



Technical Brief

NVIDIA Firewall

PC Security and Hacker Defenses





PC Security and Hacker Defenses

Introduction

Computers are an integral part of everyday life for business and pleasure. They contain all sorts of valuable information, making them prime targets for hackers. This reality is why computer security is one of the most important issues facing us today.

Computer security has three independent components: a firewall, intrusion detection, and antivirus protection.

A firewall is the core component of a computer security solution. It ensures that only data packets that comply with the defined policies get through. To provide this protection, the firewall examines each data packet that attempts to pass through, and determine if the packet has permissible attributes; if it doesn't, the packet is blocked. When firewall functionality is integrated into the driver software on the PC, unauthorized access to the PC by an intranet-based or Internet-based intruder is dramatically reduced.

NVIDIA® Firewall is the first firewall powered by the NVIDIA ActiveArmor™ secure networking engine. As a result, NVIDIA Firewall has the highest system performance at the lowest CPU utilization. At the same time, it enhances overall security by providing hardware-based deep packet inspection, instant-on protection, and tamper-resistant functionality.

Firewalls

Their Purpose

Network data is composed of *packets* whose headers contain meta-information. This meta-information allows the packet to be delivered across a subnetwork (Data Link layer header), across an internetwork (Network Layer header), and into the correct process in a host (Transport Layer header). When a machine is attached to the Internet, any other machine on the Internet can send a packet to it if the remote machine knows the target machine's IP address.

Most packets are harmless, but occasionally someone attempts to send packets that exploit bugs in the target host's protocol software or operating system. The purpose of these packets is to disable the host (known as a "denial of service" attack) or gain unauthorized access to the host.

Most corporate and home networks have a well-defined connection to the Internet. The connection consists of a limited number of connection points (DSL modem) through which packets from inside hosts can reach the Internet, and vice versa. To control which packets cross this boundary, the concept of a *firewall* was created.

How They Work

Firewalls enable network traffic to be filtered, based on a variety of criteria. The most obvious way to filter traffic is by packet type. By using the TCP or UDP port numbers in a packet, the firewall permits or denies a packet, based on rules stored in an access control table.

There are two possible scenarios for a packet-filtering firewall:

- ❑ A firewall can allow everything except a list of packets (identified by port numbers) that are considered harmful and will be dropped.
- ❑ A firewall could be programmed to block everything by default, only allowing certain safe packets to cross.

Security is all about managing risk. By defining the configuration of a firewall, users limit their risk to packets allowed into their network. In general, firewalls can be configured, so it's difficult for an attacker to determine what traffic is allowed through the firewall. This protection helps maintain a degree of stealth for the computer being protected.

Types of Firewalls

Stateless Firewall

The stateless firewall is the most basic type of firewall and has existed in one form or another since the early 1990s. In this type of firewall, a list of permit/deny rules is defined so that only packets that match a permit condition are allowed through the firewall. The rules can filter inbound and outbound traffic based on Ethernet type, IP Source or Destination Address, IP options, IP Protocol, ICMP type and code values, TCP or UDP Source or Destination port, and TCP options.

If the packet passes this test, it may pass through; otherwise it is dropped. However, every packet will be subjected to the same barrage of tests. The scaling problem in this design is that every packet must be checked against all the rules. As more and more rules are added, it takes more effort to process each packet. This additional effort reduces performance, as measured in packets-per-second, or as measured in CPU utilization to process a given amount of traffic. Stateless firewalls are most suited for certain packets, such as ICMP, which are stateless in nature.

NVIDIA Firewall supports stateless inspection. It can filter traffic based on the Ethernet type, the IP Protocol, and IP and TCP options rules. IPv4 and IPv6 are treated equally, whenever it's applicable. For example, IPv4 options and IPv6 extension headers may both be used as filter elements.

Stateful Firewall

A stateful firewall is a variant of the stateless firewall. It behaves much the same as a stateless firewall when a new connection is established because it compares the new protocol (and the source and destination of the packet) against its local policies.

The optimization in the stateful firewall is that the packets of a given flow are only examined in detail when a connection begins. Whenever a new connection is permitted, an entry is added to a connection state tracking table. Future packets that match this connection table entry can be verified against the table of permitted connections, without needing to compare each packet against the entire set of rules. The benefit of a stateful firewall is that it offers all the security of a packet-filtering firewall, but for a fraction of the CPU cycles.

NVIDIA Firewall supports stateful inspection of TCP and UDP traffic. A UDP “state” is determined by observing new UDP packets and creating states only if they pass the firewall policy rules defined by the user.

