



INTELLIGENT LAYERED SECURITY ARCHITECTURE
Better Security for the Growing Business

NOVEMBER 2005

Intelligent Layered Security Architecture: Better Security for the Growing Enterprise

Why Intelligent Layered Security?

The security landscape grows more complex and dynamic each day. A parade of newer technologies, such as instant messaging, wireless networks, and advanced Web services, are continually being deployed in businesses, presenting further opportunities for hackers to exploit. Security threats have increased in sophistication, frequency, and complexity as we have seen hacking become increasingly driven by fraud and organized crime. Today, virtually all attentive system administrators understand that the traditional stateful packet inspecting firewall is insufficient protection when working alone.

Many vendors today offer Unified Threat Management (UTM) appliances, which incorporate multiple security functions including firewall, VPN, spam filtering, antivirus, and intrusion prevention. These functions typically work independently, and do not integrate in a way that enables you to leverage information about one layer to make other layers more effective (Figure 1). Configuration can be complex, and logging information from different functions can be inconsistent. This translates into greater complexity, leading to a higher likelihood of misconfiguration, and ultimately, poorer security. Moreover, these systems are not designed with extensibility in mind; thus they can't rapidly evolve or extend defenses as new threats appear.

Traditional Layered Security




Figure 1: Unified Threat Management solutions created from separate security functions.

Businesses need an integrated security solution that by design protects against today's threats while remaining flexible enough to defend against the threats of tomorrow. Addressing this need, WatchGuard® created the Intelligent Layered Security (ILS) architecture. This technology provides the best defense possible against today's complex and rapidly changing threat environment. This paper describes our unique approach, and explains why it delivers better protection than other UTM implementations. The ILS model is currently deployed on our Firebox® X Core™ and the Firebox X Peak™ UTM appliances.

Intelligent Layered Security Architecture: Overview

The WatchGuard ILS architecture consists of six security layers intelligently cooperating with one another to dynamically detect, block, and report on malicious traffic, while passing benign traffic through as efficiently as

possible. This design results in a superior system, capable of defending networks against both known and unknown attacks without sacrificing performance.

For this discussion, a layer is a logical construct that defines a conceptual boundary between components of a network's security infrastructure. We're regarding each different type of security technology as a separate layer.

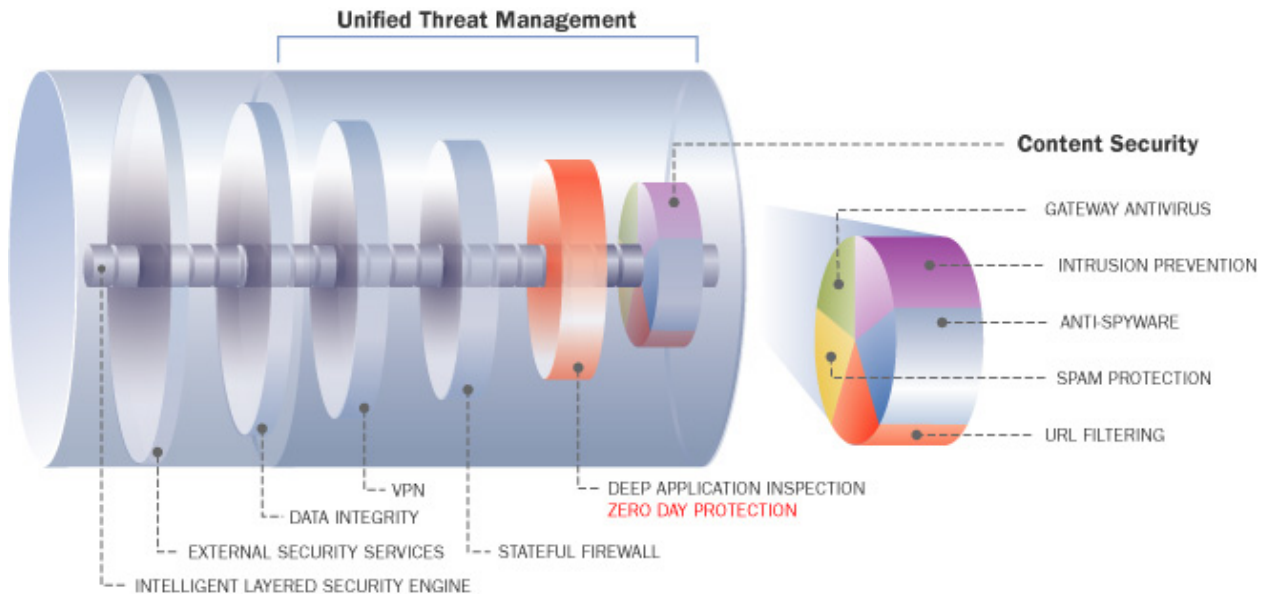


Figure 2: The Intelligent Layered Security Architecture

The ILS engine is the central nervous system of the architecture. By designing each layer to take advantage of and reinforce the capabilities of the other layers, and by exchanging information about the traffic being processed between the layers, it provides maximum protection, reliability, and performance. Let's look at an overview of each layer:

