimedia files.
- ▶ Web application protection: Provides protection against sophisticated web application attacks, such as SQL injection, cross-site scripting (XSS), and cross-site request forgery (CSRF).
- ▶ Advanced threat detection and prevention: Provides advanced intrusion prevention, including protection against potential zero-day attacks.
- ▶ Data security: Provides monitoring and identification of personally identifiable information (PII) and other confidential data over both unencrypted traffic and SSL encrypted traffic if SSL inspection is enabled.

Figure 2 depicts the key features of the PAM engine.



Figure 2 Core PAM components

Virtual Patch Technology

This module of the PAM engine detects and prevents attacks that are directed at known vulnerabilities, regardless of whether a software patch exists. In this way, vulnerabilities can be patched according to a known patch management process. In the meantime, systems are protected by PAM.

Client-side application protection

As users become aware of the risks of receiving executable style files (for example, .exe), malware writers have changed to target file types that are considered safe by users, including Adobe PDF, Microsoft Office, and other file types. The aim is to use weaknesses in the document processing systems to allow execution of system shell commands (hence the term, *shellcode*) to download malware or other types of attacks.

The X-Force Shellcode Heuristics technology provides powerful protection against these types of attacks. PAM includes heuristic-based decodes that detect shellcode in the most commonly used file and network protocols. All of these decodes detect payloads that are used by, but not limited to, Metasploit² tools and other well-known patterns that are used to attack multiple operating systems.

This module protects against attacks that are aimed at everyday client-based applications. As shown in Figure 3, the PAM module provides protection for vulnerabilities before they are discovered. In the diagram, the blue marks indicate when PAM provided the protection. The red marks indicate the vendor announcement, public disclosure of the vulnerability, and the occurrence of a zero-day attack.

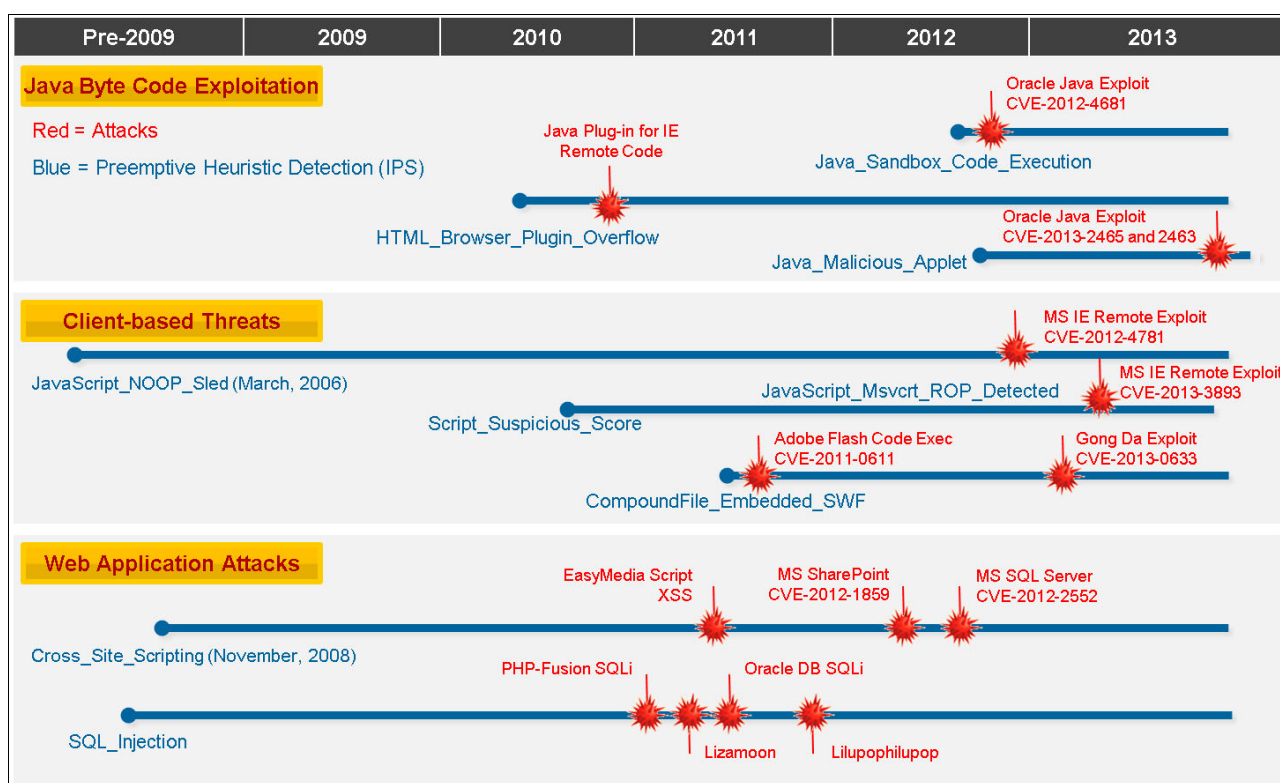


Figure 3 The pre-emptive PAM protection

² <http://www.metasploit.com>

Web Application Protection module

The *Web Application Protection* module provides protection for web applications against a multitude of web-based attacks, such as SQL injection, cross-site scripting (XSS), PHP file includes, and cross-site request forgery (CSRF). According to the IBM X-Force 2013 Mid-year Trend and Risk Report, over 1500 web application vulnerabilities were reported in the first half of 2013 alone, with more than half related to cross-site scripting vulnerabilities.

Instead of using a pattern matching approach, the IBM X-Force team developed a behavioral approach to identifying and blocking injection-related attempts to use web application vulnerabilities. A heuristic examination of the entire datastream sent to the web server makes evasion more difficult, resulting in fewer false negatives and false positives.

The X-Force Injection Logic Engine (ILE) helps preempt injection attacks by detecting unique patterns that are not usually seen in valid web requests. By applying scores for specific keywords and symbols and their resulting logical constructions, ILE can detect and then block SQL injection and other injection-related attacks without requiring new signature updates. Through its comprehensive heuristic understanding of SQL syntactic cues, the ILE helps protect systems in the following ways:

- ▶ Evaluating and scoring parameter URL query and POST data values
- ▶ Blocking requests that exceed the scoring threshold
- ▶ Flagging particular keyword combinations to identify the type of SQL injection that is occurring

Advanced Threat Detection and Prevention module

This module protects against attacks such as botnets, worms, and trojans. Using advanced heuristics techniques, the PAM engine can detect evolving threats that are modified to avoid IPS products. The detection techniques of PAM enable it to protect against vulnerabilities that are still unknown, thus providing protection against the zero-day attacks that are commonly used by APTs.³

Data Security module

The Data Security module provides the PAM engine with the ability to inspect many forms of data transfer, such as emails, HTTP, FTP, instant messaging, and Internet Relay Chat (IRC). Within these different protocols, the PAM engine can inspect various types of data, including Adobe PDF files, Microsoft Office files, and compressed files. This feature allows the PAM engine to search for customizable user strings to prevent the unauthorized transfer of PII, such as credit card numbers, names, dates, email addresses, and Social Security numbers. In addition to these predefined types, administrators can add their own customized searches that can detect certain sensitive organization information. Moreover, the Data Security Audit feature in PAM can help organizations meet compliance with a great deal of granularity, such as the usage of web email services and file sharing services.

³ <http://xforce.iss.net/xforce/xfdb/62643>

IP quarantine

IBM Security Network Protection can quarantine certain IP addresses when they match the Intrusion Prevention policy, IPS Event Filter policy, or the Advanced Threat Policy in response to Advanced Threat or IPS events that are detected by the appliance. An administrator can create customized IP quarantine objects to do granular control to specific events. In Figure 4 an administrator can specify the type of the IP quarantine object, the duration of the quarantine, and the rate limit of the quarantined IP address.

Add Quarantine

Name:
Customized_Quarantine_Object

Comment:
It is a customized quarantine object

Duration (seconds):
3600

Rate Limit (kbps):
100

Type:
Intrusion

☒ Intruder Address
☐ Intruder Port
☐ Victim Address
☐ Victim Port
☐ ICMP Type
☐ ICMP Code
☐ Issue ID

Save Configuration Cancel

Figure 4 Add a quarantine object

An administrator can view and drop the current active quarantine rules on the local management interface (LMI) and on IBM Security SiteProtector™ (SiteProtector) V3.1 and later. Moreover, administrators can also promote the quarantine rule to make it become an enterprise-wide rule if they found the event that is related to the quarantine impacts other segments. Figure 5 on page 9 shows the view of active quarantine rules on SiteProtector and how to promote them.

Enable	Applies To	Creation Date	Expiration /	Type	Source IP	Source Port	Target IP	Target Port	Target URI	VLAN	ICMP Type	IC
✓	XGS5100.demo@192.168.122.243	Jun 4, 2014 10:14 PM	Jun 4, 2014 11:14 PM	tcp	192.168.122.243		192.168.124.100			none		

Figure 5 Quarantine rules in SiteProtector

Gaining deep insights into network applications and users

In today's environment of APTs and ubiquitous web and social media applications, every organization needs deep insights into the usage patterns on their networks so that they can effectively use and secure these network assets. The IBM Security Network Protection solution provides this knowledge and security with granular application and user visibility and control.

The business cost of unauthorized social media use is significant and growing as social media applications become more pervasive. The exponential growth of social media is both good and bad for business. On one hand, use for non-business purposes reduces employee productivity and limits bandwidth availability for business applications. Additionally, it can provide new attack vectors into valuable corporate resources. However, common social media applications such as Twitter, Facebook, and YouTube are important marketing media. The IBM Security Network Protection solution allows visibility and control of over 800 applications, 1600 actions, and over 20 billion URLs. These functions are controllable at the organization, department, workgroup, or individual level.

Network visibility

Modern enterprises can have thousands or millions of connections that flow through their network at any point in time. In addition to implementing a deep packet inspection engine, organizations must be aware of what activity is occurring on their network. Continuous monitoring of the network connections on a network allows an enterprise to detect unusual patterns and abnormal behaviors.

In addition to detecting security threats, retaining detailed statistics of network connections can assist with compliance and regulatory mandates. If there is a security violation, network connection flow data can be used as part of the forensic analysis process. This security intelligence about the network of an organization is a critical component of the overall security solution for any modern IT system. The IBM Security Network Protection solution assists this process by storing and analyzing network connection flow data, providing complete network and security intelligence, for all organizations.

The IBM Security Network Protection solution can automatically collect network connection information about all network traffic. It then stores this information by using its internal storage for local on-box analysis by the LMI. Depending on the amount of network activity, the IBM Security Network Protection can store up to 30 days of network flow data. For more detailed analysis and long-term retention of data for compliance and regulatory purposes, IBM Security Network Protection can export this same network connection information to an external SIEM product by using Internet Protocol Flow Information Export (IPFIX) formatted data. This action is known as *off-box analysis*.

On-box analysis

The IBM Security Network Protection solution provides a number of network flow analysis options without using any external appliances or software. These analysis options are available through the LMI. The LMI is a web-based interface that is used for both configuration and for viewing the security analysis information.

A number of different views are available without any specialized configuration or setup as soon as the IBM Security Network Protection solution starts analyzing traffic. The following views are available with the IBM Security Network Protection solution:

- ▶ Traffic Details by Application
- ▶ Traffic Details by User
- ▶ Traffic Details by Web Category

In the *Traffic Details by Application* view, each of the network connections is analyzed by the PAM engine and categorized based on the application type. Applications that are detected are divided into two main types: non-web based client applications and web-based applications. Examples of non-web based applications include DNS, SSH, FTP, and Yahoo! Messenger. Examples of web-based applications include Gmail, Facebook, Twitter, and the New York Times website. The number of detected applications is continuously increasing as the IBM X-Force research team categorizes new applications and releases updated information to the IBM Security Network Protection solution.

The top applications can be displayed as both a pie chart and line chart to give a cumulative view of application bandwidth and time-based view of application bandwidth usage. This view is useful to see immediately any outlying applications that are using a large amount of bandwidth. An example of this view is shown in Figure 6.

