Interested in learning more about security?

SANS Institute InfoSec Reading Room

3Com Distributed Embedded Firewall

As the Internet community becomes more skilled in their use of attack tools, we are seeing an increase in the number and severity of Internet attacks. Internet neophytes and professionals ali available solutions. It is, however, not without limitations. This paper will review the benefits and limitations of this product.

3Com® Distributed Embedded Firewall

A Review

Kyle Kelliher

GSEC Assignment

Version 1.4b Option1

April 14, 2004

Abstract

The 3Com® distributed embedded firewall product is a unique approach to an enterprise class host based security solution. 3Com® has chosen to locate the packet inspection at a different point in the network path than other host based firewalls. The product offers some significant advantages over other available solutions. It is, however, not without limitations. This paper will review the benefits and limitations of this product.

Table of Contents

Introduction

Need for Host Based Firewalls

In the ongoing war, between security professionals and hackers, the stakes keep getting higher. In the early years of the war, it was sufficient to have a network perimeter firewall to protect one's network. In this day and age, perimeter firewalls alone are no longer acceptable. Today's networks no longer have clearly definable perimeters. The advent of wireless access and extranets have blurred where one organizations ends and another starts.

Even if organization and network perimeters could be clearly delineated, the security industry is constantly abuzz with the latest firewall vulnerability that has been discovered. It is naïve to assume that perimeter firewalls will ever reach a stage where they in and of themselves will provide an acceptable level of defense. It must be assumed that at some point the perimeter defenses will be comprised.

Even if the perimeter defenses are not comprised, what about the attack that begins from inside? Rogue employees and lax physical security account for far more security breaches than perimeter firewall failings. Indeed, fully, 70%-80% of security attacks now originate from with inside the network (Carr). "The perimeter firewall doesn't protect you from the bad guys inside the network," says Raphael Reich of Check Point Software." The lessons, which have been learned from these attacks, are that perimeter security is no longer satisfactory in and of itself; defense-in-depth (NSA) must be practiced.

Defense-in-depth is the 'don't put all your eggs in one basket' strategy of security. Just because someone **should not** ever reach a certain point in a network does not mean that they **will not** reach that point. The more layers an attacker has to penetrate, the less likely he or she is to reach the target.

Once an attacker has intruded upon a network, he or she should not have free range of all of the computing resources. Protecting individual components of the network serves to reduce the severity of an attack. Even if someone has intruded a network, if the individual machines on the network are independently protected, there is less risk of vulnerability.

This individual protection can be provided through the use of host-based firewalls. Host based firewalls comprise an integral part of a sound defense-in-depth strategy. They serve to limit the damage that can occur when perimeter defenses are compromised.

Many approaches have been developed to provide robust, secure, easily manageable host based firewalls.

Personal, software based, host firewalls

One of the first approaches to host based firewalls is the personal software based firewall. There are many personal software host based firewalls available; examples include products from BlackICE®, Zone Alarm®, and Windows® Internet Connection Firewall.

Personal software firewalls suffer from a number of limitations. They lack centralized management functionality, preventing corporate security staff from remotely configuring the policies and auditing the activity of the firewalls. This limitation makes them more suited to the small office / home office environment than to the enterprise environment.

In addition, they are often too confusing for the average end user to correctly configure and maintain. Also, end users frequently deactivate personal firewalls, either intentionally or unintentionally. An inactive firewall is potentially more dangerous than no firewall at all, because the perception of protection is there. A firewall is only as secure as the policy it is implementing.

Software based firewalls are comparatively simple for attackers to disable. All of the code that is doing the packet filtering and inspecting is resident on the machine being protected. This allows attackers to use rights that they have gained to get into the machine to also disable the firewall protection.

Software based firewalls, also place some load on the machines they protect. They all in some way or another replace or put hooks in the network stack to allow them to look at all packets coming and going. This additional processing takes CPU cycles.

Distributed software based host firewalls

In attempt to address some of the limitations associated with using personal software firewalls in corporate environment, a number of companies have come out with software based distributed firewalls. Examples of companies producing such products include F-secure® , Cyber Armor®, and Check Point™.