1. **External Security Services**, providing technologies to extend protection beyond the firewall, and information which empowers the end user/administrator to be more efficient
2. **Data Integrity**, validating the data packet integrity and packet protocol conformance
3. **Virtual Private Networking (VPN)**, ensuring secure and private external communications
4. **Stateful firewall**, restricting traffic to those sources, destinations, and ports which are allowed by the security policy
5. **Deep application inspection**, ensuring conformance with application layer protocol standards, blocking dangerous files by pattern or file type, and blocking dangerous commands and modifying data to prevent leakage of critical system information
6. **Content Security**, analyzing and regulating traffic for appropriate content, including services as diverse as Gateway AntiVirus (AV), Intrusion Prevention Service (IPS), spyware protection, spam protection and URL filtering

Although there are six distinct layers identified in this model, there are many functions and capabilities in each of these layers which are designed to cooperate with and pass information to other capabilities either within the same

layer or at different layers. All the layers are designed to be easily extensible as new security functions are required to handle new threats.

Whether a WatchGuard firewall with ILS is deployed at the network perimeter as an endpoint device, or at the core of the infrastructure, it provides key security capabilities vital to a protecting the network. Details of these capabilities are given in the section of this paper entitled, “Intelligent Layered Security Architecture Details.”

Intelligent Layered Security Benefits

The layers of the ILS architecture are designed to work together to provide:

Better Security

- **Zero day protection** - blocks many threats inherently, no ‘window of vulnerability’ exists for these threats
- **Proactive identification and blocking of attackers** - Identifies attacks and attack behaviors, and drops subsequent attacks from the same site
- **Minimal false positives** where content security (signature-based) technologies are used, such as Gateway AV/IPS

Greater Ease of Use

- **Deep application inspection** and other ILS layers are "always on"
- **Well-designed defaults** block most attacks without requiring complex configuration
- Proactive prevention

Better Performance

- **The layered architecture** allows attacks to be detected and blocked with the minimum amount of processing

Better Security – How Does It Work?

Zero Day Protection

Zero Day protection means defending against attacks which are not yet known, so that when a new attack emerges there is no window of vulnerability for the network being attacked. There are many new attacks launched each year; however because most of these attacks use techniques closely related to previous attacks. It is rare that an entirely new class of attack emerges. By understanding the classes of attacks, defense mechanisms can be developed which defend against whole classes of attacks. This is much more effective than the reactive, signature-based technologies which rely on fingerprinting each new attack as it emerges.

The deep application inspection layer contains core capabilities which provide Zero Day protection, including:

- **Protocol Anomaly Detection (PAD)**
Protocols define the way a given exchange of data between two systems should proceed if everything goes according to plan. Since some servers don’t deal gracefully with malformed traffic, many attacks rely on violating application layer protocols to allow the hacker to create a Denial of Service (DoS) attack, or to get root access to the server. By enforcing the protocol RFCs or standards, we can prevent this class of attack. In addition to protocol violations, this mechanism will trap illegal arguments in commands and prevent

many buffer overflows.

Examples of attacks blocked:

CAN-2004-0434: k5admind (kadmind) for Heimdal allows remote attackers to execute arbitrary code via a Kerberos 4 compatibility administration request whose framing length is less than 2, which leads to a heap-based buffer overflow. See <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-2004-0434>.

CERT CA 2003-12: There is a remotely exploitable vulnerability in Sendmail that could allow an attacker to gain control of a vulnerable Sendmail server. Due to a variable type conversion problem (char to signed int), Sendmail may not adequately check the length of address tokens. A specially crafted email message could trigger a stack overflow. See <http://www.cert.org/advisories/CA-2003-12.html>.

- **Pattern Matching**

In order for malware to be effective, it must get onto a computer and then execute. This means it needs to be embedded in a file type which allows execution on the target computer. By blocking file and mime types used to carry malware (i.e., .exe, .pif, .scr) we can prevent malware from getting into the network.

In many cases however, we will need to bring potentially dangerous file types such as .exe or .dll into our network from trusted sources. Examples of this would be software updates from Microsoft or printer drivers from Hewlett Packard. In this case, we need to be able to identify trusted sources and allow file types otherwise considered to be dangerous to be imported from those trusted sources.

Examples of attacks blocked:

SOBER.K Virus is a .pif file which is typically inside a .zip file, but the file name can change. This attack was automatically blocked by the WatchGuard pattern-matching technology. By blocking .pif and .zip files (done by default), the virus file is dropped before it can enter the network.

BAGEL.BB Virus can be one of these file types: .exe, .scr, .com, or .cpl. This attack was automatically blocked by the WatchGuard pattern-matching technology. By blocking these file types (done by default), the virus file is dropped before it can enter the network.

Spyware programs such as CoolWebsearch, Transponder/Vx2, MyWebsearch, Moneytreedialer, and huntbar are transmitted in .exe and .dll files. By blocking these file types (done by default), the spyware file is dropped before it can enter the network.