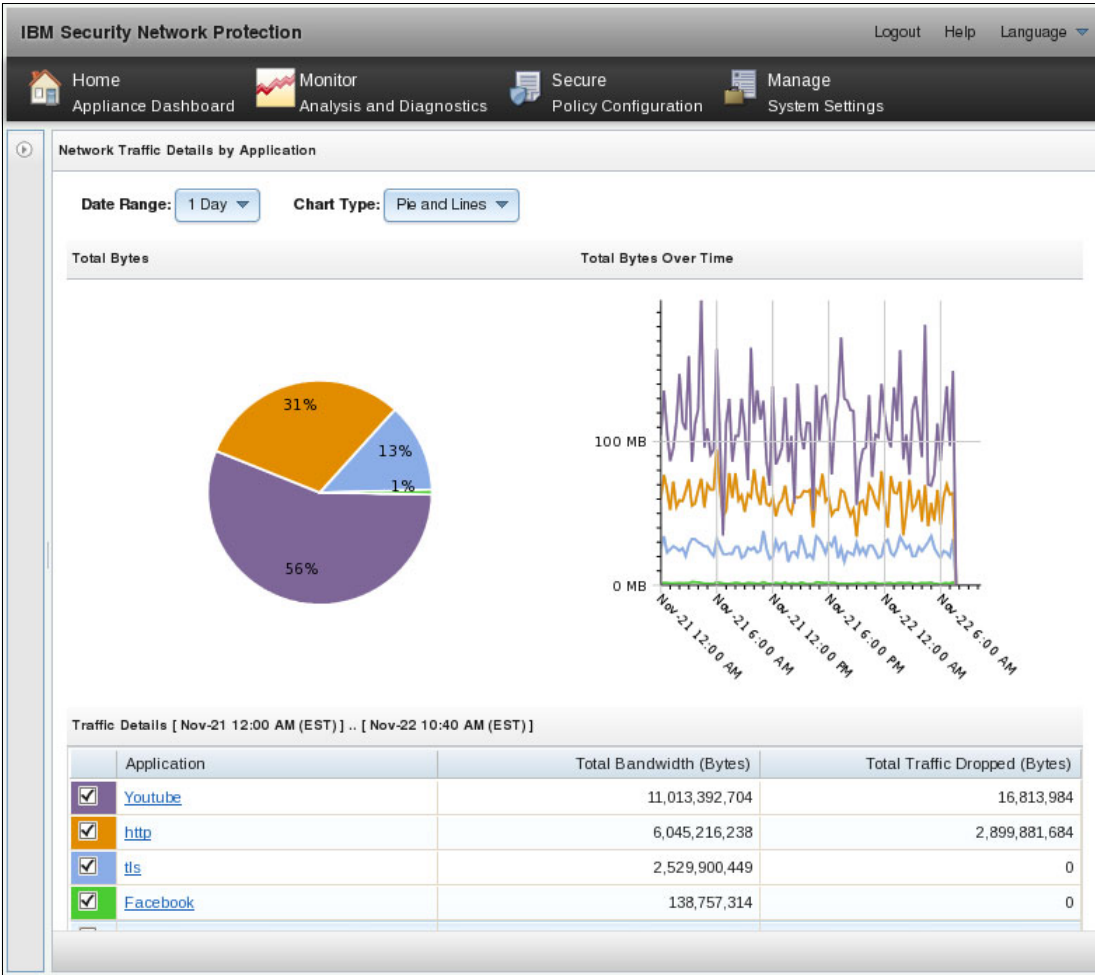


Figure 6 IBM Security Network Protection Traffic Details by Application view

Figure 7 shows the *Traffic Details by User* view, which displays all connections on a per user basis. The IBM Security Network Protection solution generates this view by correlating network traffic data with user information from an authentication server, such as a Microsoft Active Directory server.

The Traffic Details by User view enables a security administrator to see quickly whether any users are transferring an unusually high amount of traffic. This usage might indicate a number of different situations, some of which might be the result of a security breach. An unusually high amount of traffic might indicate the following security breaches:

- ▶ Malware infected the notebook of a user.
- ▶ A command and control botnet took over the notebook of a user and is sending spam.
- ▶ A user has non-compliant software that is installed, such as peer to peer file transfer, which is transferring large amounts of data.

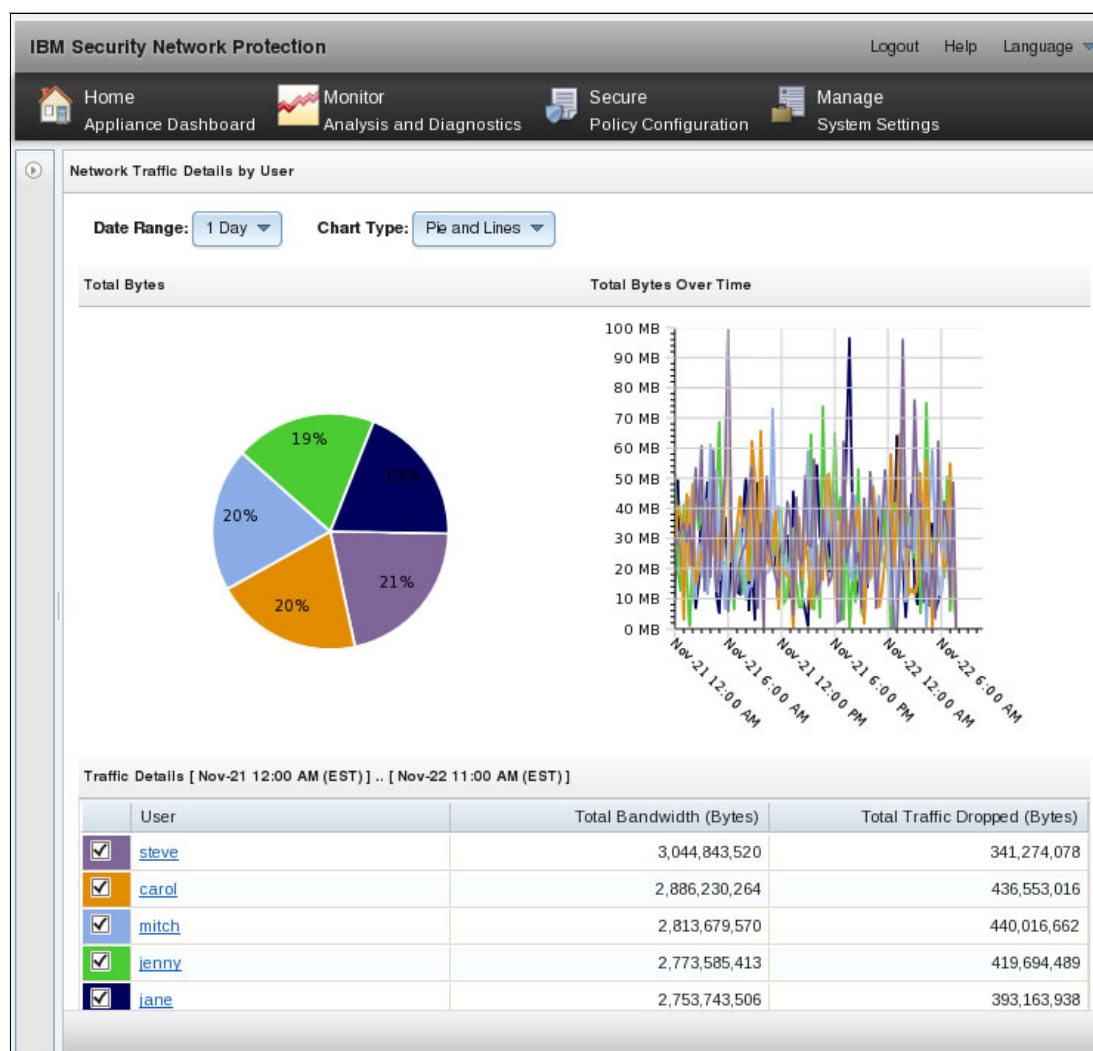


Figure 7 IBM Security Network Protection Traffic Details by User view

The other analysis view that the IBM Security Network Protection can show is the *Web Traffic by Category* view, which is shown in Figure 8 on page 13.

The Web Traffic by Category view shows the amount of network traffic that is connecting to websites of a particular category. The list of categories includes certain topics that are often required by the enterprise to be blocked for all users because they are not related to the area of business of the organization. Examples often include gambling, illegal activities, and software piracy. By using this view, a system administrator or manager can immediately see whether there are large amounts of traffic to particular web categories and act on this network activity.

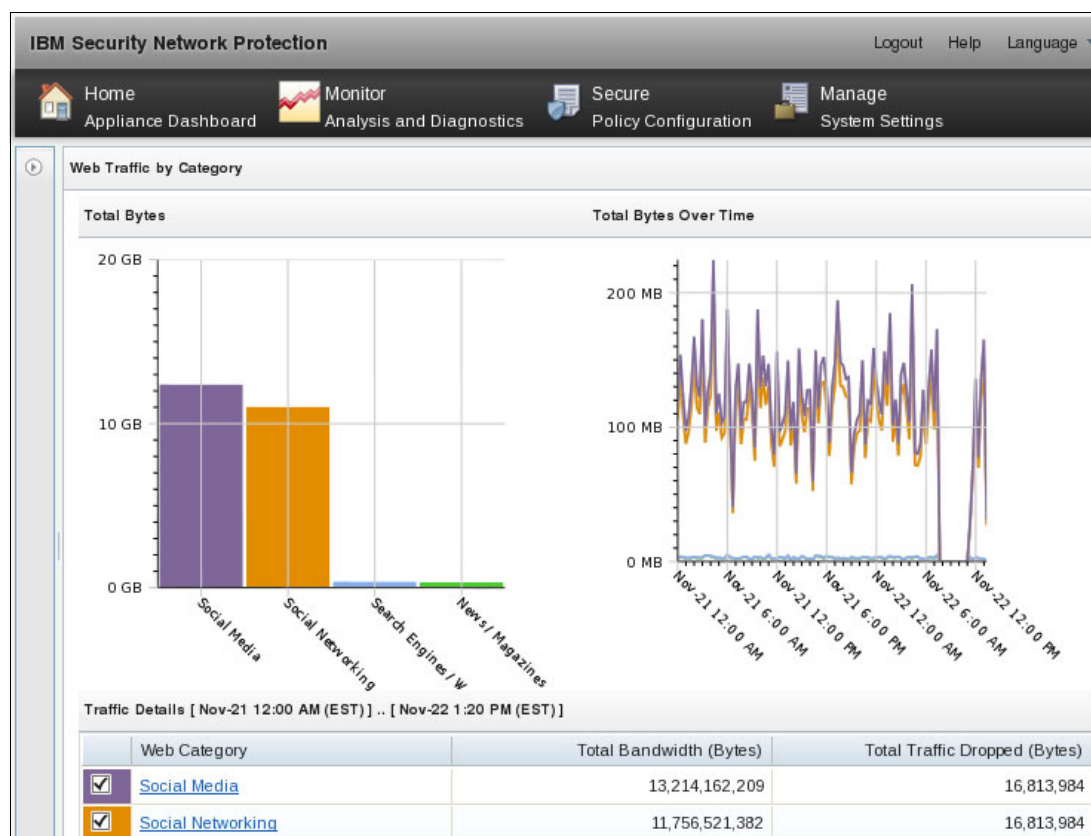


Figure 8 IBM Security Network Protection Web Traffic by Category view

Off-box analysis

The Cisco Systems Netflow Services protocol⁴ was developed to collect IP traffic information. This protocol contains information about each network flow, such as source and destination IP address, source and destination port number, and the amount of data that is transferred. The latest version of the Netflow protocol, Version 10, is standardized by the Internet Engineering Task Force (IETF) and is renamed *IP Flow Information Export (IPFIX)*.⁵

In addition to its on-box analysis capabilities, the IBM Security Network Protection solution can export the same network flow data to any external product that can understand the IPFIX network flow data format. This function is useful for many security reasons. If the external product has archival features, it can keep long-term network flow data for archival purposes if a potential breach is discovered weeks, months, or even years after the event. Further deep analysis can also be done by external products, which have the extra capability to correlate network connection data with other data sources, such as server log information. Combining multiple sources of security information provides even more detail in to the network of an organization, which allows for a higher level of security.

