Global Information Assurance Certification Paper

Daniel L. Safeer, July 20, 2001
SANS Security Essentials – GSEC Practical Assignment
Version 1.2c
Active Net Steward – Distributed Firewall

" (Firewalls) rely on the assumption that everyone on one side of the entry point, the firewall, is to be trusted, and that anyone on the out side is, at least potentially, an enemy. The vastly expanding Internet connectivity in recent years has called that assumption into question."(1)

Steven M. Bellovin
"Distributed Firewalls", *;login:*, November 1999

Recent studies have proven just how incorrect that assumption is. A Digital Research, Inc. reported, "Authorized users are by far a company's biggest security threat." (3) A study by the FBI and CSI showed 44% of respondents "reported unauthorized access by employees."(3) The report that opens the most eyes is a 1996 study by American Society for Industrial Security that reports, "A massive 75 per cent of all computer break-ins occurred internally."(4) Whether this access was malicious or simple curiosity is irrelevant, this access was possible because it was not stopped by traditional methods: firewalls and IDS. The question then becomes, how do I deal with the implied trust afforded to users who are inside of the firewall, either physically or electronically (via VPN or dialup)?

John Earl, of The PowerTech Group, answers this question quite simply; "Security is the responsibility of the **Target** system." (5) It is not enough to rely on network passwords and good faith when dealing with workstation and host level security. The SANS lesson on Host Perimeter Defense focuses on the idea of protecting the target, but in terms of preventing an outside attack or being used in a man-in-the-middle hack. Telecommuter or traveling user's home PC or laptop being compromised and then used to access corporate networks or dial-up via a rogue modem. Imagine a lost or stolen laptop with a VPN client and a valid SA. The intrusion into Microsoft was possible because a remote user who was connected to Microsoft with a VPN client had an unprotected pipe open to the Internet. Something as simple as a Personal Firewall would have prevented that breach. Study after study has shown that the real threat is the user who already on the network. If Microsoft had some form of Host Perimeter Defense in place, the implied trust that is granted to a remote employee who is using a VPN client would have been in check. However, imagine trying to administer policy on several hundred machines without some form of centralized management.

Personal Firewall products have done a good job of securing home users PC's and mobile user's laptops, but they should not be relied upon within the walls of the office. A main product feature should instantly exclude Personal Firewalls from network use, they allow the end user to change and make decisions regarding security

policy configuration. The same end users, which studies show are the people who are committing the most break-ins. From an enterprise management view, these products do not provide centralized administration, policy definition or reporting. Several manufacturers have come out with what they call “Centralized Management,” but they are nothing more than Personal Firewall products and an additional server product that only performs log scraping of alerts. Alerts notification should be real-time, not hours or days later. These products are named correctly, personal; they are suitable for home or laptop personal use. One manufacturer has even go so far as to propose a four product solution, Personal Firewall on the desktop, Management Console to scrape the logs, a Reporting product to analyze the scraped data and a Policy Management product to distribute and update a security policy that the end user is still able to change.
These first attempts at centralized management and Personal Firewalls are a step in the right direction, but they are lacking many of the features that are required for true enterprise management.

Steven M. Bellovin was the first to promote the idea of a “Distributed Firewall.” In his 1999 paper, he proposed that policy should be centrally administered, deployed via SMS or other management package and enforced at the host level. (1) The goal was to free a network from the topographic limitations commanded by the use of a firewall. Secondary to that was the protection that a distributed implementation afforded against the internal threat. Firewall makers acknowledged the internal threat, but also defended their products; "The perimeter firewall doesn't protect you from the bad guys inside the network," says Raphael Reich of Check Point Software. "But people should not be replacing perimeter firewalls with distributed ones." (2) Unfortunately people are not even adding distributed ones to fill in the gaps left by traditional methods. The ideas that Steven wrote about have become the foundation of a new breed of Host Perimeter Defense products. But, until recently there was no single product that could provide true Distributed Firewall functionality as defined in his white paper. Centrally managed, automatic policy distribution that is controlled at the host level and that is not limited by network topology.

Earlier this year Security Designers, LTD (SDL) released their Active Net Steward (ANS), Distributed Firewall product. Product development began early in 1998 with a Windows NT only proof of concept prototype. The then named NetSteward was purchased by Network Designers, LTD (NDL) in July of 1999 and made a part of their “Active” product line. In October of 2000 NDL took the completion of the newly named Active Net Steward as the opportunity to spin off a security arm, SDL, to sell and continue to develop the ANS product line. (6) In April of 2001, NDL recruited their US business partner, Waytek, Inc. to resell the SDL product line and Datacom Depot, Inc. to act as North American distributor.