These products address some of the shortcomings of the personal firewalls, by providing for a centralized method for enforcing policies and auditing performance. They work in a client server fashion, with one (or a group) machine controlling many workstations.

In general, they leave little or no control to the workstation user, overcoming the user's tendency to deactivate anything which might be considered a nuisance. They also report there activity to a central location, allowing security staff to 'see' what is going on.

This approach still suffers from some limitations. The code that is doing the protecting is still resident on the machine being protected. The protection only extends as far as when the firewall releases the packet to the link layer. There is still an opportunity for an attacker to deactivate or circumvent a software-based firewall, particularly if he or she has administrative privileges on the machine.

There is still a load imposed on the processor of the machine being protected. While this load is slight, in certain applications, it must be considered.

Stand alone distributed hardware based firewalls

Another approach to host based protection is to use separate, stand alone hardware devices to do the packet filtering and inspection. Companies producing products in this space include Lucent®, Nortel®, and Sonic Wall®.

These products address many of the limitations of software based firewalls, not however without presenting some of their own. Similar to the distributed software firewalls, these products are managed from a central console. This allows security ease of configuration and monitoring. As these firewalls are separate machines using separate processors, they do not place any load on the machine being protected.

In addition, being distinct from the host, there is less possibility of the end users trying to fiddle with the controls. However the potential exists, that an end user or an attacker could simple unplug the workstation form one side of the firewall and plug it in to the other side of the firewall, rendering the firewall useless. Also, when a mobile user takes their laptop to an internet café, the stand alone firewall back in the office is not doing much good.

A possible answer

3Com Corporation has developed a family of products, which address many of the limitations of other distributed, and host based solutions available today. This family of products is known as the 3Com embedded firewall.

The major difference between the 3Com® offering and other offerings is that the packet filtering is built into the hardware network adaptor. It is neither a separate stand-alone device nor a software-based firewall. This solution offers some distinct advantages, but also has some significant shortcomings.

The Product

Overview

The 3Com® embedded firewall product was initially released in 2001. The embedded firewall has similarities to both software and hardware based firewalls. The product works similar to many other distributed hardware firewalls in that there is a centralized location for policy management and audit data collection. It works similar to the software firewalls, in that the firewall is an integral piece of the protected machine. However rather than using software or stand-alone hardware devices to do packet filtering, the packet filtering is done by a processor built into special network interface cards.

The product family consists of several different embedded firewall network interface cards and centralized management software, called a policy server. Each protected machine requires an embedded firewall network interface card. These cards all receive their policies from and send their audit data to a policy server.

Upon startup, the network interface cards communicate, via UDP, with the policy server software, to determine what policies, they should currently be enforcing. In addition to checking with the policy server for rule sets, the cards also send data back to the policy server for any event that the policy server has asked the cards to track. If on startup, a network card cannot reach a policy, a default fallback policy is enforced. Communication between the policy servers and the devices is secured by using IPSEC.

Network Interface Cards

3Com makes several different PCI Cards available in the embedded firewall product line. Both twisted pair and fiber adaptors are available. There are also two different levels of adaptor available, desktop and server. It is also possible to buy upgrade licenses to allow you to install the embedded firewall firmware on existing 3Com 990 family NICs.

The server versions apparently have more onboard memory allowing them to handle large, more involved rule sets. Information available on the 3Com website in regards to the difference between desktop and server PCI cards was sketchy. Resources were unavailable for this review to test whether there is a performance difference between the desktop and server models.

3Com also makes available both type 2 and type 3 PC cards for mobile users.

Policy Server Software

The policy server software is the nerve center of the embedded firewall system. It has two primary functions, creating and distributing policies and collecting and analyzing audit data. It runs on top of an underlying MySQL® database.

To manage the distributed firewall, the network interface cards, referred to in the policy server as devices, are arranged into groups called device sets. Stationary, devices, i.e. a desktop or server PCI card, are members of only one device set. Roaming devices, i.e. a PC card, are members of two device sets, one is a 'home' device set, for example, for when the mobile computer is with inside the corporate perimeter firewall, and the other device set is 'roaming', for when the mobile computer is exterior to the corporate network. The policy server will automatically sense whether the PC cards are local or roaming, based upon criteria defined by the administrator.