⁴ Cisco Systems Netflow Services Export Version 9: <http://www.ietf.org/rfc/rfc3954.txt>

⁵ <http://datatracker.ietf.org/wg/ipfix/charter/>

In addition to supporting standard IPFIX, the IBM Security Network Protection can add extra fields about user and application information to the network flow data by using IPFIX extensions. This feature allows detailed analysis by user and application at the SIEM. The IBM QRadar® SIEM product⁶ is an example of an SIEM that supports IPFIX with extensions.

Figure 9 shows a QRadar summary page with data grouped by IBM Security Network Protection applications.

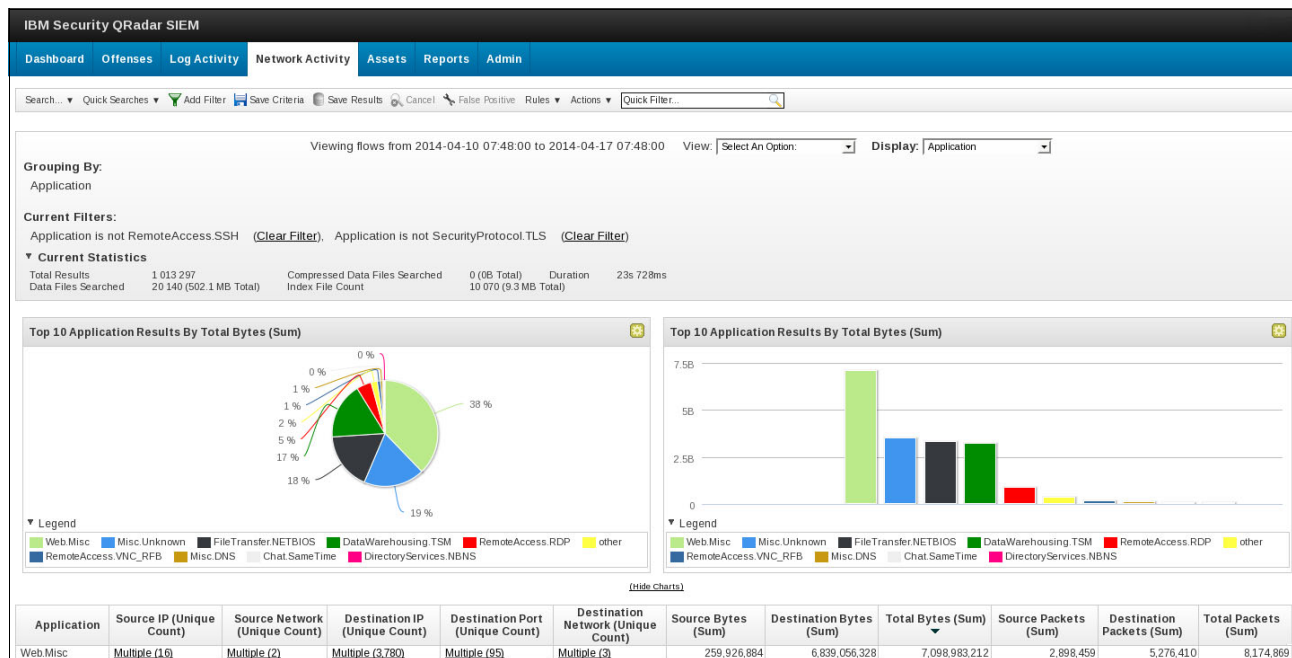


Figure 9 QRadar SIEM traffic analysis view by application

Obtaining granular control of your network applications and users

The IBM Security Network Protection solution offers visibility into user and application network traffic usage so that security administrators can design Network Access Policies (NAP) to enforce corporate security requirements.

Controlling applications usage within your network

Figure 10 on page 15 shows a user on the corporate intranet accessing both web and non-web applications through the IBM Security Network Protection. This solution can recognize over 1000 applications and actions, regardless of network address, port, and protocol. The IBM Security Network Protection solution can even detect many applications that are designed to specifically evade detection by first-generation IPSes.

⁶ QRadar SIEM: <http://q1labs.com/products/qradar-siem.aspx>

The IBM Security Network Protection solution consolidates the functions that are normally found in a URL filtering proxy. With a URL category database that contains over 18 billion URLs, visibility and control of employee Internet activity is also provided. This database allows an organization to retire their existing URL proxy solutions. Instead, the database uses the web application visibility and control function that is built into the IBM Security Network Protection solution.

Figure 10 shows the access to Internet-based applications.

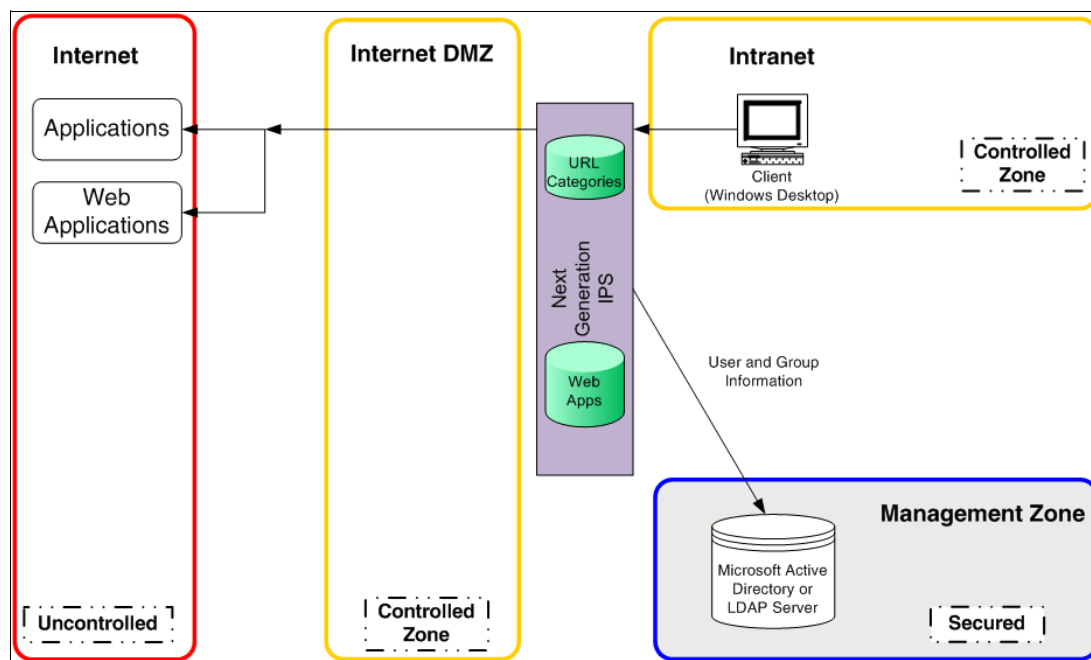


Figure 10 Controlling access to internet-based applications

The security administrator is armed with the deep knowledge of network usage that is gained from the network visibility features of the IBM Security Network Protection solution. Therefore, the administrator can design policies to allow or disallow access to any of the supported applications or URL categories. The administrator has ultimate control over both bandwidth usage and the risk of application-originated attacks.

An IBM Security Network Protection policy extends traditional firewall rules by allowing applications or groups of applications to be matched as either the source or destination. This policy gives the security administrator control over traffic that flows to or from a particular application.

As an example of granular control, an IBM Security Network Protection policy can be written to allow users the ability to view Facebook but not post updates to Facebook. That is, this solution can control specific application actions rather than only the complete application.

Controlling user activity within your network

In addition to controlling web and non-web applications, organizations need visibility and control over who is using their network resources. The IBM Security Network Protection policy allows security administrators to specify users or user groups, thus granting or denying access on a per user or per group of users basis. As shown in Figure 10, the IBM Security Network Protection solution can query corporate directory servers to obtain user and group information, for use in an IBM Security Network Protection policy.

This solution supports Lightweight Directory Access Protocol (LDAP) and Microsoft Active Directory (AD), thus minimizing the number of user names and passwords that each employee must manage. This support provides significant cost savings because a large percentage of internal help desk calls is related to user credentials. To reduce operational costs even further, the IBM Security Network Protection solution can be configured to automatically log in users (passive authentication) when they log in to their Microsoft Active Directory domain. Additionally, for sites without access to corporate directory servers, IBM Security Network Protection provides a local user and group database that allows user and group information to be entered directly.

As an example, an IBM Security Network Protection policy can be written to allow users in the marketing department unrestricted access to social media sites. Access to these sites for other users can be restricted to lunch time and outside of normal business hours.

IBM Security Network Protection network access policy

IBM Security Network Protection network access control is exercised through policy enforcement. Your corporate network access policy is built by using a series of NAP rules. An IBM Security Network Protection policy has a similar style to a typical firewall rule set, which is intuitive and easily understood by your security administrators.

The IBM Security Network Protection policy is highly flexible, giving security administrators the ability to create powerful network access policies quickly and efficiently. Figure 11 shows a simple NAP consisting of seven rules.













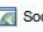
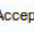






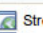




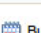


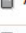
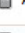

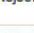








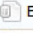









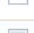





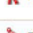
	Order	Enable	Source	Destination	Application	Action	Response	Inspection	Schedule
	1		 Any	 Any	 DHCP-DNS	 Accept		 Default IPS	
	2		 Marketing Dept	 Any	 Social Networking	 Accept	 Event Log	 Default IPS	
	3		 Any	 Any	 Streaming Media  Social Networking	 Reject	 Event Log	 Default IPS	 Business Hours
	4		 Any	 Any	 Prohibited URLs	 Reject	 Event Log	 Default IPS	
	5		 Any	 Any	 C&C Server	 Reject	 Capture Packet  Event Log	 Default IPS	
	6		 Cybercrime countries	 Any	 Any	 Reject		 Default IPS	
	7		 Any	 Any	 Any	 Accept		 Default IPS	

Figure 11 Network Access Policy

The IBM Security Network Protection solution enforces these rules by evaluating them in priority order from the first rule to the last. Each rule consists of the following fields:

- *Source* and *Destination* fields are used to identify the entity sending or receiving network traffic. IBM Security Network Protection can use the traditional network attribute to identify the network traffic, for example, IP range and IP subnet can be used in Source and Destination fields. Furthermore, IBM Security Network Protection can use user identity to distinguish the network traffic, which can bind with a local user database or a remote directory server, including the LDAP and AD. This advance feature enables the capability to design the role-base security policy. Furthermore, to do geographic location enforcement within the enterprise, IBM Security Network Protection can leverage the geographic information that is associated with IP addresses to restrict accesses to or from certain countries by constructing a Geo Location Object in NAP. Section “IP Reputation database” on page 20 provides more details about Geo Location Object.

Figure 12 on page 17 shows the source and destination types in the LMI.