- **Command Limiting**

Application protocols contain commands or verbs, some of which are used to send and receive data; however many are administrative commands which we normally would not want to be executed from outside the network. By blocking potentially dangerous commands such as FTP SITE command and SMTP DEBUG command, we prevent an entire class of attacks.

Examples of attacks blocked:

Microsoft Security Bulletin, April 12 2005: Microsoft released a security bulletin describing a critical buffer overflow flaw affecting Microsoft® Exchange Server 2000 and 2003:

- By connecting to your Exchange server and sending a specially-crafted request, any remote attacker could create a buffer overflow and gain full control of your company's e-mail server
 - According to Microsoft, the buffer overflow results from the way Exchange handles one of the Microsoft-enhanced SMTP verbs (commands)
- The command limiting capability of deep application inspection on the SMTP protocol blocks the use of this enhanced SMTP verb (command), and thus blocked the attack without requiring any configuration change or update

- **Cloaking**

Cloaking hides critical server information from hacker probes. For example, in SMTP it can masquerade domain names, hide mail server type and patch level, and remove information from Message IDs and MIME boundary strings. This prevents hackers from identifying particular servers and targeting them with attacks which are known to work against those servers.

Examples of attacks blocked:

Microsoft Vulnerability #MS05-21 gave an attacker full control of an Exchange Server. This is how this type of attack works:

- Start with an automated scan to figure out if the mail server is vulnerable, If vulnerable, the scanner runs the exploit, once the exploit is run, a root kit is installed.
 - The hacker now has complete control of the Exchange Server
- By cloaking the identity of the Exchange Server, the scanner is fooled into thinking that the server isn't vulnerable.

- **Filters/Blocks Headers**

Another class of attack relies on creating malformed headers to exploit vulnerabilities in poorly written server implementations. This mechanism allows us to block such headers.

Examples of attacks blocked:

CERT VU#897604 is an example of where a specially crafted email address can be used to take control of a Sendmail server by triggering a buffer overflow.

As well as providing Zero Day protection against a broad range of attack classes, these capabilities are based on well-defined standards and policy decisions, so they do not introduce the problem of false positives.

Proactive Identification and Blocking of Attackers

The objective of this mechanism is to identify hackers either before they launch an attack (by their behavior) or the first time they launch an attack. This gives us the ability to dynamically respond to the behavior or attack by blocking the attacker's IP address; this is known as shunning. Blocking an IP address is a very simple process which minimizes the subsequent processing effort required to defend against repeated attacks. If we can detect a behavior before an attack is launched, this mechanism can even protect us against new or unknown attacks. This is how the capability works:

- **Identify Attacks**

The strength of the WatchGuard ILS architecture is its distributed intelligence capability. As well as blocking the attacks for which it is designed, each layer has the ability to identify and report the IP address of an attacker site. This applies to many different classes of attacks. For example, a DoS attack, IP options, violation of a protocol such as SMTP (PAD-based attack), or even an attack blocked by the GAV/IPS engine.

- **Identify Attacker Behavior**

An attacker can also be identified by their behavior before an attack is launched. In fact, Tel Aviv University reports that: "there was a 96.3% match between scans and identified attackers. . .In other words, virtually every scan was followed at some point in time by an attack from the same source." (Network World, August 25, 2003) The behaviors which ILS currently identifies are:

- Port scans

- Address scans
- Use of IP options, spoofing, and source routing
- **Shun Attackers**
The shunning mechanism takes the reports of attacks or attacker behavior and dynamically blocks the reported IP addresses for a user configurable period of time. This is particularly effective in:
 - Blocking automated attack tools – The firewall and its protected systems simply disappear as far as the attacker is concerned
 - Reducing processing for second and subsequent attacks from same site to simple IP blocking

The firewall can also be manually configured to ignore all traffic on specific ports or from specific IP addresses.

Minimize False Positives

It is well known that signature-based technologies such as antivirus engines can be prone to introducing false positives (wrongly identifying legitimate traffic as attacks). Statistically, the probability of a false positive is directly proportional to the volume of data scanned and the number of signatures used, if all other factors are equal. In ILS, the layers work together to reduce both the volume of data scanned and the number of signatures used.

- **Gateway AntiVirus and Deep Application Inspection**
The deep application inspection layer detects attacks by enforcing protocol standards and blocking known dangerous file types as described previously; blocking known dangerous file types does not lead to false positives, and significantly reduces the number of files which need to be scanned by the signature-based antivirus engine, significantly reducing the probability of a false virus detection.
- **Intrusion Prevention**
Optimized with ILS In a typical standalone intrusion prevention system such as Snort, all forms of possible attacks are addressed by means of signatures. The ratio is basically one signature for each attack defended against. This results in a Snort database of around 6000 signatures.

In the ILS architecture, the majority of these attacks are blocked by other layers of the architecture, from the data integrity layer through to deep application inspection. This means that the total number of signatures required to deliver the same level of protection in the content layer IPS engine is around 1000.

In addition to this, the deep application layer is capable of identifying many of the protocols being processed and passing that information up to the IPS engine. This way, the IPS engine, which is protocol-aware, can further reduce the number of signatures which are required for a particular scan. For example, scanning SMTP traffic regardless of port means only the SMTP signature subset needs to be used.