Device sets are directly mapped to policies. Policies are made up of rules or groups of rules, called rule sets. A limited number of predefined rule sets are included as a starting point. Rules are standard access control list style allow or deny functions. Rule processing is in a straight top down fashion. Both ingress and egress filtering is supported.

Each policy server has a capacity to handle 1000 devices. A maximum of 3 policy servers may be used in a domain. This limits a single domain to no more than 2000 devices.

Usage

There are many scenarios in which the embedded firewall could be utilized. As shown in Figure 1, in a data center, the PCI NICs could be used on critical workstations and servers to prevent an attack from all services off-line. For a mobile sales force, the traveling laptops could be protected with embedded firewall PC cards to prevent roaming users from jeopardizing the corporate network. Effective policies, in this case could prevent roaming users from getting malware while on the road and prevent that malware from attacking other machines on the corporate network when they come back to the office. Roaming policy could also be configured to only allow a mobile machine to make a VPN connection to the home office when taveling.

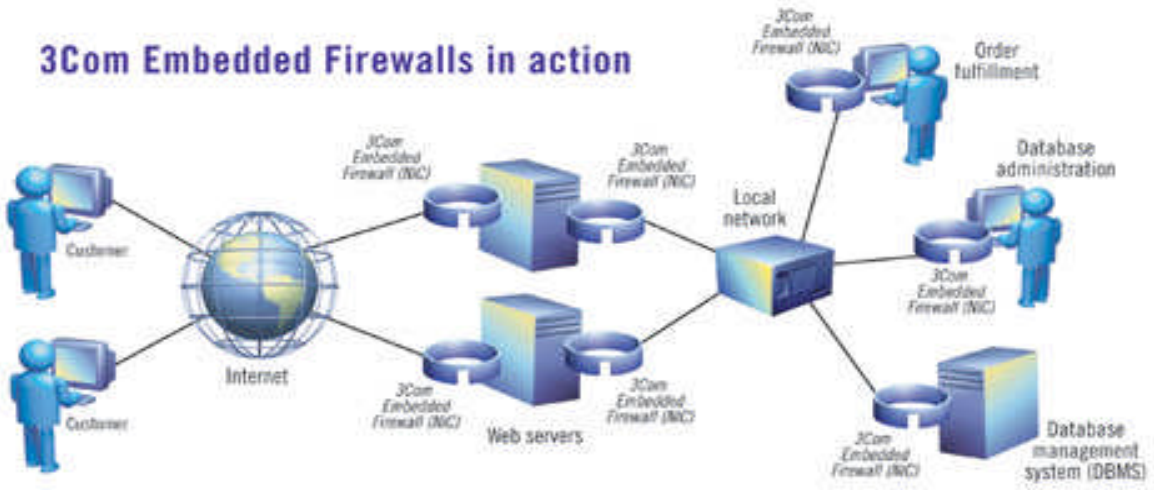

Figure 1

A typical data center configuration, Courtesy of Secure Computing

Installation & Operation

Device Registration

Installation did not go seamlessly. After installing the firewall firmware on one interface card, the card would not register with the policy server. After discussions with 3Com tech support, the card was replaced at no cost. The replacement card immediately registered with the policy server, as is seen in Figure 2, the device information page from the firewall management console.

Outside of the difficulty with the one interface card, installation was quite easy. However, care must be taken to proceed in exactly the order specified in the handbook. It is quite easy to render the interface inoperable by such things as installing the diagnostics after the firewall firmware.