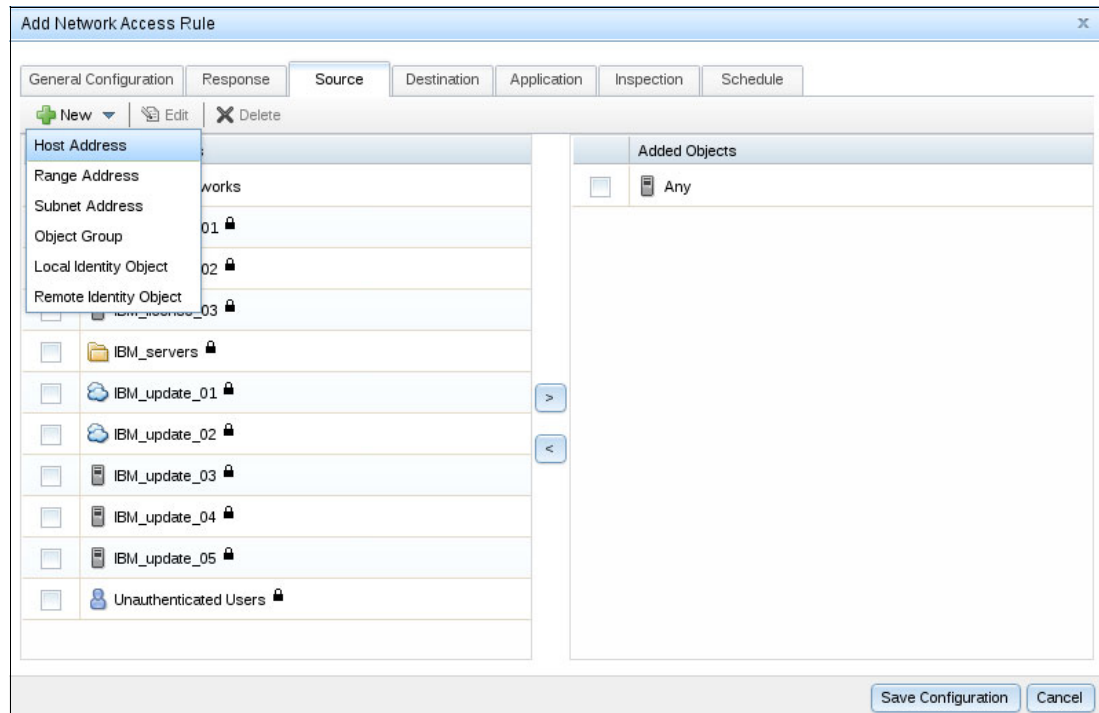


Figure 12 Network Access Policy - source and destination types

- Application fields are used to identify the type of the network traffic. IBM Security Network Protection classifies applications into four categories:
 - Web Application objects

Let an administrator control access to web-based applications and to control the user behavior on the website.
 - Non-web Application objects

Let an administrator control the access to the applications that use the network to exchange information.
 - URL Categories and List objects

Let an administrator control access to certain categories of websites.
 - Domain Certificate Categories and List objects

Let an administrator restrict the types of domain certificates that users can access.
 - IP Reputation object

Lets an administrator control or monitor traffic based on the classification of source IP or destination IP with a configurable threshold. The detailed configuration of the IP Reputation object can be found in the “IP Reputation database” on page 20.

Figure 13 shows these different application types in the LMI.

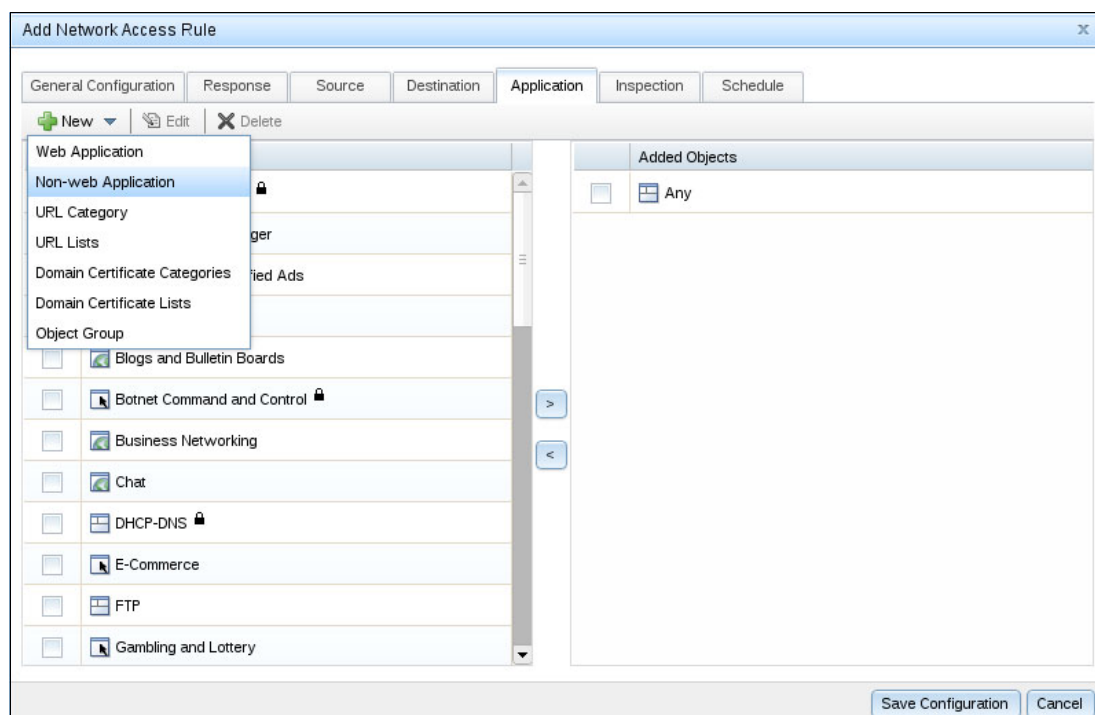


Figure 13 Network Access Policy - application types

- ▶ The *Action* field defines what the IBM Security Network Protection does with network packets that match this rule.
- ▶ The *Response* field defines how the IBM Security Network Protection communicates the fact that this rule is matched. Security administrators can put one or more response objects in the Response field. There are five types of response object:
 - Local Log: Sends an event to a local log. A security administrator can view the events in a NAP event console in LMI and in the SiteProtector console.
 - Remote syslog: Sends an event to a remote syslog server by using LEEF format. Security administrator can enable the QRadar format to send the event to QRadar SIEM.
 - SNMP: Generates an SNMP trap when the rule is matched.
 - Email: Sends an email to the appropriate security administrator when the rule is matched.
 - Packet Capture: Records the traffic that is associated with the matched NAP rule. A security administrator can analyze recorded packets to understand the attacks and the violations to the security policies. A security administrator can manage and download the captured files in the LMI.
- ▶ The *Inspection* field specifies the IPS policy to enforce.
- ▶ The *Schedule* field allows rule enforcement to be on a time of day basis.

The following scenarios explain some of the examples that are used in Figure 11 on page 16:

- ▶ *Rule 2* allows members of the user group “Marketing Dept” access to social networking applications at any time. Rule 2 records these accesses in the event log.
- ▶ *Rule 3* is lower priority than Rule 2, and is evaluated only for traffic that does not match Rule 2. Although Rule 2 allows the marketing department access to social networking applications always, Rule 3 denies all users access to streaming media and social networking applications during business hours.
- ▶ *Rule 4* blocks all traffic to a set of prohibited URL categories.
- ▶ *Rule 5* blocks and records all traffic to or from a command-and-control server if the score of the source IP or destination IP that are in the IP reputation database exceeds the user-defined threshold.
- ▶ *Rule 6* blocks all traffic that originates from a country that has active cybercrime.
- ▶ *Rule 7* is the default rule, that is, it is the last rule in the policy. It allows all traffic that does not match the previous rules.

SSL inspection

The IBM Security Network Protection solution provides inspection of both outbound and inbound SSL connections, which reduces the need to implement a separate SSL inspection appliance. Enabling SSL inspection reduces overall throughput. For the current performance metrics and testing criteria, see the IBM Security Network Protection data sheet (the URL to the data sheet can be found in “Other resources for more information” on page 36).

Outbound SSL inspection

The IBM Security Network Protection solution can inspect outbound SSL traffic. The on-box SSL inspection reduces the administrative impact compared to other solutions by eliminating the need to maintain an additional SSL inspection appliance. Outbound SSL traffic is the type of traffic that is initiated from internal clients to remote web applications. It is critical to inspect outbound SSL traffic because web applications increasingly rely on SSL encryption to secure the connection between application and user. It is also a requirement to have this feature to do granular control of web applications.

When the outbound SSL inspection policy is enabled, IBM Security Network Protection establishes a man-in-the-middle (MitM) session between the client and the destination server when the client initiates an SSL connection. After the MitM session is created, the appliance uses the self-signed certificate (also known as a proxy certificate) to establish the SSL connection to the client, therefore allowing it to encrypt the SSL packets from the client. After analyzing the decrypted data, the appliance then re-encrypts the packet and sends it to the server. Clients see a warning message for the certificate when connecting to the remote web application. The administrator can export the IBM Security Network Protection CA certificate and then import it to a client’s environment to suppress the warning messages.

One of the key features of the IBM Security Network Protection SSL inspection is the ability to be selective about which SSL sessions are decrypted. Figure 14 shows a sample policy. In this case, any outbound SSL sessions to privacy-sensitive websites, or to IBM update servers, are not inspected. The policy-based SSL inspection is an important feature of the IBM Security Network Protection solution. It allows an organization to bypass decryption of SSL sessions that are protected by law, or where the organization wants to protect the privacy and security of its users.

The screenshot shows the 'IBM Security Network Protection' web interface. At the top, there are navigation tabs: 'Home' (Appliance Dashboard), 'Monitor' (Analysis and Diagnostics), 'Secure' (Policy Configuration), and 'Manage' (System Settings). The 'Secure' tab is active. Below the tabs, the 'SSL Decryption Policy' section is displayed. It includes a toolbar with 'New', 'Edit', and 'Delete' buttons. A table lists four policy rules. The first three rules are set to 'Ignore' and the fourth is set to 'Inspect'. The table has columns for Order, Enable, Source, Destination, Domain, Action, and Comment.

Order	Enable	Source	Destination	Domain	Action	Comment
1	<input type="checkbox"/>	Any	Any	Privacy-sensitive Information	Ignore	Do not inspect encrypted traffic to websites that contain privacy-sensitive information.
2	<input type="checkbox"/>	Any	IBM_servers	Any	Ignore	Do not inspect IBM update server traffic that passes through the protection interfaces.
3	<input type="checkbox"/>	Marketing	Any	Privacy-sensitive Information	Ignore	
4	<input type="checkbox"/>	Any	Any	Any	Inspect	Inspect all encrypted traffic that is not ignored by other rules.

At the bottom of the table, it shows '1 - 4 of 4 items' and pagination controls for 10, 25, 50, 100, and 200 items per page.

Figure 14 SSL inspection policy

Inbound SSL inspection

IBM Security Network Protection can also inspect inbound SSL traffic. Inbound SSL traffic is encrypted traffic that is initiated by a remote client to a server within the enterprise. For example, an HTTPS connection between a client and the internal web server hosting a company’s official website is an inbound SSL connection. There is no need to create a MitM session for an inbound SSL connection because an administrator can store the certificate and private key for the protected server in the appliance to enable this feature. The appliance can use the private key and server certificate to decrypt the SSL packets then analyze them.

IP Reputation database

IBM Security Network Protection integrates an IP Reputation database from IBM X-Force Research and Development. The database provides the reputation and geographic location information for both source and destination IP addresses. IBM Security Network Protection can therefore use the additional information to provide a more accurate analysis of the on-going network traffic in the corporate network. For example, if a protected machine accesses an IP address that is categorized as a botnet command-and-control server, it is likely that the protected machine is compromised with botnet client malware.

With the geographic information in the IP Reputation database, IBM Security Network Protection can base decisions on the location of the IP address to identify the originator as being from a country that is rated as high for malicious activities, such as spam and malware spreading.

Moreover, the security administrator can design a more secure Network Access Policy that is based on the IP reputation information in the traffic. For example, IBM Security Network Protection can block the traffic originating from an IP address that is classified as an anonymous proxy server by using an IP Reputation Object in the NAP to stop the attackers hiding behind the proxy servers. In addition to geographic location information, an organization can use the Geo Location Object in the NAP to block traffic that is sent from a certain country that has no business relationship with the organization.

The data in the IP Reputation database is based on the analysis of massive amounts of worldwide network activity data that IBM X-Force Research and Development continuously collects.

As an example of use of the IP Reputation data, in the IPS event detail, an administrator can see the IP Geo Location and the IP reputation categories in both source and target IP address. In Figure 15, the IPS event detail shows that the source IP is rated as high (86%) for an anonymous proxy, which means that an attacker might be using this proxy to attack the corporate web server anonymously.

