Addressing the Limitations of Deep Packet Inspection with Complete Content Protection

White Paper

January, 2004

Abstract

Network threats have evolved from relatively simple, connection-based attacks to more complex content-based attacks such as viruses, worms, and Trojans. At the same time, organizations are struggling to cope with other content-based threats, such as email spam and inappropriate Web content that reduce productivity and expose them to substantial liability. These new content-based attacks are not detected or stopped by the Stateful Inspection firewalls that have been deployed by many companies, causing a search for newer, more effective technologies. Recently, many firewall vendors have been touting the benefits of a technology called Deep Packet Inspection, promising better protection against content-based threats. While Deep Packet Inspection is more effective than Stateful Inspection for certain types of attacks, it falls far short of being a complete solution for protecting network and computing systems. Specifically, Deep Packet Inspection cannot detect a substantial portion of active viruses, Trojans and worms, and is completely ineffective for dealing with inappropriate Web content and email spam. A more effective technology, called Complete Content Protection, can detect and prevent the full range of content attacks in the network before they reach desktops, laptops, and servers. With an appropriate hardware-based platform, Complete Content Protection technology can be deployed in high speed networks without impacting the performance of network applications.

This paper discusses the characteristics and limitations of Stateful and Deep Packet Inspection technologies and explains the benefits of Complete Content Protection for providing comprehensive network security.

Fortinet, Inc.
920 Stewart Drive
Sunnyvale, CA 94085
USA

Contents

How Network Threats Have Evolved

Corporations and other organizations have come to realize the intense value of the proprietary data and intellectual property required to operate and be successful. As an example, in 2003 the total cost of proprietary information theft in the United States alone was estimated to be approximately more than $70 million, with an average reported loss at approximately $2.7 million per entity. [Source: 2003 CSI/FBI Computer Crime and Security Survey, based upon responses of 530 computer security practitioners in U.S. corporations, government agencies, financial institutions, medical institutions and universities.]

The proliferation of public and private networks and the increasing sophistication of network protocols and applications have driven a rapid escalation in the number and severity of attacks against computing systems, as shown in Fig. 1 below.

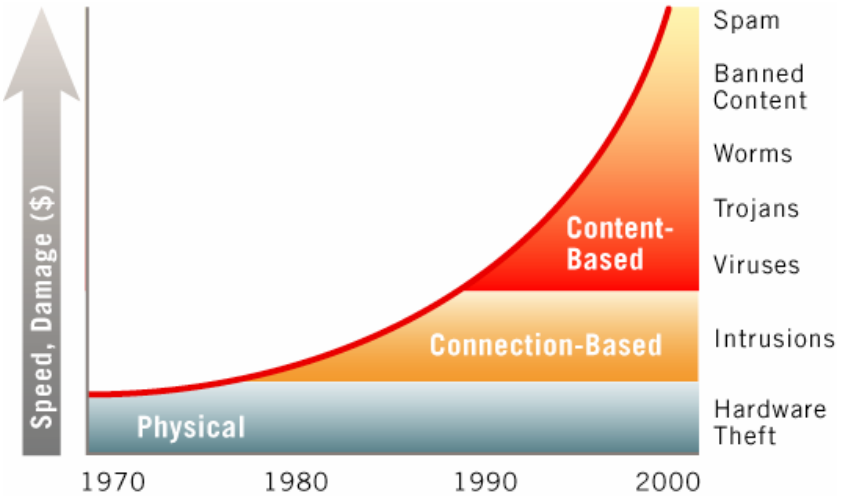

Early network protocols, such as Telnet, RPC and FTP, were relatively simple and required action by a dedicated hacker with a sustained connection to a remote system to launch an attack. The first incidences of this kind were identified by military organizations as intrusions to obtain classified information. The response to these types of attacks was the development of connection-oriented security systems, called Stateful Inspection firewalls, which control access to computing resources on a network by selectively allowing or denying the establishment of remote connections based primarily on the identity of the sender and receiver.

In the last ten years, applications have become much more complex, and protocols are used to carry much richer content. These changes have been exploited by attackers to develop more effective, content-based threats that circumvent connection-oriented security and that also have the ability to reproduce and spread automatically. Content-based threats are able to bypass connection-oriented Stateful-Inspection firewalls because they are typically delivered via connections that are inherently “trusted”. Content-based

threats include viruses, Trojans, worms, banned content and spam, and are readily propagated through email, web pages and other real-time communications applications.