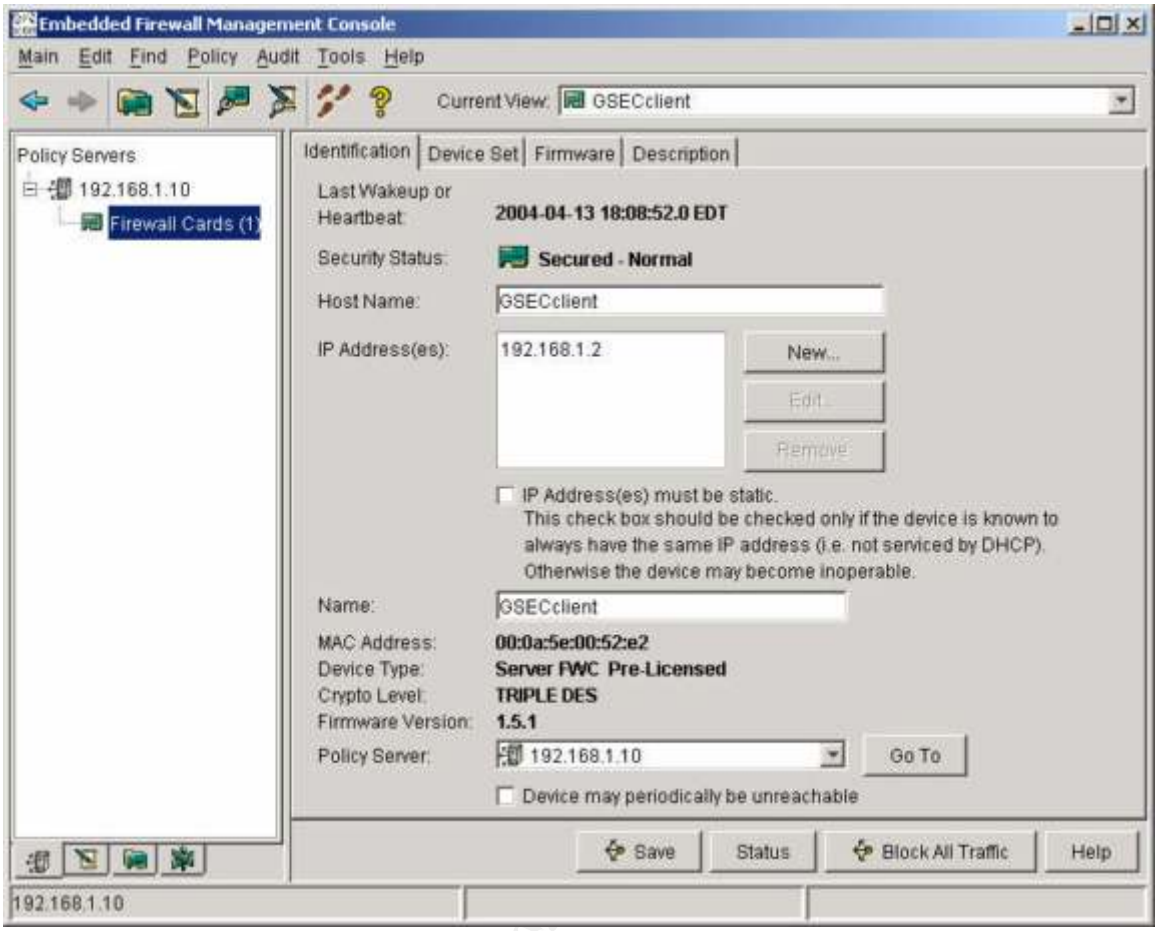

Figure 2

Management Console showing Stationary Device Information

Policy and Rule Definition

Defining rules is straightforward. It is quite similar to writing rules for other firewalls or router access control lists. Figure 3, gives an example of the policy-editing console. New rules are written via wizard, which walks the user through creating rules. There are convenient buttons for adding and deleting rules, combine rules in rules sets, and importing and exporting rule sets and policies.