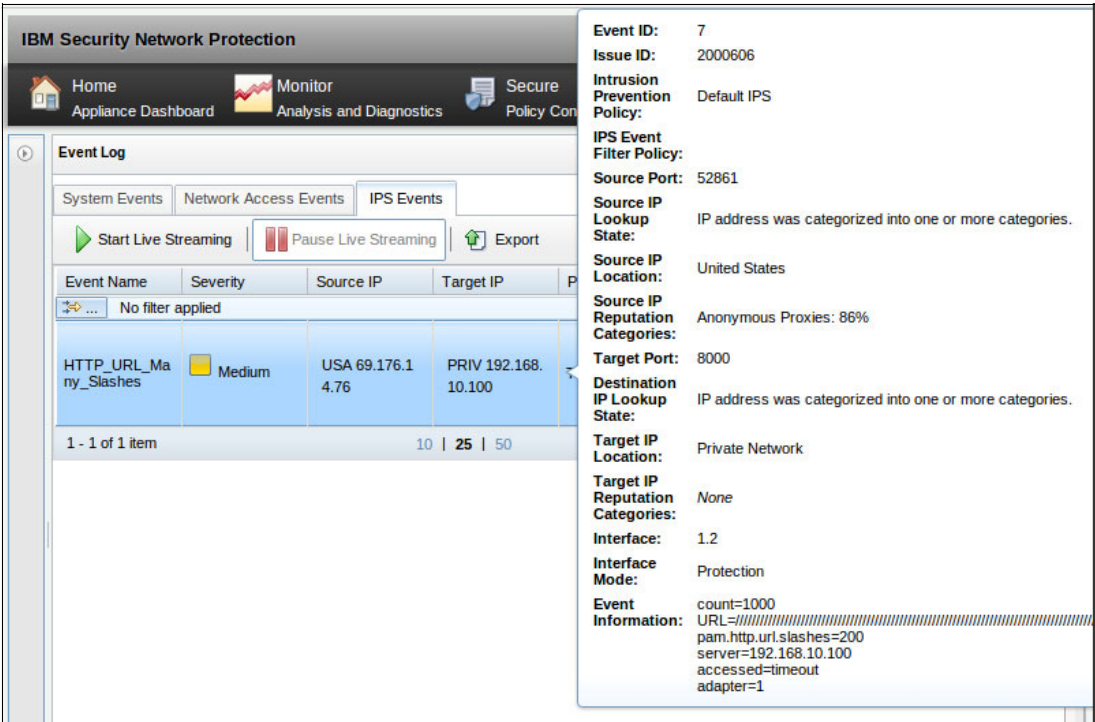


Figure 15 An IPS event identifying access from an anonymous proxy

Figure 16 shows an IP Reputation Object that matches traffic if the source IP or destination IP is classified as a command-and-control server and the threshold to the reputation score is set to 50.

Edit IP Reputation Category Object

Name:

Comment:

This object matches traffic if the source IP or destination IP is classified as command-and-control server and the threshold to the reputation score is set to 50.

Included	Categories	Thresholds
<input type="checkbox"/>	Spam	<input type="text" value="50"/>
<input type="checkbox"/>	Anonymous Proxies	<input type="text" value="50"/>
<input type="checkbox"/>	Scanning IPs	<input type="text" value="50"/>
<input type="checkbox"/>	Dynamic IPs	<input type="text" value="50"/>
<input type="checkbox"/>	Malware	<input type="text" value="50"/>
<input checked="" type="checkbox"/>	Botnet Command and Control Server	<input type="text" value="50"/>

1 - 6 of 6 items
10 | **25** | 50 | 100 | 200

Figure 16 Defining a command & control object

Figure 17 on page 23 shows a Geo Location Object that matches traffic to or from the countries in North America.

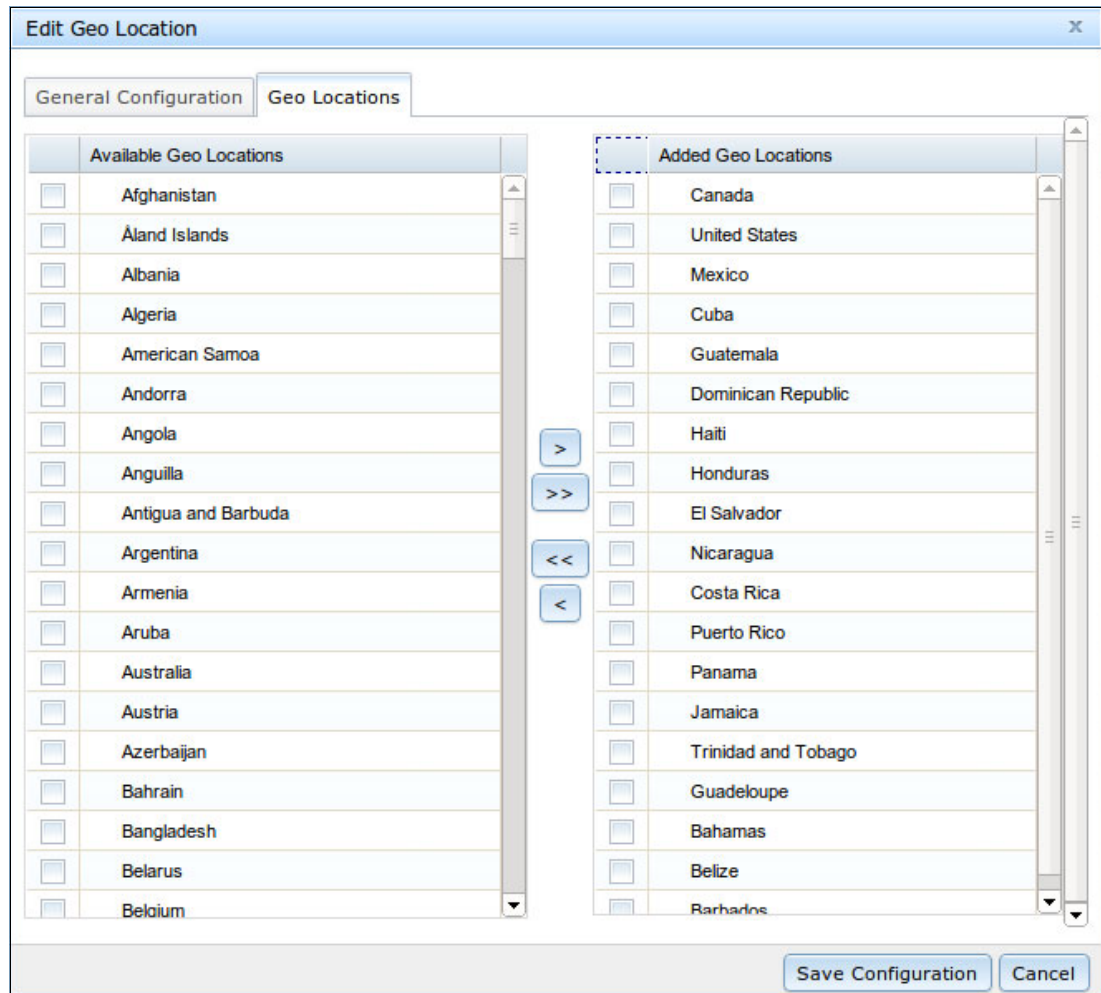


Figure 17 Defining an object for North America countries

How IBM Security Network Protection deals with spear phishing attacks

IBM Security Network Protection supports a number of new features that enable it to detect and block user-targeted attacks.

This section clarifies which features are required to enable IBM Security Network Protection to protect users from *spear phishing* attacks. Traditionally, these attacks took the form of an email containing a rogue URL; now, these rogue URLs can be delivered to the target user through social media, such as Twitter, Facebook, or LinkedIn.

To prevent users from connecting to the target URL that is provided in a spear phishing email, you must enable blocking of the following URL categories within the NAP:

- ▶ Spam
- ▶ Phishing
- ▶ Malware
- ▶ C&C Server

In addition, to ensure that https://* (and http://*) connections are blocked to the above categories, the SSL inspection feature also must be enabled.

The ability to block the above connection types at a network level is useful for organizations that allow, for example, external consultants who might be using a device that is infected with malware to connect to organization's network.

Advanced Threat Protection

Figure 18 shows the Advanced Threat Protection (ATP) system. This feature allows the IBM Security Network Protection solution to receive alerts from IBM and third-party software components and to act on these alerts. It does this by using the IPS Quarantine function that allows both IP addresses and URLs to be blocked for a specific period.

Examining Figure 18, the IBM or third-party solution sends an alert to the IBM Security Network Protection solution. An example alert is that a host in the network was compromised by an attacker (perhaps identified by inappropriate traffic). The alert is parsed by the ATP system and the ATP agent's policy is consulted. Depending on the alert and the policy, both an event and a quarantine might be created. Taking the host compromise as an example, the host IP address might be added to the quarantine list so that communication to and from that host is blocked.

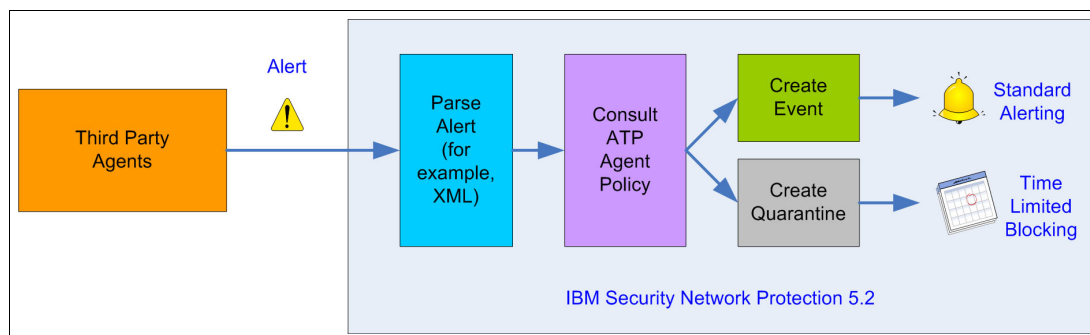


Figure 18 Advanced Threat Protection system

FireEye WebMPS Integration

Figure 19 on page 25 shows an integration of the FireEye WebMPS⁷ and IBM Security Network Protection ATP system. In Figure 19 on page 25, the FireEye appliance is connected to a span port on a switch so that it can monitor all traffic from and to the user network segment. The FireEye WebMPS can capture files in transit, for example, downloaded from a website, and using its sandboxing system attempt to identify and profile any files that contain malware. It does this by having the malware run in its sandbox. The FireEye appliance can identify calls to C&C servers from internal hosts and provides other functions.

⁷ To learn more about the FireEye offerings visit <http://www.fireeye.com/>.

The diagram illustrates the architecture of IBM Security Network Protection 5.2. At the top, a Firewall is connected to the Internet. Below the Firewall is the IBM Security Network Protection 5.2 appliance. To the right, the IBM Security Site Protector 3.1 is shown, which is connected to the Network Protection appliance via a Management Network. The Management Network is represented by a switch icon. Below the Network Protection appliance is another switch icon, which is connected to the FireEye Web MPS 6.2. The FireEye Web MPS 6.2 is connected to the User Network Segment, which consists of three laptops. The diagram shows the flow of traffic and management connections between these components.

25

Figure 20 shows the IBM Security Network Protection ATP for the FireEye WebMPS appliance. It shows that for different alert types and priorities that the action around events and quarantines can be defined.

The screenshot displays the 'Advanced Threat Policy' configuration page. On the left, a sidebar titled 'ATP Objects (Drag/Drop on Rules)' contains a tree view with 'Alert' and 'Quarantine' categories. The 'Alert' category includes 'SNMP', 'Email', 'Log', and 'Remote Syslog'. The 'Quarantine' category lists various threat types like 'ATP-Compromise-Host', 'ATP-Exposure-Endpoint', 'ATP-Exposure-Host', 'ATP-Intrusion-DDOS', 'ATP-Intrusion-Intruder', 'ATP-Intrusion-Origin', 'ATP-Intrusion-Trojan', 'ATP-Intrusion-Worm', 'ATP-Malware-Intruder', 'ATP-Malware-URI', 'ATP-Malware-Victim', 'ATP-Reputation-Host', and 'ATP-Reputation-URL'. The main area shows a table with columns: 'Enable', 'Agent Type', 'Alert Type', 'Alert Severity', 'Current Responses', 'Default Protection', and 'User Overridden'. The table lists 10 rules for 'FireEye WebMPS 6.2' agents, covering 'Compromise' and 'Malware' alert types with 'High', 'Medium', and 'Low' severities. Each rule specifies default protection actions like 'ATP-Compromise-Host', 'ATP-Malware-Intruder', and 'ATP-Malware-URI' with associated event logs.