The propagation speed of content-based threats and the resulting damage they can cause is considerable. For example, a recent email virus (called MyDoom) released on a Monday in North America accounted for nearly 30 percent of worldwide email traffic by the following Wednesday, just two days later. Other industry reports stated that nearly 3.4 million copies of the virus had spread worldwide in two days, which accounted for one of every 12 messages. [Source: CNET News.com, January 28, 2004]

Deep Packet Inspection – A Step in the Right Direction

As mentioned above, most firewalls utilize Stateful Inspection technology, which works at the network layer to track each connection traversing all interfaces of the firewall to ensure validity. Decisions on whether or not to accept the packets are based upon policies established by the network administrator related to which senders are allowed to reach designated computing systems on their internal network, and which protocols they can use to exchange information. While this filtering is useful, it is not adequate to determine the difference between, say, a valid email message or a message infected with a virus, because Stateful Inspection does not check the actual contents of the packets to distinguish malicious content from valid content.

As shown in Figure 2 below, Stateful Inspection examines only the "headers" of data packets, which contain information such as the sender's and receivers' addresses and the type of protocol and data contained in the packet "payload". However, much as one might try to determine the value of a postal letter based only on the addresses on the outside of the envelope, the contents of the packet payloads themselves are not examined. As a result, Stateful Inspection technology cannot tell the difference between valid and harmful data if it originates from an otherwise "trusted" source such as an ISP's email server or any public Web site. Stateful Inspection is therefore effective only for preventing simple intrusions and other connection-based attacks.

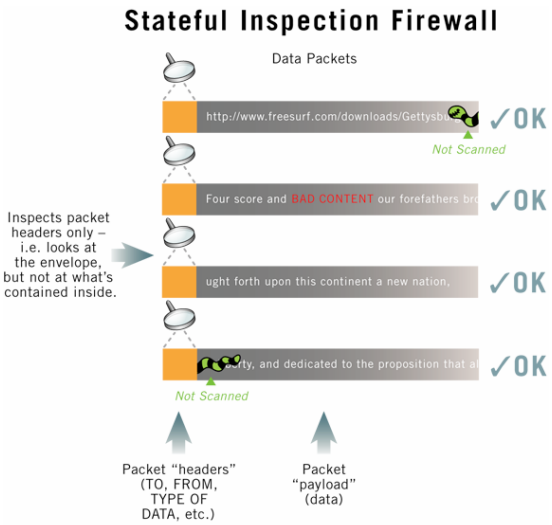

To address the limitations of Stateful Inspection, a technology known as Deep Packet Inspection (DPI) was developed. DPI goes further than Stateful Inspection to examine the payloads, or contents contained in packets in addition to the headers. As long as an attack can be contained in just a few packets, DPI can be effective in detecting and ultimately preventing the attack. As such, DPI technology is effective against buffer overflow attacks, denial of service (DoS) attacks, sophisticated intrusions, and a small percentage of worms that fit within a single packet.

The primary limitation of DPI technology is that it generally cannot detect threats that require many packets to transmit across the Internet. In general, the largest payload that can be delivered in a single Internet packet is approximately 1,500 bytes in length. Most viruses and worms are measured in dozens of kilobytes, and further, these threats may be embedded in files (documents, programs, etc.) that may be millions of bytes in length, requiring hundreds or thousands of packets. As a result, the likelihood of detecting most viruses and worms by analyzing the contents of just a few packets at a time is quite small. It is as if a terrorist were trying to ship a missile to another terrorist by cutting the missile up into 500 small pieces and shipping each in a separate package, along with pieces from something harmless, like a car. By themselves, none of the individual pieces might be recognizable as a part of a missile even if each package were unwrapped and examined individually. As a result, the missile would likely get through undetected.

Complete Content Protection: A Better Approach

The limitations of DPI can be addressed by using a more sophisticated approach to network security called Complete Content Protection, or CCP. The key aspects of CCP involve the reassembly of packet payloads into application-level objects, such as files, documents, and programs, followed by the scanning and analysis of the objects to detect content-based threats. Content reassembly assures that critical threats such as viruses and worms that are often embedded in large files are not missed. As shown in Figure 4 below, CCP technology can detect the full range of threats – including all viruses, worms, Trojans, inappropriate Web content, and email spam – regardless of the length of the threat or the length of the "host" file used to carry it.