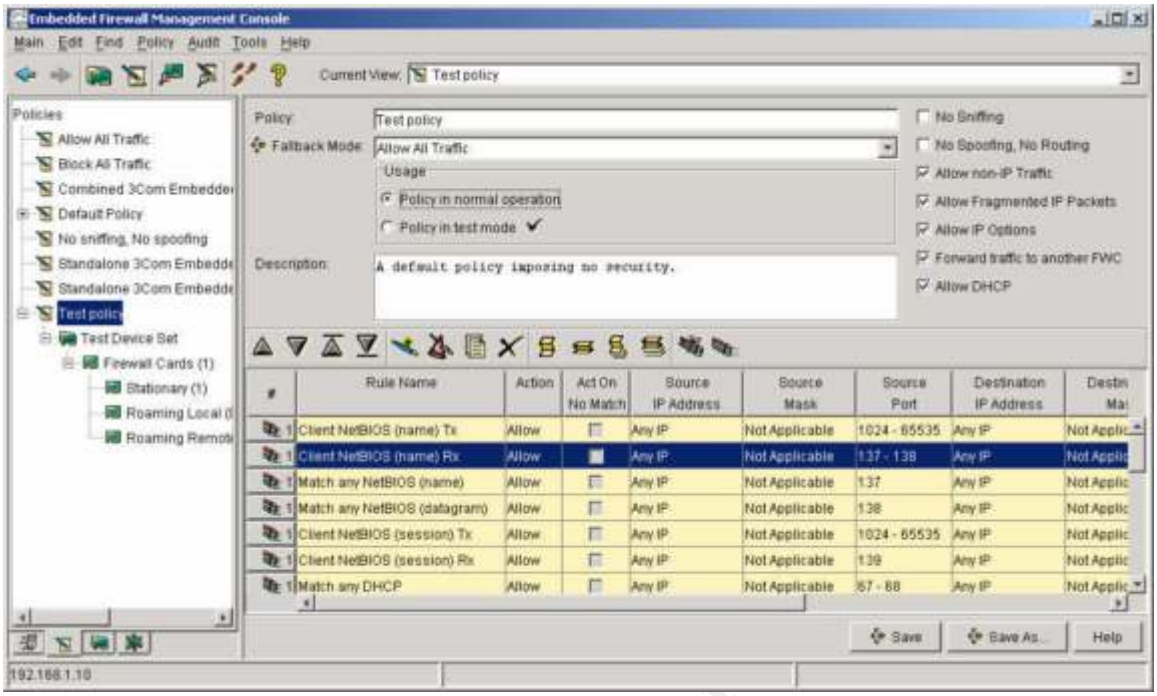

Figure 3

Policy editing console

Policies, rule sets, and rules are exported in a straightforward XML format, as is shown in Figure 4. It is certainly possible to export the policy in XML format, edit the policy in an XML editor, and import the policy again, instead of using the built-in policy editing console. The policy still must be imported to apply it to the device set and distribute it to the cards.

As the 3Com does not support stateful packet inspection, explicit rules must be defined to allow ingress and egress traffic on all ephemeral ports. This is seen in Figures 3 and 4 by the source and destination port ranges 1024-65535.

```
<Rule name="Client HTTP Tx" sourceHostID="Any IP"
sourceMask="255.255.255.255" sourcePortRange="1024-65535"
destinationHostID="Any IP" destinationMask="255.255.255.255"
destinationPortRange="80" direction="out" action="Allow"
ipProtocol="tcp (6)" enabled="true" audit="false"
testMode="false" ruleNegated="false"
allowTcpConnectInit="false"><Description>Generates HTTP
requests.</Description> </Rule>

<Rule name="Client HTTP Rx" sourceHostID="Any IP"
sourceMask="255.255.255.255" sourcePortRange="80"
destinationHostID="Any IP" destinationMask="255.255.255.255"
destinationPortRange="1024-65535" direction="in" action="Allow"
ipProtocol="tcp (6)" enabled="true" audit="false"
testMode="false" ruleNegated="false"
allowTcpConnectInit="false"><Description>Accepts HTTP
replies.</Description></Rule>
```

Figure 4

Rules in native XML Format

Audit Browser

The audit browser is also accessible from the firewall management console. In the policy editor, individual rules can be specified to be audited. If a rule in a policy is marked as audit or test, any interface cards that are members of device sets connected to that policy will send events to the policy server whenever a packet matching the rule is encountered.

Figure 5 shows an example of the browser that is used to view the events received from the firewall cards. The browser offers a number of options for sorting and filtering the data. The complete list of fields that the policy server collects from the firewall cards follows:

> Unique ID, Date & Time, Audit Code, Category, Device, Policy Server, Device Set, Policy Name, Policy Version #, Rule #, Source MAC Address, Destination MAC Address, MAC Type, IP Protocol, Source IP Address, Destination IP Address, Source Port, Destination Port, TCP Flags, ICMP Type, ICMP Code, Test, Action, Packet

The audit browser allows the audit data to be exported in comma separated value format for use with other analysis tools.