The lookup technique involves computing a hash value based on several key fields in the packet header. Key fields might include the source and destination IP addresses, the IP protocol (which indicates whether TCP or UDP or some other transport layer protocol is in use), and the source and destination transport layer ports. Computing a hash function over these five values takes a fixed (small) amount of time on a per-packet basis.

The complexity of the policy rules for the firewall does not affect the packet validation speed of the firewall. By comparison, the stateless firewall must apply all its rules (or enough rules so that a definitive go/no-go decision can be reached) for each packet. Plus, its packet analysis time increases linearly as the number of policy rules is increased, resulting in packet-forwarding performance that decreases linearly as the number of rules increases.

Application-Level Gateways

An application-layer gateway, or transport-layer bridge, is a special-purpose computer that runs proxy services for each application allowed to pass through. These proxy servers must be exceptionally stable and bullet-proof; otherwise, the proxy server will have its own vulnerabilities. No packets ever directly pass through an application-layer gateway. After a packet is received, all its headers are stripped, the contents are examined, and a new series of packets is created on a new connection to the destination host.

An application-layer gateway is just as transparent as the packet-filtering firewall, except there may be more delay in the examination process. The advantage of this approach is that there is a logical “air gap” between the two networks, but only for protocols the gateway understands.

The biggest limitation of the application-layer gateway is that in order for a certain type of traffic to pass through, a proxy server must exist for that protocol. Proxies for common protocols like SMTP, FTP, HTTP, and TELNET are readily available, but proxies for more exotic protocols may not be available. For limited applications, though, these gateways are the best choice to ensure that only valid data is allowed through the firewall.

Application-layer gateway firewalls are usually at the edge of the network and require dedicated hardware. NVIDIA Firewall, an end-point firewall, does not support application-layer gateway functionality.

Firewalls As Hacker Defenses (Antihacking)

A spoofed IP packet has an illegally generated value in its IP Source Address field. By intentionally using an incorrect IP address, a hacker can build certain kinds of attacks. The most notorious is a distributed denial-of-service (DDoS) attack, which is also one of the most common types of attacks that use IP spoofing. These DDoS attacks depend on two things: 1) an Internet-connected “zombie” device, often a PC, that has been compromised; and 2) the ability to command the zombie PC to send packets with spoofed IP source addresses.

Firewalls have always been able to filter based on an IP address, but the detection of spoofed packets involves a more subtle distinction. For example, based on a given packet’s IP source address, should that packet have arrived on the interface that received it, given what the firewall knows about the routing table? An intermediate device cannot easily detect that a given packet is spoofed.

The best approach to preventing spoofing is to block spoofed packets at their source—the zombie PC. By embedding the antispoofing capability directly into the PC’s networking hardware/software infrastructure, you prevent the PC from using any IP address other than its statically assigned address or its DHCP-assigned address.

Other Important Security Capabilities

The firewall provides one “layer” of protection, and it is usually considered the foundation layer. However, a complete security solution is multilayered.

These additional capabilities are not provided by NVIDIA Firewall, but may be obtained by selecting the best-of-breed components according to the user’s requirements.

Anti-Intrusion Protection

Intrusion detection is the capability to analyze all incoming traffic for patterns of behavior that correspond to known attacks or to precursors of known attacks. For example, in order to attack a vulnerable piece of network application software, an attacker may first scan all possible ports in search of a known example of a vulnerable piece of software. Thus, detecting a “port scan” may indicate that an attack is about to happen, and defensive measures may be taken before any damage occurs.

With prevention intrusion, however, various known attacks are directly detected and thwarted before they can harm a system.

In both cases, the anti-intrusion software depends on access to a library of known attacks. These products are usually not able to detect new attacks, because no “signature” has yet been established for that attack.

Antivirus Protection

Antivirus capabilities protect a user's PC from executing code that has known viruses or Trojans. As in the case of anti-intrusion software, antivirus products are based on a library of attacks that a product knows how to defend.

In addition, certain antivirus software can alert users to suspicious activity, even if it does not match a known virus.

NVIDIA Firewall

NVIDIA Firewall is now powered by the ActiveArmor secure networking engine, which makes it the industry's first true hardware-based PC firewall. Because of the secure networking engine, NVIDIA Firewall does not rely on the CPU for any overhead.

The NVIDIA ActiveArmor secure networking solution—a combination of NVIDIA Firewall and the ActiveArmor secure engine—improves network throughput to full gigabit Ethernet speeds, lowers CPU utilization, performs deep packet inspection, and improves overall network security (Figure 1).