ANS meets all of the requirements of a true Distributed Firewall. Policy is created, maintained, policed and distributed from a central location, the policy is automatically pushed to the clients and because of its client-server design there are no limitations on network topology. Administrators are given real-time reporting of all

TCP/IP activity on their workstations and servers. Appropriate use policy can be policed and enforced from a central management console. (4) Network health check information is also collected enabling administrators to resolve issues before serious performance degradation occurs. ANS can also be configured to simply observe and collect, giving administrators the information that is needed to develop a security or appropriate usage policy that is based on actual activity.

ANS is a client-server product made up of the management console and client software. The current version of the product requires that each workstation be visited for client installation and a reboot is required. Future releases will integrate with the Active Directory, allowing the client to be installed remotely. (6) The clients run at the lowest possible level, as virtual NDIS drivers in kernel mode. They see everything going onto and coming off of the network. Running in kernel mode also makes the client more secure and less likely to be interfered with by other processes. (7) There is no user interface or application to disable, uninstall is supported, but is protected with a password. The problem of rogue modems is also addressed by the ANS; the client can be set to disable modem access to prevent the user from bypassing security. The current client only supports TCP/IP traffic across a network interface card. Future releases will police modem traffic as well as support a portable rule set for mobile users. (6)

The Active Net Steward management console is comprised of two components. A service that runs under Windows NT and 2000 and the management console application that is used to review activity and administer site access rules and client security level profiles. Upon installation, ANS functions as a network monitor for the installed clients, recording details about all inbound and outbound traffic. This information can then be used to develop an educated security policy based on actual network use, unlike most other firewall products that block everything and require hours of configuration to return your network to a useable state. The management system listens for clients to register. During registration, security settings that have been defined for the physical device or for the user logged into the device are passed to the client. The clients and the ANS service communicate with a proprietary protocol, which is optimized to minimize its effect on the network. The client registers with the service the moment the network drivers are loaded during the OS boot. This gives unprecedented protection to machines that sit idle, i.e. public libraries or college computer labs. Computers can be placed into a stealth mode, pings are ignored and rejected traffic is dropped, would be attackers will not even know the machine is there. In the event that the client is unable to communicate with the management system, predefined policy can direct the client to register with a different system, failsafe and block all traffic or continue unprotected and attempt to establish communication with the management system.

Clients that are being policed have priority over those that are just being monitored. (7) In every way ANS is designed to have as little impact on overall network performance as possible. When a site access rule is created to block access to a specific application or URL, the first attempt to access that application or URL is

checked by requesting a ruling from the management system. Again optimization is the goal, so only the packet that starts a network connection is policed, not all of the packets that comprise a connection. ANS provides three policing levels, four if you include allow all. Level 1 – Casual control, at the same time the ruling is requested the connection is allowed to continue, if denied then the connection is dropped. This saves time for approved use, but does allow traffic that would have been blocked onto the network. Level 2 – Ridged control, the ruling request is made and if approved then the connection is allowed. The savings here is that no unauthorized traffic reaches the file. Finally, Level 3 – Secure control, same operation as Level 2, except that traffic is now encrypted. This will add considerable processing overhead, so make sure you have the CPU power before choosing this level. Once the ruling is passed back to the client it is stored in a local cache and the event is logged in the central database. The centralized database stores the connection information: time, user, destination, source, etc, that will be used for auditing and forensic purposes. The next time that same request is made; it is processed locally and the event is logged in the central database. By applying the rules locally from the cache, the traffic never gets on the wire, with several hundred users making several hundred requests the amount of saved bandwidth will add up quickly. Site access rules can be assigned to a specific MAC address, user and group of users or to all users. Rules can be created to block inbound or outbound traffic on specific ports, URL's, IP addresses or network ranges. ANS comes with a set of preconfigured rules that protect against Trojans, SpyWare and AddWare. This list is updated quarterly to keep up with current threats.

In addition to site access rules, ANS employs client service levels to control the behavior and reporting options of the clients. Wizards are provided to automate the service levels or to place a client into diagnostic mode for troubleshooting. Using the service control window the administrator has access to a list of currently connected clients and from that list view and change the security policy for any user or device. These changes can be applied immediately or scheduled for the next registration. Policy changes can be made for the user security policy, the device security policy or both.

A default service level is applied to all connections. It is in the service level configuration that the policing level is selected. There are also, fixed controls that determine if the client will respond to pings, accept unsolicited TCP packets, block all inward connections or stop the use of vulnerable applications such as TFTP or NFS. Separate reporting configurations for outbound and inbound traffic are also possible. ANS tracks both connection events such as repetitive connections, the duration of the connection and protocol events such as DNS requests, ARP requests and other broadcast events. Making it possible to quickly detect a workstation that has a DNS configuration error or in the case of a broadcast storm, the event can be detected and blocked. It is in the service level configuration that health checks options can be selected. By distributing the collection of health information across the clients in the network an extremely accurate picture of the overall health of the network is quickly created. Each client could be running with a different service level and site access rules. Changes to a client's service level can be pushed to the client immediately or

delivered the next time the client registers.