Overall this reduces not only the amount of traffic scanned, but also the number of signatures used for each scan – significantly reducing false positive rates.

Better Performance

Real-world performance in a UTM appliance is difficult to measure, as it will be dependent on the traffic mix, the complexity of firewall rules, and number of VPN tunnels, as well as number and type of services enabled, such as Gateway AV and anti-spam.

For this reason, UTM vendors give performance specifications which are typically best case for each service, or function so that the customer can understand the maximum expected performance. For instance, Gateway AV performance will typically be measured with a minimal firewall configuration and all other services disabled.

This makes it difficult to compare the real-world performance of UTM appliances. While the performance of a particular service or function can be compared from the manufacturers' specifications, the effect of cooperation between various services and functions can significantly influence overall performance, and is not measured.

WatchGuard ILS is designed to optimize the real-world performance of the UTM appliance through careful design of how the security layers interact. The three key design principles are:

Processing Order

The ILS engine is designed so that the minimum amount of processing is performed on each traffic stream in order to detect an attack. Simple attacks, such as malformed packets and DoS which take minimal effort to detect, are eliminated first. This means the amount of traffic which reaches the more resource-intensive processes, such as the intrusion prevention engine, is greatly reduced.

Example:

BAGEL.BB Virus can be one of the following file types: .exe, .scr, .com, or .cpl. Most UTM appliances would use a gateway antivirus scan to detect this virus, which is very processor-intensive. WatchGuard ILS can discard a file attachment based on its file type. Files such as .exe, .scr, .com, and .cpl are blocked by default, so the WatchGuard Firebox X would discard the virus without needing to examine the file contents or running the Gateway AV scanning engine, thus greatly reducing the processing required.

Another benefit of correct processing order is a reduction in the number of signatures needed for effective intrusion protection. Signature matching is one of the most computationally expensive functions in a UTM.

Example:

SNORT, a well-known intrusion detection system, uses around 6000 signatures. WatchGuard deep application inspection and other ILS layers automatically block the majority of attacks for which SNORT requires signatures. Therefore, WatchGuard UTM appliances only require around 1000 signatures to deliver an equivalent level of protection.

Information Exchange between Layers

In many UTM solutions, the security functions work independently, not integrating in a way that enables the information discovered in one function to be leveraged to make the other functions more effective. This means that good traffic is often unnecessarily processed many different times by different security functions. Intelligent Layered Security passes information between its layers to reduce and fine tune the processing required by the security functions.

Example:

The **deep application inspection layer** can pass protocol information to the IPS engine, which then uses a subset of signatures applicable to that protocol, rather than using the whole signature set. This makes IPS scanning much more efficient. For SMTP scanning, less than 20 of the approximately 1,000 signatures in the IPS signature set need to be used.

Dynamic Blocking (or Shunning)

Most real-world attacks are launched from automated attack tools. These tools typically scan a network for vulnerabilities before launching an attack. With the exception of DDoS or spam, the attacks which follow are typically from a single IP address.

The ability of ILS to detect behaviors such as port scans and address scans in addition to actual attacks, and then to use this to shun attacking sites, reduces the traffic load on the system tremendously when under attack. When the first scan or actual attack is processed, the shunning mechanism is triggered. All subsequent traffic from the attacking site is then blocked for a user-determined period of time - even on ports which are open to legitimate users. Additional processing to analyze the attacks is not required.

Intelligent Layered Security Architecture Details

Now that we've introduced the WatchGuard concept of ILS and described the way the distributed intelligence of the system provides better security, let's look at each specific security layer in the WatchGuard architecture to understand what it does.

Layer 1 – External Security Services

In order for the network to function at peak efficiency, it is necessary to supplement its protection with external security services, such as vulnerability assessment and desktop antivirus, while simultaneously assisting the administrator with making astute choices regarding the configuration of the firewall and its associated systems. In the WatchGuard model, this concept is represented by a “layer” of security which exists outside the firewall, adding to the network's capabilities by preempting what the firewall would otherwise have to do itself.

Emphasizing preparation, prevention, and off-box services, the external security services layer enables the administrator to address the network as a single entity, its parts working together efficiently and securely to meet the organization's goals.

Layer 2 – Data Integrity

The data integrity layer is the first line of defense on the Firebox itself. This layer validates the data coming onto the network, assuring that it conforms to packet protocol specifications. All traffic coming to the network passes through this layer. Computationally, it is the cheapest and best place to stop attacks. Traffic filtering can occur very rapidly here because most of the verifications performed are straight yes/no propositions: For instance, does the packet conform to RFC standards? Is the header information longer than the standard allows? If so, the packet is dropped without further processing. The key capabilities of this layer are:

- Traffic Normalization
 - Validating IP checksums protects your network by stopping any malformed TCP/IP traffic before passing the data to the next layer
 - Detecting and dropping traffic associated with DoS, DDoS, and fragment overlap attacks before passing the good traffic to the next layer via the WatchGuard patent-pending DDoS technology with sophisticated rate control mechanisms to allow more legitimate traffic to pass through even while under attack. Capabilities include defenses for IPSEC, IKE, ICMP, UDP and SYN flood attacks
- Detecting and blocking traffic associated with:
 - Port scans
 - Address space probes
 - Spoofing attacks

- Blocking traffic from a specific source (using the shunning mechanism), based on detection of attacks by other layers of the architecture

Normalizing the traffic at this layer and blocking traffic from known attackers improves overall system performance because ILS can process data at this layer very efficiently and other layers subsequently receive only properly-formatted, predictable packets.