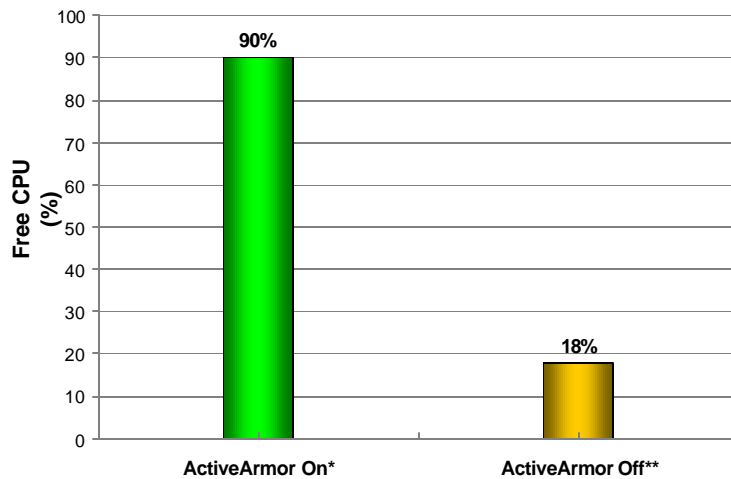


Figure 1. NVIDIA ActiveArmor Delivers Maximum Performance and Minimum CPU Usage

NVIDIA Firewall incorporates both firewall and antihacking technologies. And, it supports stateless and stateful inspection, Web-based management, predefined security profiles, port block filtering, Intelligent Application Manager, remote administration, and an easy-to-use wizard. In addition, NVIDIA Firewall has antihacking features such as anti-IP-spoofing, antisniffing, anti-ARP cache-poisoning, and anti-DHCP server—important security controls for corporate network environments.

In a corporate setting, an end-point firewall (such as a desktop firewall) with antihacking capabilities can reduce the internally originated security breaches, and can inhibit desktops from generating unauthorized traffic. The result is improved overall security, with reduced requirements from the IT staff.

Advanced Management Features

NVIDIA Firewall offers many advanced management features such as remote access, configuration, monitoring, command line interface (CLI), and WMI scripts. And, it is easy to set up through a user-friendly wizard.

These advanced management features make NVIDIA Firewall flexible, easy to use, and very powerful (Figure 2).

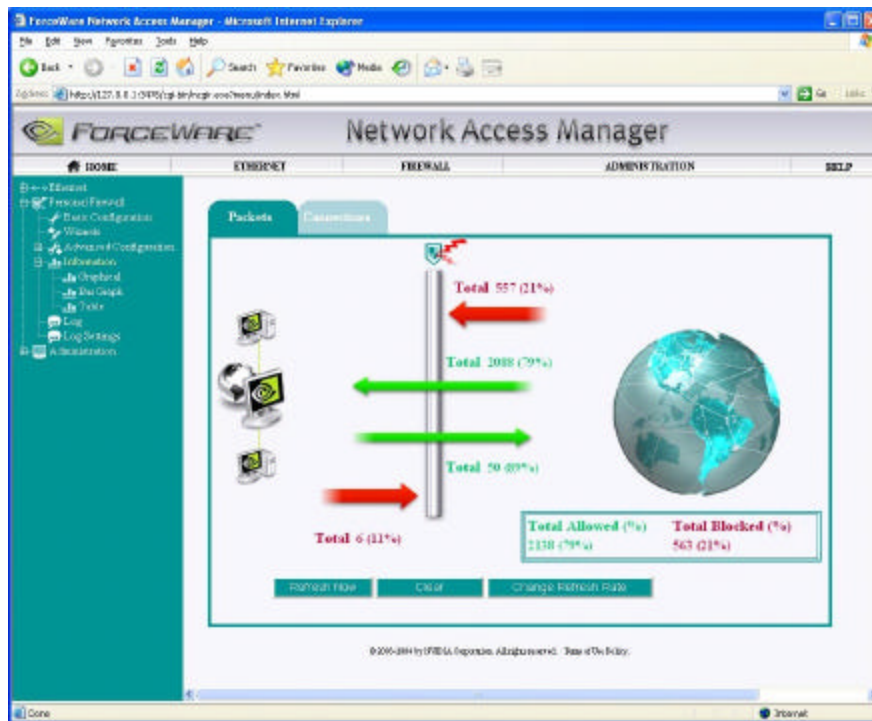


Figure 2. Easy Configuration with Web-Based Browser Interface

Intelligent Application Manager (IAM)

The Intelligent Application Manager is an addition to NVIDIA Firewall that adds application-based filtering to the already comprehensive firewall filtering capabilities. IAM extends the NVIDIA Firewall policy management elements to provide filtering based on applications, whether they are acting as clients or servers. IAM puts users in the loop, enabling them to decide what's safe to let in or out of their computer. Once an application is allowed, it can open up ports without specific configuration by the user (Figure 3).