Complete Content Protection and Network Performance
The compelling benefits of CCP technology come at a cost in terms of computing power. As shown in Figure 5 below, CCP can require one hundred to one thousand times more processing per packet than Stateful Inspection or DPI.

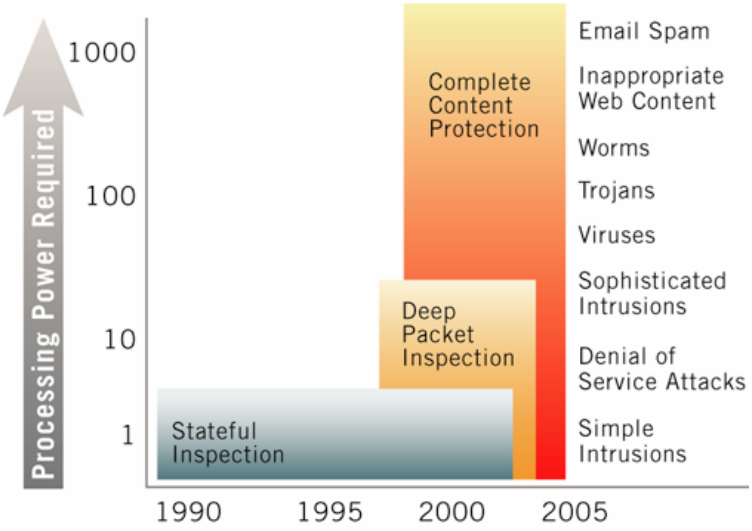

As a result, deploying CCP technology on standard computing systems (e.g. servers) or networking products (e.g. routers, firewall/VPN gateways, etc.) can dramatically reduce performance. Indeed, even executing DPI on standard computing and network security hardware can reduce performance by 75% or more. Therefore, in order to provide CCP technology without compromising network performance, a new type of hardware and software architecture is required.

A Unique Architecture for Complete, Real-Time Network Protection

Fortinet has designed a unique architecture that delivers CCP in a network-based solution with real-time performance. By blending the right core technology with the required applications and a world-class threat response system, Fortinet provides a powerful, cost effective solution for real-time network protection.

Core Technology
Fortinet's FortiGate Antivirus Firewalls are based on an integrated hardware/software architecture system that performs real-time, application-level content processing in the network, along with network level security functions. Fortinet's ABACAS™ (Accelerated Behavior and Content Analysis System) technology is the only such platform that can deliver application-layer services such as virus detection and content filtering in real-time at data rates greater than one gigabit/second.

The proprietary FortiASIC chip incorporates a hardware scanning engine, hardware encryption and real-time content analysis processing. The chip, which is designed by

Fortinet and available exclusively in FortiGate Systems, provides acceleration for the following four functions:

- Checking packet headers to make sure that they are from valid sources (firewall)
- Encrypting/decrypting and authenticating VPN packets (DES, 3DES, MD5, and SHA-1)
- Assembling packets into content and searching for attacks and banned material (signature and heuristic scanning)
- Counting packets and measuring flows (traffic shaping)

As shown in Figure 6 below, the FortiASIC processor provides much greater processing power to ensure CCP can be done using conventional servers or ASIC-based firewall/VPN appliances.

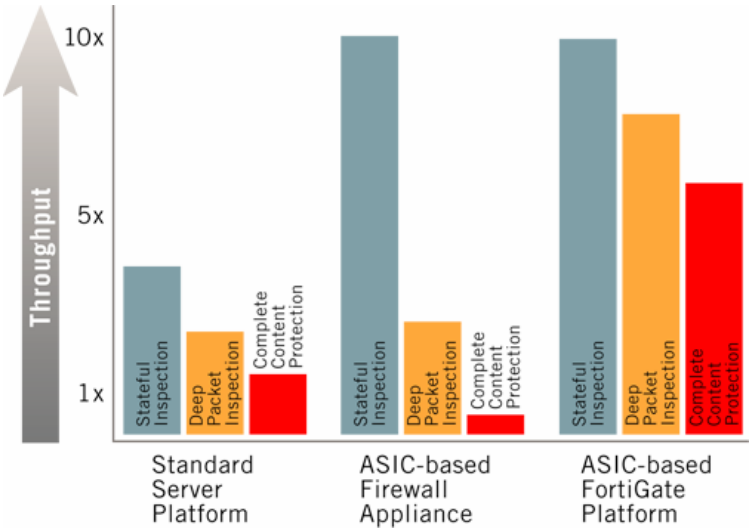

The FortiOS™ operating system is the high-performance, robust and reliable operating system that provides stateful inspection, DPI and CCP in a single platform. The platform deters content-based attacks overall by combining antivirus, intrusion detection and intrusion prevention.