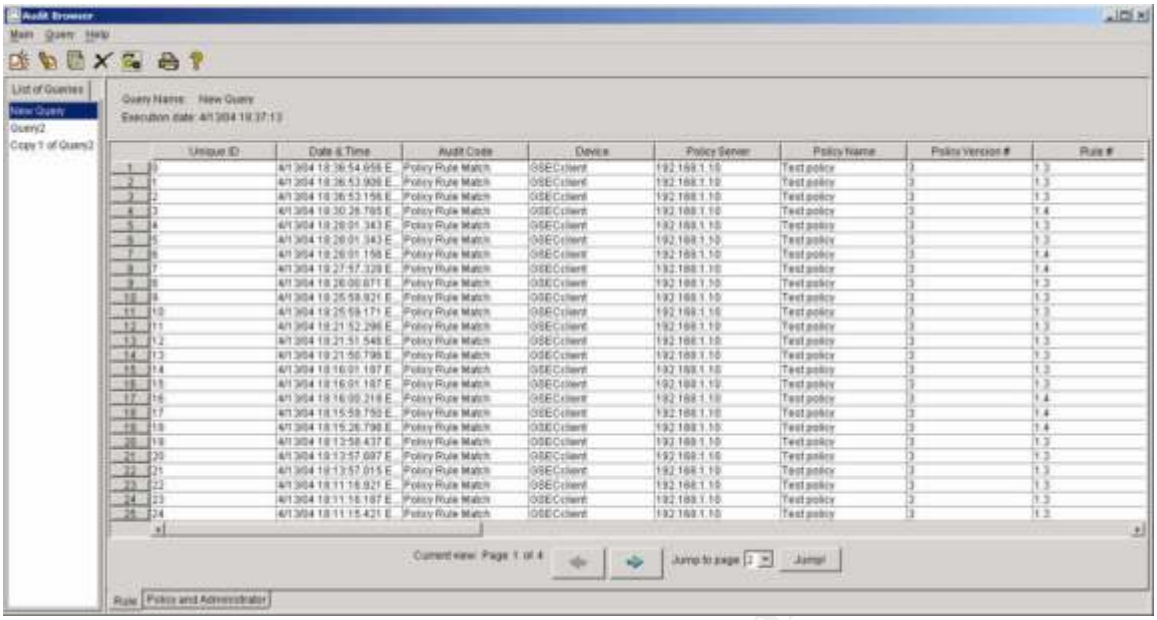

Figure 4

Sample Audit Browser Query Output

The Benefits

3Com's embedded firewall offers some significant benefits.

There has not to date been a great a security solution for road warriors. 3Com's idea of having the PC card auto sense whether the laptop is 'home' or roaming and applying an appropriate policy makes life easier for everybody. The user doesn't have to worry about configuring different settings each time he plugs into a different network. The administrator can be comfortable that the user won't be able to disable the firewall and that that laptop will not bring malware into the corporate network.

In addition, to the second device set and second policy for roaming cards, there is also a default policy which will go into effect if a protected machine cannot contact a policy server.

It is convenient for security administrators and auditors to have the entire log data for all of the devices collected in one place. The data can be easily filtered and sorted using the audit console. In addition, the data can be exported for analysis with other tools.

Administrators don't have to worry every time new operating system vulnerabilities are discovered about whether their machines are protected. The interaction between the operating system and the firewall NIC is very minimal. (The cards do look at an initial configuration file to determine what policy server to contact at startup.)

It is easier to write rule sets for distributed firewalls. With the policy server, rules are only written for one or a small group of machines at a time, It is easier to mentally parse a small set of rules to determine if only what is needed is being allowed than to parse the rules for everything that might be passing through the perimeter.

The policy server also allows applying the policy in test mode. In this mode, audit data is collected to show what would happen if the policy went live but nothing is actually blocked. This is very handy to be able to give a policy or rule a dry run, to make sure critical services will still function.

The policy cannot be intentionally or inadvertently disabled or altered in any way from the protected computer, only from the policy server. The policy server also offers the ability to disable all policy enforcement instantly. It is questionable whether this is really a desirable feature or not.