IAM eliminates the possibility that a rogue application on the user's PC could send out traffic that happened to pass through the firewall; the outbound traffic is only permitted if it comes from an application that the user deems safe. IAM can even track existing applications and determine if they have been altered—for example, by a virus or Trojan attaching itself to the executable, or by an application renaming itself to imitate a known application.

IAM is also useful for protecting the PC against incoming packets. It limits the ability of Trojans or other Spyware to set themselves up as servers on the PC, preventing them from receiving traffic from outside the PC. Not only is IAM able to filter based on ports, it can also prohibit the server from opening any sockets, effectively preventing it from receiving any traffic at the application layer.

IAM offers full protection against attacks, protecting the PC from being attacked by outside entities as well as preventing the PC from attacking other PCs.

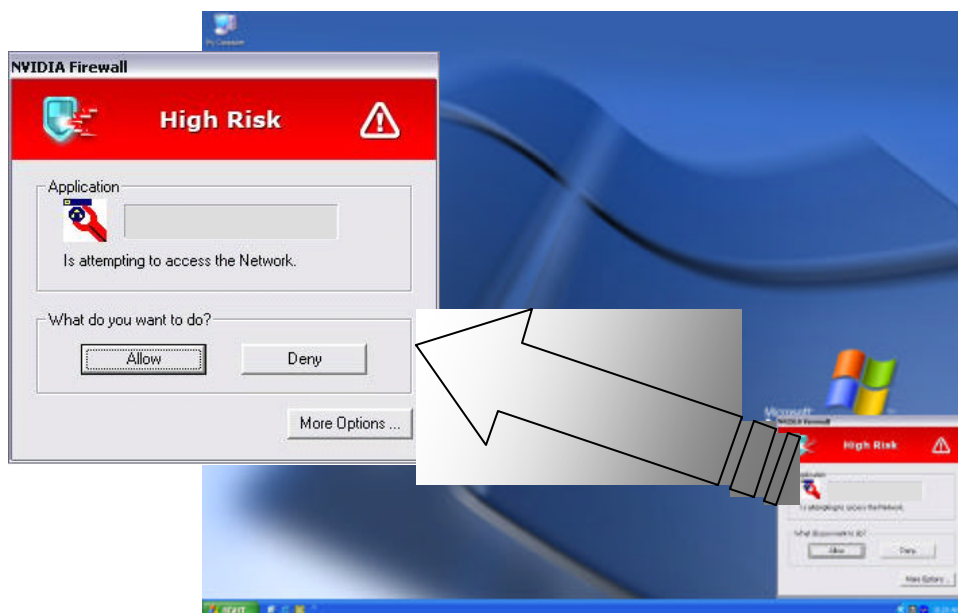


Figure 3. IAM Alerts You When Unknown Applications Attempt to Access Network

Why Choose NVIDIA Firewall?

Most PC firewalls are software-based add-ons, but NVIDIA Firewall is the industry's first true hardware-based PC firewall. And NVIDIA's ActiveArmor secure networking solution, which consists of NVIDIA Firewall and the ActiveArmor engine, improves overall network security.

NVIDIA Firewall has unique features. It offers the Intelligent Application Manager (IAM), advanced management functionality—remote access, configuration, and monitoring—and is easy to use and set up through a user-friendly wizard. Plus, it can be deployed in corporate settings, including an end-point firewall such as a desktop firewall. Or it can be used in home settings, as when the PC is attached to a broadband Internet connection, to protect the PC from unauthorized access.

NVIDIA Firewall technology may be used as a powerful baseline policy enforcer. For full protection, users should augment the protection of the NVIDIA Firewall with leading antivirus and intrusion-detection software to achieve the best total PC security solution



Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, and ActiveArmor are trademarks or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2004 by NVIDIA Corporation. All rights reserved.



NVIDIA.

NVIDIA Corporation
2701 San Tomas Expressway
Santa Clara, CA 95050
www.nvidia.com