Enable	Agent Type	Alert Type	Alert Severity	Current Responses	Default Protection	User Overridden
<input type="checkbox"/>	FireEye WebMPS 6.2	Compromise	High	ATP-Compromise-Host Event Log	ATP-Compromise-Host	No
<input type="checkbox"/>	FireEye WebMPS 6.2	Compromise	Medium	ATP-Compromise-Host Event Log	ATP-Compromise-Host	No
<input type="checkbox"/>	FireEye WebMPS 6.2	Compromise	Low	Event Log		No
<input type="checkbox"/>	FireEye WebMPS 6.2	Compromise	Unknown	ATP-Compromise-Host Event Log	ATP-Compromise-Host	No
<input type="checkbox"/>	FireEye WebMPS 6.2	Malware	High	ATP-Malware-Intruder ATP-Malware-URI Event Log	ATP-Malware-Intruder ATP-Malware-URI	No
<input type="checkbox"/>	FireEye WebMPS 6.2	Malware	Medium	ATP-Malware-Intruder ATP-Malware-URI Event Log	ATP-Malware-Intruder ATP-Malware-URI	No
<input type="checkbox"/>	FireEye WebMPS 6.2	Malware	Low	Event Log		No
<input type="checkbox"/>	FireEye WebMPS 6.2	Malware	Unknown	ATP-Malware-Intruder ATP-Malware-URI Event Log	ATP-Malware-Intruder ATP-Malware-URI	No
<input type="checkbox"/>	FireEye WebMPS 6.2	Reputation	High	ATP-Reputation-Host ATP-Reputation-URL Event Log	ATP-Reputation-Host ATP-Reputation-URL	No

Figure 20 FireEye WebMPS Advanced Threat Policy

QRadar right-click integration

Figure 21 on page 27 shows a screen capture of the QRadar right-click integration with IBM Security Network Protection. In the figure, an administrator right-clicks a source or destination IP or port and an alert is configured to be sent to the IBM Security Network Protection appliance. This function allows an administrator to act quickly on an alert that is raised in QRadar. For example, QRadar might determine a host is compromised and raise an alert. The administrator can then act on this alert by blocking traffic to or from that host.

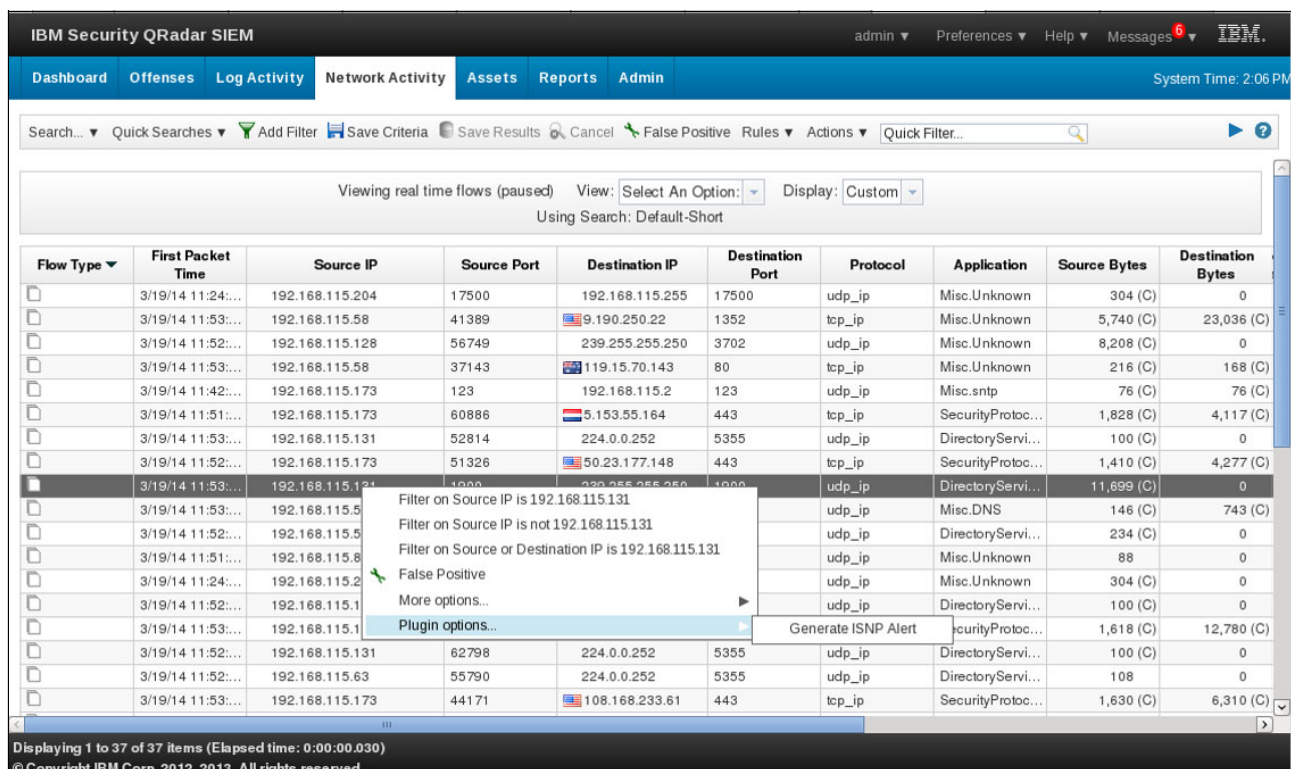


Figure 21 Right-click integration on QRadar to send an alert to the ATP system

There are a number of different scenarios that are supported:

- ▶ **Compromise:** If the source IP is right-clicked, this IP address is sent to the IBM Security Network Protection solution. This function can be used when the host is infected with malware.
- ▶ **Reputation:** If the destination IP is right-clicked, this IP address is sent to the IBM Security Network Protection solution. The destination IP represents a malicious server, such as a C&C server or one that hosts malware.
- ▶ **Intrusion:** If a source port is right-clicked, this IP address and port combination is sent to the IBM Security Network Protection solution. This action can result from that client system attacking a server.
- ▶ **Exposure:** If the destination port is right-clicked, this IP address and port combination is sent to the IBM Security Network Protection solution. This function can be used in a case where the service has a vulnerability.

Centralized management using SiteProtector

The SiteProtector system provides centralized management of a security policy to allow an administrator to design the policy and deploy it to multiple agents (for example, IBM Security Network Protection) that are managed by SiteProtector. Moreover, it also collects the events that are sent from these agents and analyzes them in a unified event console. SiteProtector allows efficient security policy and event management for even the largest enterprise deployment, with hundreds of agents deployed. The IBM Security Network Protection solution can be managed from SiteProtector V3.0 and later releases.

The SiteProtector system consists of the following components:

- ▶ The SiteProtector Console:
 - Monitors the status of all agents.
 - Creates, saves, and prints analysis views.
 - Monitors and filters event alerts.
 - Manages agent policies.
 - Generates and schedules reports.
 - Configures SiteProtector Database maintenance options.
 - Creates SiteProtector tickets.
- ▶ The X-Press Update Server: Retrieves the update packages from the central IBM update repository and deploys them to agents, which allows agents to sit in an isolated network and still receive the latest updates. An administrator can also download the update packages from the IBM update repository and then manually deploy them to the X-Press Update Server; as a result, the entire SiteProtector system can be isolated from the public network.
- ▶ SiteProtector Database: Stores the agent data, including the events and statistics. SiteProtector can then use the data for analysis and reporting.
- ▶ Event Collector: Receives the data that is sent from agents or Agent Manager in real time and stores them in the SiteProtector Database for analysis.
- ▶ Agent Manager: Mediates the agent data transmission between the Event Collector and agents. Therefore, there might be multiple Agent Managers in one SiteProtector system to share the workload and provide high availability.

Figure 22 shows a diagram of SiteProtector management of the IBM Security Network Protection policy. The right pane of the figure shows the policies that are managed by SiteProtector. The left pane is used to logically group appliances with the same policy configuration so that their policies can be managed in one place.

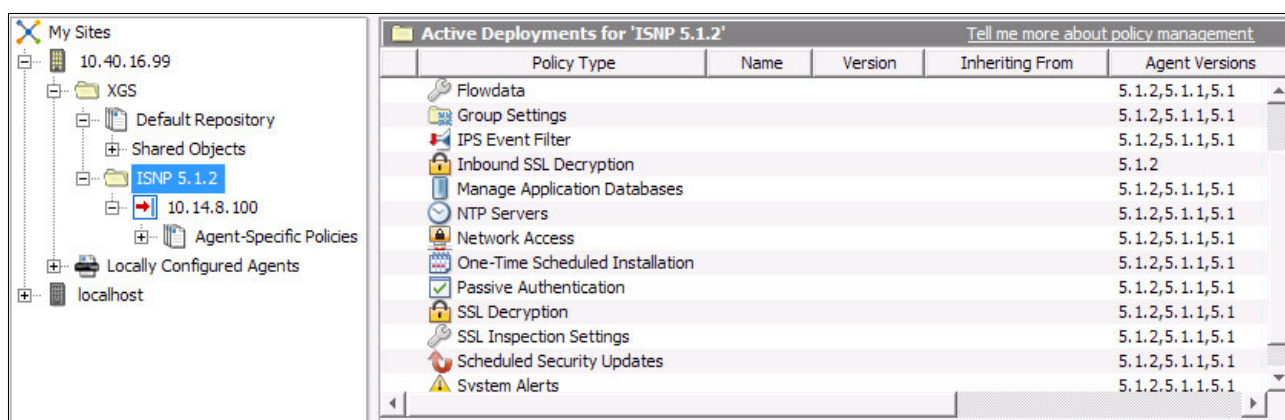


Figure 22 Centrally managing the IBM Security Network Protection solution from SiteProtector

Flexible hardware configuration

There are three hardware models that are available for IBM Security Network Protection. All are 1U in depth. The first is the XGS 5100, which is shown in Figure 23. It provides an on-board serial port, two on-board management ports, and four on-board copper 1G network interfaces. Additionally, it provides two bays for network interface modules (NIMs). In Figure 23 and Figure 24, the NIMs that are shown each have eight copper 1 Gbps interfaces.



Figure 23 IBM Security Network Protection - XGS 5100 model

Figure 24 shows the XGS 4100 model. It has the same on-board network interfaces, and one bay for a network interface module.



Figure 24 IBM Security Network Protection - XGS 4100 model

Figure 25 shows the XGS 3100 model. It has the same on-board network interfaces, but no bay for a network interface module.



Figure 25 IBM Security Network Protection - XGS 3100 model

Table 1 provides more details about the different options for the NIMs. A range of copper and fiber NIMs with and without bypass are available.

Table 1 IBM Security Network Protection Network Interface Modules

Item	XGS 3100	XGS 4100	XGS 5100
Number of supported NIMs	N/A	1	2
NIMs with integrated bypass	N/A	8 x 1GbE TX (100/1000) 4 x 1GbE SX 4 x 1GbE LX 2 x 10GbE SR 2 x 10GbE LR	8 x 1GbE TX (100/1000) 4 x 1GbE SX 4 x 1GbE LX 2 x 10GbE SR 2 x 10GbE LR
NIMs with external bypass optional	N/A	4 x 1GbE SFP 2 x 10GbE SFP+	4 x 1GbE SFP 2 x 10GbE SFP+

Flexible performance licensing

Table 2 shows some of the performance characteristics of the XGS 3100, XGS 4100, and XGS 5100 models. In row two, the Flexible Performance Levels (FPLs) for each appliance type are shown. These FPLs allow an organization to license the appliance at a lower level than its rated capacity, and update to a higher FPL when the traffic in the organization increases.