The management console provides the administrator with a view of network events and the tools to enable configuration changes and to initiate housekeeping functions.[7] It uses the familiar Windows Explorer tree view. Activity is listed by device and user, clicking on any user name gives a listing off all events that are related to that user, regardless of what machine they login to. Using the device drill down will show the events of a particular machine, this is helpful to audit if a user logins into a machine that they have no business using. Activity is also broken out by events and applications. These branches combine the data from all users and machines and sort the activity by type. A blocked event branch gives quick access to rule violations. Several canned graphs have been included such as Top Ten Users or Top Ten Applications. It is not possible to print these graphs, so they are of little use. The next release of ANS will be built on a SQL backend. This will allow for more robust reporting capabilities.[6] The database keeps a listing of all destination address and tools are provided to resolve those addresses and build a listing of destination domains. Users cannot hide their destination by using IP addresses instead of common names. Automated housekeeping can be used to keep that events database down to a manageable level. Different purge frequencies can be set for each type of data. Blocked traffic and health checks are of more importance than passed traffic, so this information can be retained longer. Alerts can also be set to notify the administrator if an unacceptable number of blocked connections occurs during a single session or excessive health check warnings. The product is designed to be a total intranet TCP/IP management center.

The Distributed Firewall market is still in its infancy. Several other products are currently on the market, but in one way or another they still allow the end user too much control. ANS is designed around the singular idea that end users are not to be trusted. It is sad that it we have reached this level of cynicism, but again, “A massive 75 per cent of all computer break-ins occurred internally.”[4] The FBI says cyber crime is the fastest growing areas.[5] It is amazing that will all the reports of hacking, identity theft, web defacement, embezzlement and the embarrassment that it causes, businesses still work under the “security through obscurity” mindset. New government regulations regarding financial data and health care information security is forcing many people to get their head out of the sand. Public access could potentially be every person with a computer connected to the Internet. It is no longer enough to put up firewalls and IDS devices; the real threat is the people who are already inside. Distributed firewalls represent the final stage in a complete security solution. Once every workstation, laptop, file server and host is protected it is very difficult to get much more granular. ANS represents the first in a new breed of product, the true Distributed Firewall.

Paper References:

(1) Steven M. Bellovin, "Distributed Firewalls", *;login:*, November 1999, pp. 39-47.
http://www.research.att.com/~smb/papers/distfw.html

(2) Ellen Messmer, “Second line of defense: Distributed firewalls”, *Network World,* June 6, 2000, http://www.nwfusion.com/news/2000/0605defense.html

(3) Security Designers, LTD. “Recognizing the Enemy Within.” URL:
http://www.securitydesigners.com/pdfs/TheEnemy.pdf

(4) Security Designers, LTD. “The Technology.” URL:
http://www.securitydesigners.com/pdfs/TheTechnology.pdf

(5) PowerLock Technical Training
John Earl, The PowerTech Group, Kent, WA, August 10, 2001
Handouts and Conversations

(6) Correspondence with David Lancaster, ANS Designer, Security Designers, LTD.

(7) Security Designers, LTD. Active Net Steward Sales Guide – V1.1 Security Designers, LTD. 2001

Questions:

Multiple Choices

1. Security is the responsibility of the _____
    a. Firewall
    b. Target
    c. User
    d. VPN Client

Answer – b – Every target should have some means by which to evaluate the connections coming into and going out of it. Relying on employees to behave themselves is no form of security.

2. This level of control uses encryption to secure communication between the client and the management console
    a. Level 1 - Casual
    b. Level 2 - Ridged
    c. Level 3 - Secure
    d. Allow All

Answer – c – Level 3 add encryption to the client/server communication

3. ANS security rules are stored
    a. Rules are note stored
    b. On each client
    c. On a remote SQL server
    d. In the central database

Answer –d – ANS security rules are stored in the central database

4. ANS clients are protected
    a. Upon OS boot
    b. Once a user logs into the network
    c. Only when browsing the Internet
    d. From e-mail born viruses

Answer – a – ANS clients register with the management console during OS boot

5. The ANS client run as a
    a. Protocol
    b. Application
    c. Virtual NDIS drivers
    d. Service

Answer – c – The ANS client runs as a Virtual NDIS driver in kernel mode.

True/False

1. Most computer break-ins are committed by unknown hackers.

Answer – False – studies show that employees commit the most offences

2. ANS will protect clients against known Trojans.

Answer – True – ANS comes with a preconfigured set to rules to protect against known

Trojans

3. ANS clients report all activity in real-time to the management console
Answer – True – ANS clients report in real-time.

4. Administrators must wait until the next time a client registers to apply service level changes
Answer – False – changes can be applied immediately

5. ANS checks the entire packet stream and then makes a ruling
Answer – False – only the initiating packet is evaluated