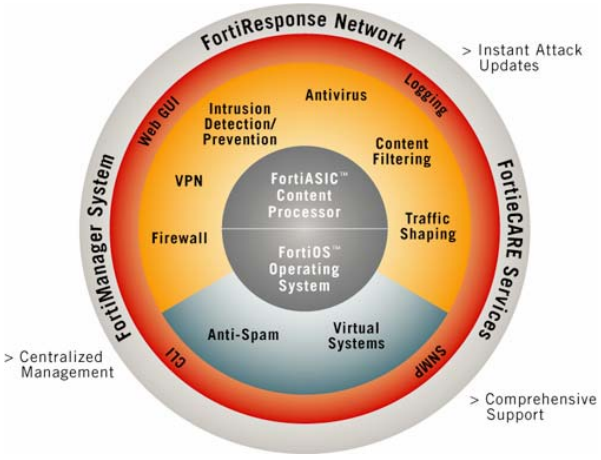

By adding stateful inspection, protocol analysis and DPI to CCP, the result is a complete system for recognizing and eliminating problematic threats. This chart provides a quick overview of the types of threats, examples and what system will detect them.

TYPE OF THREAT	EXAMPLES	DETECTED BY
Connection-based intrusions	Telnet attacks	Stateful inspection
Protocol attacks	SYN flood, ICMP flood	Protocol analysis
Packet-level content attacks	Buffer overflow, "probe phase" of some worms	DPI
File-level content attacks	Viruses, most worms, Trojans	CCP
File-level content threats	Inappropriate Web content, spam	CCP

FortiProtect™ Services for Real-time Response

FortiProtect Services (FPS) are a critical element of Fortinet's gateway-based network protection solution. Through FPS, Fortinet provides up-to-date network security threat information and timely virus and intrusion attack definition updates to FortiGate Antivirus Firewalls worldwide. Fortinet's FortiProtect Center Web portal provides up-to-the-minute information that enables Fortinet customers to stay on top of the latest security threats and the FortiProtect Distribution Network provides FortiGate Antivirus Firewalls with the data necessary to detect and prevent new threats. The system consists of three key components:

FortiProtect Center

The FortiProtect Center, available at www.fortinet.com , provides a complete overview of current network threats, information about specific viruses and vulnerabilities, and detailed definitions of the threats covered by the latest FortiGate virus and intrusion databases. The FortiProtect Center information portal is updated daily, is easy to read and navigate, and reflects Fortinet's strong commitment to staying up-to-date with the latest security threats.

FortiProtect Threat Response Team (TRT)

The TRT is a pivotal component of the FortiProtect infrastructure. The TRT's security experts research and develop the information for the FortiProtect Distribution Center and work around the clock monitoring emerging threats. This experienced team of network security specialists, led by Joe Wells, founder and president of the Wild List organization, collects and analyzes virus samples and develops virus signatures to update the current Fortinet antivirus (AV) definitions. The team also develops network vulnerability signatures and updates Fortinet network intrusion detection and system (IDP).

FortiProtect Distribution Network (FDN)

Taking the virus and intrusion signature definitions developed by the TRT, the FDN provides automated, timely and reliable updates, ensuring that FortiGate Antivirus Firewalls worldwide have the most current AV and NIDS protection available.

Summary

Complete Content Protection is the optimal solution for combating today's content-based network security threats. With the FortiGate Antivirus Firewalls, businesses can fully secure their proprietary information and intellectual property from destruction, inaccessibility or theft. By reassembling and analyzing content using the FortiASIC processor and FortiOS system, the FortiGate Antivirus Firewalls provide a level of security and performance unmatched by other systems.

More Information

Please visit our web site at www.fortinet.com or email us at info@fortinet.com.