Table 2 IBM Security Network Protection Flexible Performance Characteristics

Item	XGS 3100	XGS 4100	XGS 5100
Inspected throughput	Up to 600 Mbps	Up to 1 Gbps	Up to 5 Gbps
Flexible Performance Levels (FPLs)	Up to 300 Mbps (FPL 1) Up to 600 Mbps (FPL 2)	Up to 500 Mbps (FPL 1) Up to 1 Gbps (FPL 2)	Up to 2 Gbps (FPL 1) Up to 3.5 Gbps (FPL 2) Up to 5 Gbps (FPL 3)
Inspected throughput (with SSL)	Up to 100 Mbps	Up to 400 Mbps	Up to 2.5 Gbps
Maximum throughput	3.5 Gbps	3.5 Gbps	6.5 Gbps
Average latency	<250 μ s	<150 μ s	<150 μ s
Connections per second	10,000	15,000	50,000
Concurrent sessions (maximum rated)	150,000	300,000	2,000,000

Implementing intrusion prevention and application control by using IBM Security Network Protection

An enterprise scenario is used as an example of using the IBM Security Network Protection solution to protect incoming and outgoing connections to the Internet.

Scenario description

In this scenario, an organization has a head office campus with more than 1000 network users (employees and contractors). The corporate data center houses all of the applications and databases, including web servers, web application servers, and core business applications.

The Chief Information Security Officer (CISO) of the enterprise identified the need to protect the enterprise from external threats by doing the following tasks:

- Control external user access and protect the web applications of the enterprise from external requests.
- Provide access control and monitoring of internal network users to Internet web and non-web applications.

Figure 26 shows the various zones in the enterprise. Secure access to the enterprise's web applications by Internet users is through a web reverse proxy. *Intranet* users are allowed access to both Internet-based web and non-web applications. The *Production Zone* houses the critical production applications and databases. The *Management Zone* includes user directories and system and security management applications.

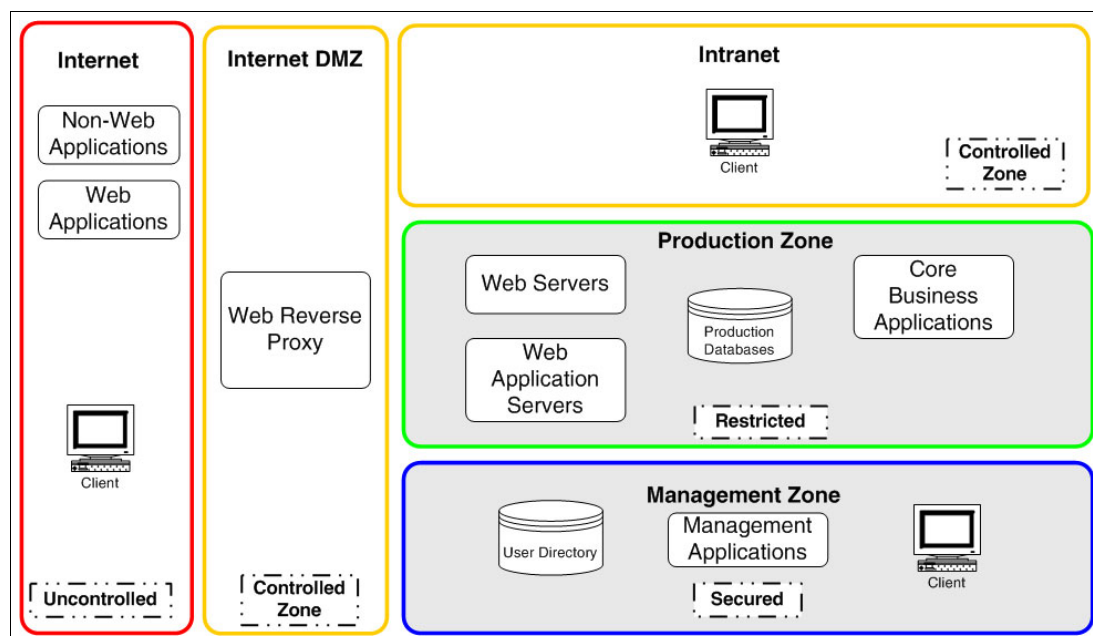


Figure 26 Enterprise network zones

The CISO determined that primarily the perimeter of the organization needs greater security protection because of the sophisticated nature of Internet threats. For example, the enterprise acknowledged the increasing variety of web-based application attacks (SQL injection and cross-site scripting being two common examples) and also the growing incidents worldwide of Internet sites and applications that are used to provide malware to unsuspecting user computers. The CISO recognizes the large number of social media applications, file sharing, online video, VoIP, WebMail, and other applications that are being accessed from internal network users. Additionally there is a recognition of a growing percentage of the user's Internet application access being over SSL. The CISO wants to understand better and control this application usage.

The CISO wants to consolidate some of the existing security products. Currently, the organization has web URL filtering proxy appliances, intrusion prevention appliances, SSL inspection appliances, and web application firewall (WAF) appliances. Each of these appliances has a different administration interface and they are not integrated. The cost of managing these disparate devices is becoming prohibitive.

Solution

To fulfill the CISO's requirements, the IBM Security Network Protection solution can be deployed at the perimeter. This solution is shown in Figure 27, with an IBM Security Network Protection appliance used on both sides of the DMZ to provide control of incoming and outgoing requests.

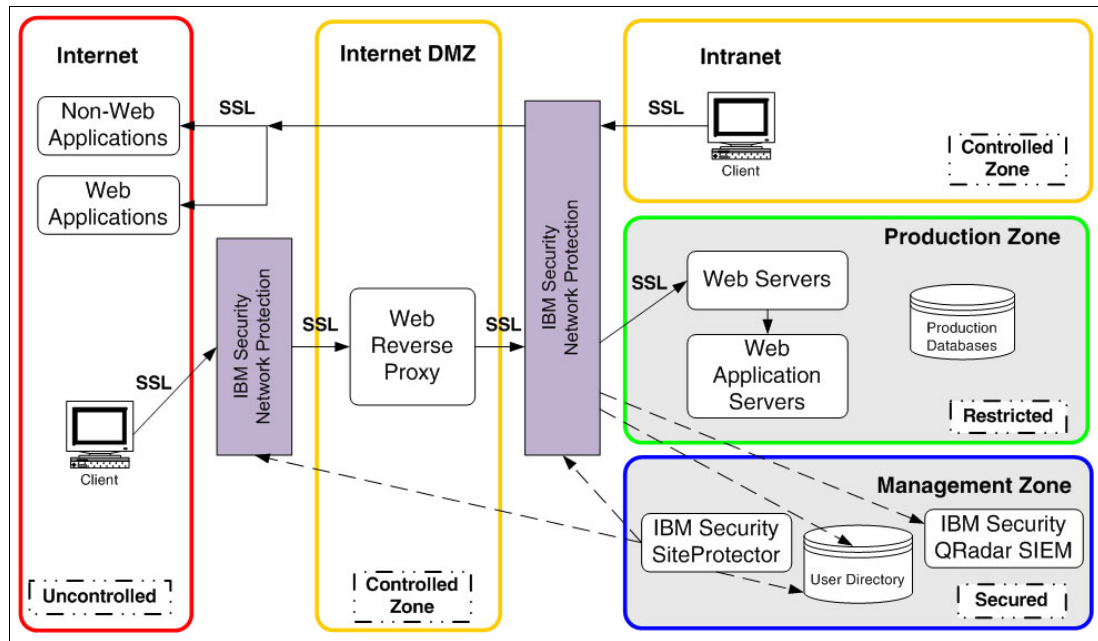


Figure 27 Using the IBM Security Network Protection solution for protection of the perimeter

- ▶ The *Web Application Protection* module of the IBM Security Network Protection protects the internal web applications from various web-based attacks, such as injection attacks.
- ▶ The *Virtual Patch* module of the IBM Security Network Protection protects against attacks that are targeted at the web server and web application servers.
- ▶ The *User Based Application Control* module of the IBM Security Network Protection solution allows users access to web and non-web applications to be identified and controlled for each user.
- ▶ The *SSL inspection* modules of the IBM Security Network Protection solution allow both incoming and outgoing SSL connections to be selectively decrypted for inspection.
- ▶ The *IBM Security SiteProtector* management application allows centralized policy management of the IBM Security Network Protection appliances.
- ▶ The *IBM Security QRadar SIEM* allows for capturing of both security events and connection flow data for analysis.

The flow data export feature of the IBM Security Network Protection solution allows exports of all flow data (in IPFIX) to the enterprise SIEM (QRadar) for behavioral analysis and reporting, including sending standard IPFIX data and IPFIX extensions for user and application data.

The organization decided to retire their existing web URL filtering proxy, intrusion prevention, SSL inspection, and WAF appliances. These appliances are no longer required because all of these functions are included in the IBM Security Network Protection solution.

Managing IBM Security Network Protection policy

A security administrator that works within the Management Zone manages the IBM Security Network Protection policy by using the SiteProtector management system. This management includes a policy that governs user-based application access control, SSL inspection, and intrusion prevention rules.

The IBM Security Network Protection intrusion prevention policy is set up by using the IBM X-Force default policy, also known as the *X-Force Virtual Patch Policy*. The IBM X-Force policy is a set of default security events that are preconfigured to protect against the most serious threats. Additionally, the policy is configured to use the WAF components to protect against web application attacks.

The IBM Security Network Protection NAP rules are then configured to provide individual and group-based application access rules. The IBM Security Network Protection needs access to the user directory (in this case, Microsoft Active Directory) so that it can obtain the configured users and groups for use in the NAP. Additionally, SiteProtector is used to manage the SSL inspection and IPS policies.

Figure 28 shows policy management only for the IBM Security Network Protection that is deployed on the inside of the DMZ. However, similar policy management is used for the IBM Security Network Protection that is deployed on the outside of the DMZ. In this case, this IBM Security Network Protection is used in IPS mode only, and is not used to enforce user-based application control.

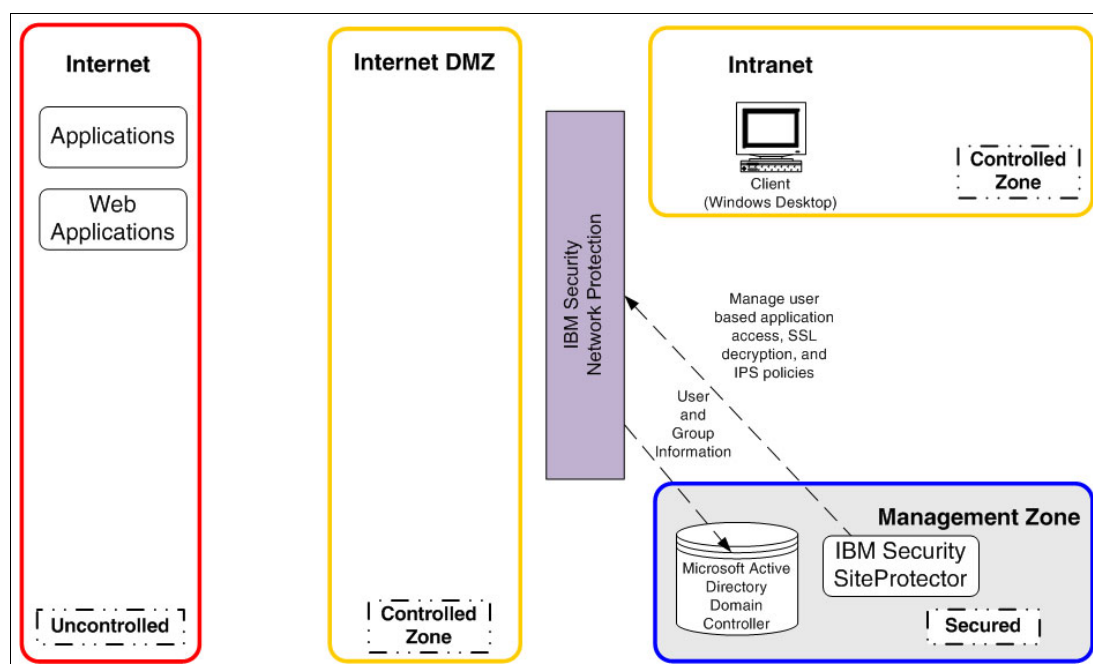


Figure 28 Using the IBM Security Network Protection policy management

User authentication

The requirement of the CISO is to provide application identification and control at the individual user level. The need is to provide different access for employees and contractors, and different groups of employees. Therefore, the IBM Security Network Protection appliance must be able to identify the user's traffic that comes across the wire. However, the CISO does not want to burden internal users with another login.