As the embedded firewall runs on independent hardware, there is no loading on the protected system's CPU. This should have a positive impact on overall throughput, particularly, for high volume servers. The firewall interface cards can also off load all IPSEC encoding from the main system processor. This should also have a positive impact on overall system performance.

The Limitations

3Com's embedded firewall has some significant limitations.

At present, neither the interface card firmware installation nor the policy server software is available on any operating systems other than Windows®. Particularly in large data centers, where there are frequently mixed environments, this will limit the products' appeal. Having to use one set of products for your Windows® systems and a different set of products for your *nix systems, adds needless complexity. Additionally 3Com only supports running the policy server software on Windows® Server operating systems.

There are at present no wireless adaptors available in this product family. With the rapid expansion of public wireless access points, road warriors will not be content with conventional wired adaptors.

The firewall is only able to filter IP packets. Other protocols can be rejected or ignored, but only on an all or nothing basis. The firewall does not perform any stateful packet inspection. Rules must be specifically written to allow outgoing traffic on all ephemeral ports. Lack of stateful packet inspection may cause protocols, such as FTP and other that include port information in the payload to not work correctly.

If logging is set to a high level, and there are a large number of clients, the audit data traveling from the interface cards to the policy server will consume a lot of bandwidth.

As the industry is rapidly moving to a 'log everything' stance, this could present a problem.

As all of the audit communication (and most other communication) is UDP based, there is no certainty that an audit event will ever reach the policy server. A flooding attack or network outage could prevent a policy server from collecting audit data for a NIC, potentially masking an attack against the NIC.

It would seem that some sort of a polling system for the audit data would remedy both of these problems. It would however require more on card memory.

It is possible to physically remove the firewall interface card from the protected machine and thereby the associated policies. It must be certain, that any other interface cards are disabled and users do not have local administrator rights to install other cards.

Given that the policy and audit server is based on the open source database MySQL®, it would seem that the policy server and audit logs are susceptible to any new MySQL® vulnerabilities that arise. 3Com does not seem to have provided a mechanism for administrators to keep their MySQL® patched without waiting for 3Com to publish patches or new versions of the policy server.

Perimeter firewall ports must be opened to allow for communication between exterior cards or DMZ cards and the policy servers. Alternatively, you can place a policy server in the DMZ but then you have to open ports for policy server to policy server communication.

The relationship between devices, device sets, and policies is strictly one-to-one. A device can only be a member of one device set and a device set can only have one policy applied to it. This seems to be somewhat limiting. A system where device could have membership in numerous groups each having their policy would b more flexible. Suppose it was known that a given machine needed access to servers x and y, but not z, the device could just be added to the x and y groups but not to the z group.

It is possible to export the audit data in a comma separated value format. There, doesn't seem to be anyway to automate the export of the data for use with a security management console.

Testing

Running port scan, packet sniffers, and the like against the firewall showed the firewall to work as advertised. It correctly dropped and logged any packets that were denied and audited. It passed any packets that were allowed. There certainly were not any obvious faults in the policy enhancement. As the product has not been on the market long, there are no published vulnerabilities.

Further work may reveal some vulnerability. Areas to consider for testing are:

- Attempting to spoof a policy server
- Fragmentation attacks
- Overloading the NIC processor or memory
- Attacks using other protocols than IP
- Attacks against the MySQL database.

Conclusion

The 3Com distributed firewall has some positive points and some negative points. In certain situations, the positive points will outweigh the negative points; in others they won't. It seems that 3Com intends the product to be used by a number of different types of customers.

3Com is clearly trying to sell to the enterprise with many mobile users. They have a clearly thought out strategy for mobile users. However, the solution still depends on administrators to lock the machine downs to make sure that users can't install other interface cards in the machine. Additionally, the absence of any wireless adaptors in this product family makes the product significantly less attractive. With the rapid expansion of wireless access points in conference centers and hotels, mobile users are going to be increasingly resistant to not have wireless networking available.

For the enterprise data center, however, the lack of wireless adaptors is a non-issue, as is the ability for users to install different interface cards (One can assume in a data center that there would some level of communication between the networking staff and the security staff). Even the lack of stateful packet inspection could be potentially not an issue in a data center where information flow patterns are well defined. The lack of support for operating systems other than Windows® could be an issue however.