Layer 3 – Virtual Private Networking

Once the traffic has been validated and normalized, ILS determines whether the traffic is an encrypted stream associated with a known VPN connection from a partner, customer, remote employee, or branch office. If so, the VPN layer decrypts the traffic and passes it on. If the traffic is encrypted with an unknown key, traffic is dropped. If the traffic is unencrypted, and there is nothing for the VPN layer to do, the traffic cuts through to the next layer for a policy decision.

For outbound traffic, the VPN layer efficiently and securely extends the reach of your appliance to mobile users, branch offices, and external partners around the globe. The VPN layer implements mobile user and branch office VPNs through PPTP or IPSec protocols. By properly configuring the VPN, you'll ensure secure and private communication via the public Internet.

Layer 4 – Stateful Firewall

This is the layer at which the firewall administrator specifies what will and will not be allowed to pass through the firewall based on source IP, destination IP, and ports. The ILS NAT capabilities are also implemented at this layer.

Though many kinds of attacks rely on using malformed packets to evoke responses from targeted machines, an individual packet can comply with all RFC standards, yet still have a malicious purpose. For example, a hacker gathering information about your network can try to sneak a packet into your network by setting the packet's "Reply" flag, thus disguising it as information the target server requested. A non-stateful packet filter will do the inspections described in Layer 2 above, then admit the packet, to which the targeted server will respond, aiding the hacker mapping your network. A stateful packet filter however, will know that a "Request" packet was never sent to the hacker's IP address, understand that a "Reply" packet makes no sense if there was never a "Request," and will discard the packet.

ILS provides this level of stateful protection, but goes one step further. The stateful firewall layer tracks the port and protocol information on all connections, as well as the state of those connections, and can also trigger the shunning mechanism when an attack is detected, or when a port which is blocked by policy is probed if the "block packets not handled" setting is activated. By doing so, ILS is able to defeat attacks designed to penetrate less sophisticated appliances, and reduces the burden on the firewall due to repetitive attacks from the same source.

Layer 5 – Deep Application Inspection

Traffic that passes the checks of the stateful firewall layer may be routed to the deep application inspection layer, where ILS determines if it is 'appropriate.' If further inspection is not required, the traffic can simply be 'cut through' to optimize performance. At the deep application inspection layer, TCP connections are terminated and new connections are built on each side of the firewall. The outgoing packets are completely normalized, preventing any attack which relies on packet characteristics from succeeding and data may be processed as a stream allowing detection of attacks which may span many packets.

Deep application inspection detects and manages or prevents/denies:

- Protocol anomalies

- Buffer overflows
- Unauthorized connections
- TCP hijacking
- Leakage of network information
- Dangerous attachments, viruses, and worms based on MIME type or patterns (examples: *.bat, *.cmd, *.com, *.exe, *.hta, *.inf, *.pif, *.scr, *.wsh, etc.)
- The use of potentially dangerous commands

In the “Better Security – How Does It Work?” section, we looked at the core mechanisms of the deep application inspection layer which provide protection against the above attacks. These are:

- Protocol Anomaly Detection
- Pattern Matching
- Command Limiting
- Cloaking
- Filters/Blocks Headers

As well as providing Zero Day protection against a broad range of attack classes, these capabilities are based on well-defined standards and policy decisions, so they do not introduce the issue of false positives. Let’s look at how these capabilities translate into specific protections for the core protocols of SMTP, HTTP, DNS, and FTP, and how they interoperate with other layers of ILS.

SMTP – Incoming or Outgoing

Some of the most devastating attacks we have seen have been blended threats -- worms that use multiple infection and propagation methods to spread. Many of these have relied upon the Simple Mail Transport Protocol (SMTP, or email) as a transport for propagation. The WatchGuard SMTP protocol handler stops:

- Potentially harmful email attachments
- Non-essential SMTP commands
- Protocol violations

It adds sites that transmit malformed traffic to the blocked sites list. In doing so, the SMTP protocol handler has proved to be extremely effective in neutralizing these types of attacks.

With our strong default SMTP filtering logic, we include basic anti-viral capabilities through file pattern matching and attachment blocking. The SMTP protocol handler is capable of:

- Blocking traffic that fails strict RFC standards for SMTP messages
- Pattern Matching to filter email attachment types by name or MIME type
- Restriction of SMTP commands (verbs)
- Cloaking critical server information in message IDs, server replies, and mime boundary strings, which gives protection from mail server profiling – a common precursor to attack
- Control of allowed and denied email headers
- Controlling SMTP email size and line length

- Restricting maximum number of recipients (helps against spamming)
- Restricting address length
- Controlling bat/CHUNKING, ETRN, and 8-bit or Binary MIME in ESMTP
- Controlling ESMTP authentication types
- Controlling length and allowed characters in SMTP greetings
- Blocking Source Routing and 8-bit characters
- SMTP relay protection
- Protection against email address harvesting by masquerading domains
- Source and destination address white- and black-listing
- Passing protocol information to the (optional) GAV/IPS module
- Relaying traffic to (optional) GAV/IPS module
- Triggering ILS shunning for any detected attacks, reducing the load caused by subsequent attacks from the same site

HTTP-Client

The HTTP-Client protocol handler offers fine-grained control over what sort of Web traffic can reach your users' Web browsers and other HTTP clients. Administrators using WatchGuard firewalls can:

- Block traffic that fails strict RFC standards for HTTP traffic
- Use pattern matching on content types as a tool to combat viruses, worms, and trojans
- Remove or deny cookies, applets, form submissions, ActiveX, and unknown headers in HTTP requests
- Limit HTTP request methods – the HTTP equivalent of command limiting
- Cloak critical server information
- Control authorization methods
- Restrict request and response header types to prevent malformed or suspicious header types
- Control allowed body content types
- Specify unsafe URLs and paths with regular expressions, empowering you to efficiently deny the download of executables and other dangerous content such as DLL files, which often lead to spyware infections
- Pass protocol information to the (optional) GAV/IPS module
- Relay traffic to (optional) IPS module
- Trigger ILS shunning for any detected attacks, reducing the load caused by subsequent attacks from the same site

HTTP- Server

The HTTP-Server protocol handler offers fine control over what sort of Web traffic can reach Web servers protected by WatchGuard Fireboxes with ILS. The capabilities of the HTTP-Server protocol handler are the same as HTTP-Client (see previous), however the default settings required to protect a server are different than protecting a client.

DNS

Some Domain Name Service exploits (like TSIG, nxd, iquery, infoleak, zxf, etc.) turn the transport layer that conveys DNS requests and answers into an attack tool capable of granting root-level access to your DNS server, or keeping anyone from using it. Such attacks often use malformed DNS requests to convey malicious code. The DNS protocol handler monitors the headers of the DNS requests and blocks queries where the header class, type, or length is abnormal. The WatchGuard DNS protocol handler:

- Blocks traffic that fails strict RFC standards for DNS
- Cloaks critical server information
- Checks DNS packet headers and discards all packets that are incorrectly structured
- Controls DNS Opcodes, Query types, and Query Names
- Relays traffic to (optional) IPS module
- Triggers ILS shunning for any detected attacks, reducing the load caused by subsequent attacks from the same site

FTP

The WatchGuard FTP protocol handler gives the administrator the ability to combat abuse of FTP network resources by:

- Blocking traffic that fails strict RFC standards for FTP server and client traffic
- Forcing session timeouts
- Restricting allowable FTP commands
- Cloaking critical server information
- Setting maximum username, password, filename, and command line lengths
- Restricting file types allowed for download
- Restricting upload access to files, directories, or file names/types by pattern
- Passing protocol information to the (optional) GAV/IPS module
- Relaying traffic to (optional) IPS module
- Triggering ILS shunning for any detected attacks, reducing the load caused by subsequent attacks from the same site

TCP

The generic TCP protocol handler performs the basic functionality of terminating TCP connections and building new connections on each side of the firewall. This means that packets are normalized and data may be processed as a stream, allowing detection of attacks and processing of data which may span many packets. The TCP protocol handler can also detect if the traffic is HTTP, and if so, behaves exactly like the HTTP protocol handler with the same controls that was described previously.

Layer 6 – Content Security

The Content Security layer looks at the actual data traffic as distinct from the protocols used to transport it. At this layer, we have optional security services as diverse as Gateway AntiVirus, Intrusion Prevention Service with anti-spyware, spam protection, and URL filtering.

Gateway AntiVirus Service

WatchGuard Gateway AntiVirus service identifies and blocks worms, spyware, and trojans within email attachments, blocking threats from entering your network and executing dangerous payloads. WatchGuard recommends using Gateway AV in conjunction with desktop AV, as it provides two important benefits:

- Because it is gateway-based, it is not susceptible to being disabled by new viruses in the way that desktop AV is
- Gateway AV and desktop AV working together provide faster average response time for signature updates than Gateway AV or desktop AV alone; this is because the response time will always be that of the first responder, which could be the Gateway AV or the desktop AV

The integration of Gateway AV with other security layers in ILS provides some important benefits, namely:

- **Efficiency** – The Gateway AV service only scans files that have not been blocked by the deep application inspection layer’s pattern-matching capabilities, greatly reducing the number of files which are scanned
- **More granular control** – Gateway AV finds viruses in file types which are allowed by deep application inspection, such as .zip, .doc, etc.

The WatchGuard Gateway AV service enables infected files to be allowed, denied, or locked. File locking is a unique approach to dealing with the issue of users wanting to retain files which are known to be infected with viruses, an issue that is problematic for most system administrators.

The most common solution is quarantining, where the system administrator sets up a separate quarantine server and directs the Gateway AV system to send infected files to the quarantine server. This means, however, that the system administrator has to set up a separate server, decide on a retention policy (how long should infected files be saved?), and constantly monitor disk space to make sure the server is not brought down by a virus outbreak.

WatchGuard file locking provides an intelligent alternative. When a file is found to be infected, the file can be encrypted but still sent on to the end user. This prevents accidental execution by the end user; however it does give them the opportunity to decrypt and clean the file with a special tool they can obtain from the system administrator. This takes the burden off the administrator, since the user can decide whether to simply discard the file, how long to retain the file, and whether they want to attempt to clean it.

The WatchGuard Gateway AV database contains virus, spyware, worm, and trojan signatures, including both WildList and “zoo” viruses. A broad range of compression/decompression algorithms is supported, including ZIP, RAR 2.0, TAR, GZIP, ARC, and CAB files. Signature delivery is automatic, and signature update checks can be programmed for any desired interval. The targeted threat response time is 8 hours, which is significantly better than industry average.

Intrusion Prevention Service with Anti-spyware

The WatchGuard Intrusion Prevention Service (IPS) provides inline protection from attacks that comply with protocol standards but carry malicious content. It is a signature-based service designed to protect against a broad range of attacks, including cross-site scripting, buffer overflows, and SQL injections.

The two main problems with most inline intrusion prevention systems are speed of execution and false positives. Tight integration of the WatchGuard IPS with the other layers of ILS provides significant benefits in both of these areas.

Because other layers of ILS block up to 70-80% of attacks (deep application inspection being particularly effective in this area), signatures for those blocked attacks are not required. This reduces the overall number of signatures and increases processing speed while simultaneously reducing the chance of false positives (the chance of a false

positive increases statistically with the size of the data scanned and the number of signatures used). For example, Snort requires around 6000 signatures, while WatchGuard IPS requires only around 1000 to achieve an equivalent - or better - level of protection.

Efficient Processing for Optimal Performance

As the deep application inspection layer is protocol-aware, it can let the IPS know what protocol is being processed. This means that the IPS needs only to scan the traffic with signatures applicable to that particular protocol. Instead of using the full 1000 signatures, the number of signatures can then be significantly reduced, possibly to just 6 for SMTP, again speeding processing and reducing the chance of false positives.

As the IPS can trigger ILS shunning, only the first attack needs to be analyzed from a particular hacker's IP address. This means that an attacker launching multiple attacks can be blocked by shunning, which does not significantly load the firewall. This is unlike other offerings in which each attack must be analyzed, wasting valuable processing power and slowing throughput.

Peer-to-Peer and IM Blocking

The WatchGuard IPS can also selectively block IM services, specifically AIM, Yahoo, IRC, and MSN Messenger. This protects against IM-based security threats, including those which allow the attacker to gain control of a machine running an IM client, and infections by viruses transferred in files over IM.

Peer to Peer (P2P) applications such as Napster, Gnutella, Kazaa, Morpheus, BitTorrent, eDonkey2000, and Phatbot can also be blocked. P2P presents two challenges: 1) it uses up valuable bandwidth better used for business purposes; and 2) it is a well-known vector for transmitting spyware (Kazaa in particular). By blocking P2P usage, WatchGuard easily solves these two problems.

Spyware Protection Methods

Spyware is propagated in many ways outside of P2P applications, including embedded files, cookies, and drive-by downloads. Spyware can log your keystrokes, rifle through your files for password and identity data, fill desktop screens with ads. It also slows PCs to a crawl, and depletes network bandwidth. The WatchGuard IPS service includes both signatures and unique scanning schemes to block spyware at different points in its lifecycle, including installation, reporting back to the propagation/host server, and post-installation activity of the spyware application. This is accomplished through a number of cooperative procedures:

- **Site Blocking** – The IPS engine will attempt to block access to known spyware-hosting or -download sites that deploy known spyware-bundling programs during an end user's HTTP session
- **Signature-based Content Inspection** – The IPS engine will also apply continuously updated attack signatures to all configured traffic to identify and block incoming spyware downloads – including covert drive-by downloads
- **Configuration Shutdown** – Successful spyware configuration usually requires the application to communicate an installation result report and retrieve initial configuration data from the host server; IPS identifies and blocks this initial configuration communication
- **Post-Install Shutdown** – Should an infected PC enter a secure network, spyware applications will utilize the network connection to create a communication channel for additional activities; IPS will attempt to identify and block these processes, which may include information stealing/hijacking, additional spyware installation, and unsolicited advertisements

The WatchGuard proprietary IPS engine integrates tightly with other firewall functions and produces comprehensive log messages which are fully integrated with the logging system. This enables the system administrator to easily detect any machines found to be infected with spyware, and remove it.

Spam Protection Service

Spam accounts for more than 63% of all email today, and represents a major problem for most companies. The WatchGuard spamBlocker Security Service utilizes Commtouch® Recurrent Pattern Detection™ (RPD) technology to give you real-time protection from spam outbreaks with 99.95% accuracy - without using signatures or filters.

Rather than evaluating keywords and content, this technology analyzes large volumes of Internet traffic in real time to identify the repetitive component (or DNA) of each outbreak as soon as they emerge. Close to 500 million messages per day are sampled and advanced algorithms detect, identify & classify new outbreaks typically within 1-2 minutes. These algorithms are also capable of distinguishing solicited bulk email from spam. spamBlocker utilizes this technology to give you up-to-the-minute protection from spam attacks by comparing suspected spam directly with the Commtouch Detection Center (which has around 20,000,000 spam classifications) in real time. This technology provides four key benefits:

- Extremely fast response to new outbreaks.
- Near zero false positives – Best in the industry at distinguishing legitimate communication from spam
- High spam detection rate – blocking 97% of unwanted emails
- Language agnostic – Spam is blocked regardless of the language, content, or format of the message.

spamBlocker uses the fundamental characteristics of mail traffic to identify outbreak patterns and remove 97% of spam at the network gateway - within minutes of distribution. By focusing on the bulk nature of the message instead of the individual content, language, or format, spamBlocker provides real-time identification of global spam – including phishing attacks - and enables continued high throughput for other network traffic.

URL Filtering Service

The WatchGuard WebBlocker URL filtering service enables you to configure not only who gets Web access and who doesn't, but also what type of Web access is available. Using an intuitive set of controls built into WatchGuard System Manager (WSM), you can quickly select which categories of Web pages users get access to - and what time of day they get access.

WebBlocker utilizes a site database and engines from the global Web-filtering leader SurfControl to ensure the most accurate categorization and complete coverage. WebBlocker uses numerous categories to help you block content you don't want to allow on your network:

- Blocking known spyware sites and sites known to contain malicious content helps protect your network assets
- Blocking pornography can assist in enforcing company policy on sexual harassment in the workplace
- Blocking sports content may increase workplace productivity

With customizable exceptions lists, per-user authentication, and provisions for different access policies according to the time of day, WebBlocker enables you to efficiently enforce IT policies. WebBlocker will also help keep your network and end users secure from viruses, worms, and spyware by keeping them from reaching sites that are known distribution points for these malicious applications.

Conclusion

Although the concept of layered security has been talked about for years, small- to mid-sized enterprises have never had the resources to deploy it appropriately. UTM appliances promise to solve this problem; however UTMs which rely on a collection of disparate technology functions on a single appliance don't provide the level of security, ease of use, or performance that customers require. The WatchGuard ILS architecture provides an intelligent, multi-layered defense in which each of the layers cooperates with and passes information to other layers, creating a level of security that far surpasses other solutions. Because the security landscape is constantly changing, the ILS architecture was designed with extensibility in mind, so future functionality can be added seamlessly and quickly to combat emerging threats.

Firebox® X: SME Security Architecture




Figure 3: Network security architecture after deploying Firebox X.

What Does This Mean For You?

- **Better security through Zero Day protection** at the deep application inspection layer and the active cooperation between multiple security layers
- **Easy to use**, strong, and sensible default configurations help you to quickly get new installations up and running, and to know at a glance what kind of traffic is passing through your Firebox

- **Better performance** because malicious traffic is efficiently dropped at the layer which is least expensive from a computational standpoint
- **Lower false positives** for GAV and IPS, because the tight integration of GAV and IPS capabilities and other ILS layers minimizes the amount of signature-based scanning performed, and the number of signatures required
- **More (and more accurate) intelligence** regarding your network from the innovative WatchGuard approach to external security services

WatchGuard Intelligent Layered Security provides the best security and remains highly cost-effective through the most efficient use of processing power available. WatchGuard integrated security appliances are the ideal solution for your Unified Threat Management needs, both today and tomorrow.

For more information about WatchGuard security solutions, visit us at www.watchguard.com or contact your reseller.

ADDRESS:

505 Fifth Avenue South
Suite 500
Seattle, WA 98104

WEB:

www.watchguard.com

U.S. SALES:

+1.800.734.9905

**INTERNATIONAL
SALES:**

+1.206.613.0895

ABOUT WATCHGUARD

WatchGuard provides network security. The company's Firebox X family of upgradeable appliances delivers the performance, functionality and security strength to meet the needs of organizations of any size. WatchGuard's Intelligent Layered Security protects against emerging threats and provides the platform to integrate additional services offered by WatchGuard. All WatchGuard products include a LiveSecurity Service subscription for vulnerability alerts, software updates, expert security instruction, as well as individualized and self-help customer care. WatchGuard is headquartered in Seattle, Washington, with offices throughout Europe and Asia.

No express or implied warranties are provided for herein. All specifications are subject to change and any expected future products, features or functionality will be provided on an if and when available basis.

©2006 WatchGuard Technologies, Inc. All rights reserved. WatchGuard, the WatchGuard Logo, Firebox, Firewall, Peak, Core, LiveSecurity, and Stronger Security, Simply Done are either registered trademarks or trademarks of WatchGuard Technologies, Inc. in the United States and/or other countries. All other trademarks and tradenames are the property of their respective owners. Part. No. WGCE66298_081806