To accommodate this requirement, the Identity Management team of the enterprise chose the Passive Authentication deployment model with the IBM Security Network Protection solution. In this case, the IBM Security Network Protection Logon Event Scanner is installed on the Active Directory Domain Controller, as shown in Figure 29.

When a user logs in to Active Directory through their Windows workstation, the event is recorded and forwarded to the IBM Security Network Protection. A session is created for the user that identifies their IP address and Active Directory group memberships (obtained from the directory). From this point, the IBM Security Network Protection solution can determine for the individual users what their allowed access is. Additionally, the SSL inspection can be specified based on user and application.

Figure 29 shows an example of passive authentication of users by using Microsoft Active Directory.

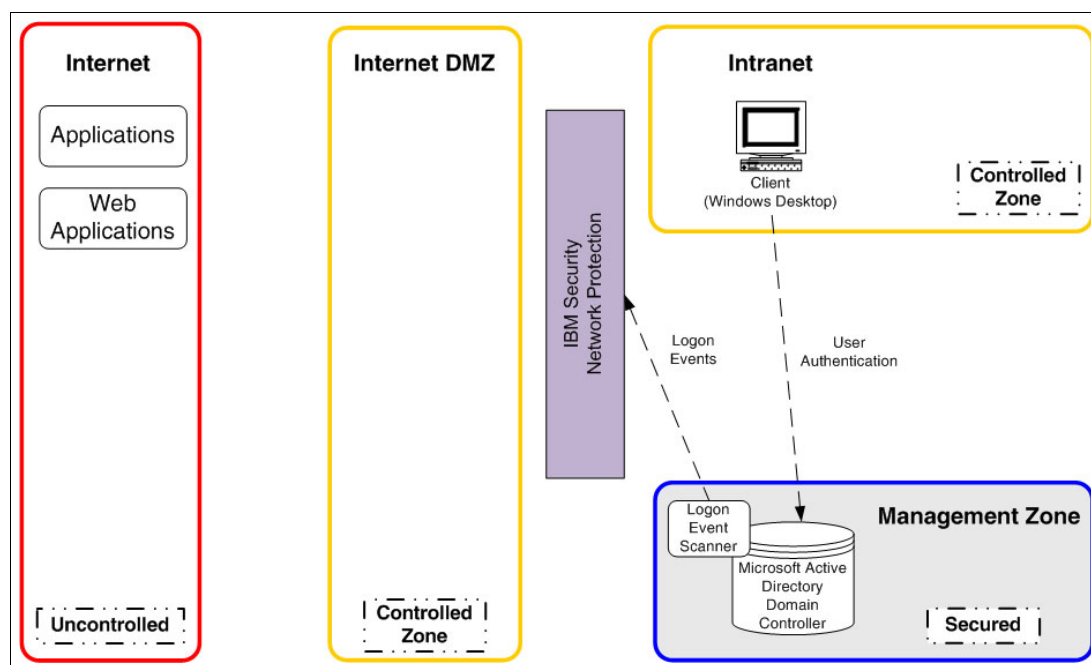


Figure 29 Using the IBM Security Network Protection passive user authentication

Recording flow and event data for user access

The CISO also is required to analyze in some detail access by users to Internet-based web and non-web applications. This analysis includes capturing all flow and event data, providing behavioral analysis on the data, and being able to characterize the IP reputation of the accessed sites. The IBM Security QRadar SIEM product is deployed in the management zone to collect both flow and event data from the IBM Security Network Protection. The QRadar also collects similar information from other security products, such as firewalls.

There is an advantage to the CISO using QRadar in this deployment. QRadar allows the enterprise to detect unusual events, drill down to understand every flow for a user, and be able to determine the IP reputation of those sites that are being accessed. Additionally, QRadar provides extensive reporting capabilities. Using the IBM Security Network Protection solution and QRadar together provides an incredible amount of security intelligence to the CISO.

Figure 30 on page 35 shows the process of recording Flow Data and Event Data at the SIEM.

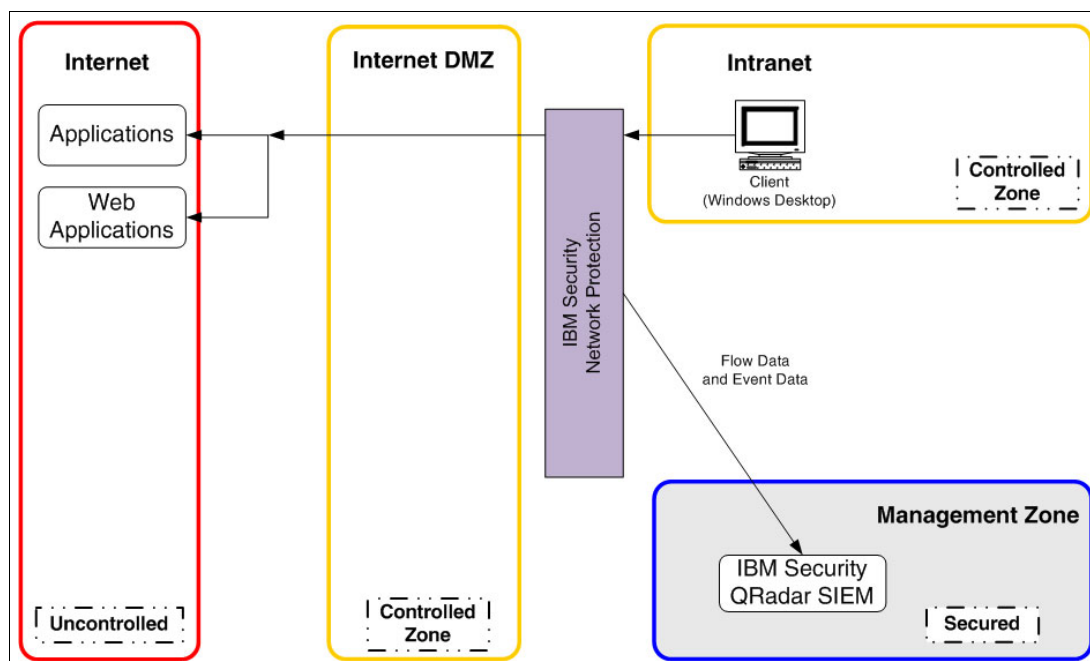


Figure 30 Recording flow data and event data at the SIEM

Conclusions

With the emergence of new and sophisticated threats, the CISO of the organization was required to provide an improved perimeter-based solution. The CISO's requirements included protecting the organization from Internet-based attacks, providing protection for the organization's Internet-facing web applications, and implementing granular user and application visibility and control. Most important for the CISO was being able to provide the application visibility and control even to SSL-encrypted sessions.

The organization selected IBM Security Network Protection in combination with IBM Security QRadar SIEM to provide these functions. The CISO now can view all traffic flows with views by user, application, and URL category. The CISO now can implement a network access policy to limit access to applications by individual user or group of users, even in SSL-encrypted connections. The CISO also can use the IBM Security QRadar SIEM functions to drill down on individual flows and respond quickly to any incidents.

An additional benefit for the CISO is that disparate security products, including IPS, Web Filter, WAF, and SSL inspection appliances can converge into one appliance, which saves on considerable on-going management and infrastructure costs.

Summary

The usage of web and non-web applications within the networks of organizations is widespread and is growing rapidly. It is becoming increasingly difficult to determine which applications are allowed and which must be blocked, as the distinction between business and non-business usage becomes more blurred. In addition, it is now commonplace for certain users and groups to require access to some applications for business purposes, while others do not. Combine these issues with an increasing number of vulnerabilities in IT software and an ever-increasing number of attacks from the Internet, and the goal of securing the network of a modern organization is exceedingly difficult.

The IBM Security Network Protection solution addresses these issues. This solution can perform deep packet inspection to protect against the latest and most complicated threats by using the PAM engine, which is backed by the IBM X-Force Research and Development team. The IBM Security Network Protection solution delivers deep insight into the bandwidth usage of the organization, which is categorized by application and by user to show who and what is using the network. The ability to enforce granular control of network usage through policy rules gives visibility and control to the security administrators and IT managers, which enable them to simply and efficiently secure their organization's network.

Other resources for more information

The following publications are useful as further sources of information:

- ▶ *Stopping Internet Threats Before They Affect Your Business by Using the IBM Security Network Intrusion Prevention System*, REDP-4683
- ▶ *Network Intrusion Prevention Design Guide: Using IBM Security Network IPS*, SG24-7979
- ▶ *IBM Security Solutions Architecture for Network, Server and Endpoint*, SG24-7581
- ▶ The IBM Security Network Protection data sheet is available at this website:
http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=SP&infotype=PM&appname=SWGGE_WG_WG_USEN&htmlfid=WGD03017USEN&attachment=WGD03017USEN.PDF
- ▶ The main IBM product web page is at the following website:
<http://www.ibm.com/software/tivoli/products/security-network-protection/>
- ▶ The official IBM Security Network Protection product documentation is at this website:
http://pic.dhe.ibm.com/infocenter/sprotect/v2r8m0/index.jsp?topic=%2Fcom.ibm.alps.doc%2Fconcepts%2Falps_intro_page.htm

Authors

This guide was produced by a team of specialists from around the world working at the IBM International Technical Support Organization (ITSO).

Paul Ashley is a Senior Technical Staff Member for the IBM Security Systems Division. He has 24 years of IT experience, with the last 20 years focusing on Information Security. His experience includes working in Identity and Access Management, Privacy Management, service-oriented architecture (SOA) security, mobile security, and Advanced Threat Protection. He has worked with IBM clients in Asia, US, Europe, and the Middle East. Paul holds a PhD in Information Security from the Queensland University of Technology and is an IBM Master Inventor. Paul is a member of the IBM Academy of Technology.

Chenta Lee is an Advisory Software Engineer with IBM Security Systems Division. His expertise includes emerging cloud technologies, with five years of experience in cloud security products, and experience in software-defined networking, virtualization, and advanced threat protection. Chenta is a member of the development team of IBM Security Network Protection. He currently focuses on network security in the cloud.

Craig Stabler leads IBM worldwide security sales enablement on Threat Protection. Craig has 25 years of experience in various international roles at Spider Systems, Shiva, Nortel, Internet Security Systems, and IBM. Craig completed CISSP and CCNA certifications and has co-authored the IBM Redbooks® series on network, endpoint, and server security. Craig has a Master of Engineering degree from Heriot-Watt University, Edinburgh, Scotland.

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

Matthew Elsner, Brian Fitch, Nathan Frith, Paul Griswold, Craig Knapik, Bob Mullins, Jack Smith, Ron Williams

IBM Software Group, Security Systems

Thanks to the authors of the previous editions of this guide.

Authors of the first edition, *Providing Next Generation Intrusion Prevention Functionality by Using the IBM Security Network Protection System*, REDP-4826, published in February 2013, were:

Paul Ashley, Greg McCane, Andrew Sallaway

IBM Australia

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

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 grant 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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-4826-01, was created or updated on July 16, 2014.




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:

IBM®
QRadar®
Redbooks®

Redguide™
Redbooks (logo) ®
SiteProtector™

Virtual Patch®
X-Force®

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

QRadar, and the Q1 logo are trademarks or registered trademarks of Q1 Labs, an IBM Company.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

FireEye is a trademark of FireEye, Inc.

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