For all customers, 3Com needs to improve the access to the audit data. With the increasing use of security management portals, the ability to aggregate data is crucial.

In the right sets of circumstances, the embedded firewall is an excellent choice as one part of a defense-in-depth strategy. With a few enhancements, it could really be a marked improvement over other options available, today.

References

"3COM Embedded Firewall Solution" PC Magazine . April 22, 2003 <http://pcmag.dit.net/reviews/review.php?id=EpVpAAFkVkFDbQUUJZ>

3Com Corporation Product Home Page <http://www.3com.com/security/efw_info.html>

Bellovin, Steven M. "Distributed Firewalls". ;login: November, 1999. pp. 37-39. <http://www.research.att.com/~smb/papers/distfw.html>

Carr, Jim. "Strategies & Issues: Thwarting Insider Attacks." Network World. 4 September 2002. <http://www.networkmagazine.com/article/NMG20020826S0011>

Cronje, Gerhard. "Choosing The Best Firewall." April 10, 2001 <http://www.sans.org/rr/papers/index.php?id=951>

Fratto, Mike. "3Com Embeds Firewall in PC NICs. " Network Computing. November 15, 2002. <http://www.nwc.com/1324/1324sp2.html>

Messmer, Ellen. "Second line of defense." Network World . Sans Institute. June 5, 2000 <http://www.nwfusion.com/news/2000/0605defense.html>

National Security Agency. "Defense in Depth *A practical strategy for achieving Information Assurance in today's highly networked environments.*" <http://nsa2.www.conxion.com/support/guides/sd-1.pdf>

Safeer, Daniel L. "Active Net Steward – Distributed Firewall." Sans Institute. July 20, 2001<http://www.sans.org/rr/papers/index.php?id=793>

Secure Computing Product Home Page <http://www.securecomputing.com/index.cfm?skey=953>

Last Updated: November 18th, 2017

Upcoming SANS Training

Click Here for a full list of all Upcoming SANS Events by Location

SANS San Francisco Winter 2017	San Francisco, CAUS	Nov 27, 2017 - Dec 02, 2017	Live Event
SIEM & Tactical Analytics Summit & Training	Scottsdale, AZUS	Nov 28, 2017 - Dec 05, 2017	Live Event
SANS Khobar 2017	Khobar, SA	Dec 02, 2017 - Dec 07, 2017	Live Event
SANS Munich December 2017	Munich, DE	Dec 04, 2017 - Dec 09, 2017	Live Event
European Security Awareness Summit & Training 2017	London, GB	Dec 04, 2017 - Dec 07, 2017	Live Event
SANS Austin Winter 2017	Austin, TXUS	Dec 04, 2017 - Dec 09, 2017	Live Event
SANS Frankfurt 2017	Frankfurt, DE	Dec 11, 2017 - Dec 16, 2017	Live Event
SANS Bangalore 2017	Bangalore, IN	Dec 11, 2017 - Dec 16, 2017	Live Event
SANS Cyber Defense Initiative 2017	Washington, DCUS	Dec 12, 2017 - Dec 19, 2017	Live Event
SANS Security East 2018	New Orleans, LAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS SEC460: Enterprise Threat Beta	San Diego, CAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
Northern VA Winter - Reston 2018	Reston, VAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SEC599: Defeat Advanced Adversaries	San Francisco, CAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, NL	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Dubai 2018	Dubai, AE	Jan 27, 2018 - Feb 01, 2018	Live Event
SANS Las Vegas 2018	Las Vegas, NVUS	Jan 28, 2018 - Feb 02, 2018	Live Event
Cyber Threat Intelligence Summit & Training 2018	Bethesda, MDUS	Jan 29, 2018 - Feb 05, 2018	Live Event
SANS Miami 2018	Miami, FLUS	Jan 29, 2018 - Feb 03, 2018	Live Event
SANS London February 2018	London, GB	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Scottsdale 2018	Scottsdale, AZUS	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Southern California- Anaheim 2018	Anaheim, CAUS	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Secure India 2018	Bangalore, IN	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS London November 2017	OnlineGB	Nov 27, 2017 - Dec 02, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced