Global Information Assurance Certification Paper

Andrej Vnuk
SANS Security Essentials
GSEC Practical Assignment version 1.4b Option 2
July 3, 2004

How to replace a firewall in a live business environment with minimal impact for users?

A Case Study of changing perimeter defence from Altavista to Cyberguard Firewall

Abstract

Everyone agree that perimeter defense itself is not satisfactory. The concept that must deployed is Defence in-depth. On the other hand, the strength of a chain depends on its weakest link. Firewalls are the base building blocks for all network security strategies, therefore they must be reliable. If they begin to leak or become a choking point for Internet traffic, then a major change is needed.

My company had been having connectivity problems due to this reason for a longer period of time. During last year these problems began to affect bussiness operations, which was a clear sign that a new perimeter defense is needed. I was given a task to search for available products, evaluate them and present the most suitable one to the senior management. If the Board agree, I would have to deploy the new firewall as soon as possible with minimum impact on users.

This paper describes the actions taken to fulfill this task, step by step. I am reporting strictly on the firewall change, deliberately not mentioning our security policy or other elements of the defence in-depth, deployed in our company, because they are not essential in this case. Replacement was designed to improve perimeter security and gain new functionality without changing existing security strategy.

Before

My company's perimeter defense had been continous issue of discussions for many years. Being a system integrator and service level agreement maintainer has a collateral affect – our customers always have a priority, postponing our needs to a near, never defined future. We have been using Digital Altavista Firewall 98 on Digital Unix platform. Compaq (who inherited the whole DEC product line) sold the product to Axent Technologies in 1999, which were adopted by Symantec soon afterwards. The product itself was unsupported since 2000, as Axent allready had his own good firewall (Raptor).
The growing number of different problems called for a new firewall. The facts are that the current firewall was outdated, slow, unefficient and unsuitable for current corporate needs. These facts had been known for quite some time, but senior management refused to invest in new equipment, stating we could continue to live

with what we got. Serious problems arrised during last year that made them changed their minds.

Most critical ones were:

- Occasion DNS or mail relay attacks caused firewall to drop its connection to the Internet by shutting down its outside interface for that protocol. That usually meant loosing Internet connection for all users.
- Mail stopped to flow in either direction, sometimes several times a day. The reason was over helming of firewall's mail queues for embedded mail server, mainly because of large amount of unsolicited mail caused by virus outbreaks or Spam.
- Occasionally firewall shut down because of overflowing its log file when collecting events from the collateral activity of specific viruses.
- Lack of VPN capabilities turned out to be a serious drawback

All these lead to a decision to buy a new firewall. Being a security manager and technical support engineer in the company, I was assigned a task to search for appropriate firewalls available on the market, evaluate their specifications and choose the most suitable for us. After acquisition of the right one, I should make it operational as soon as possible. The means and methods on how to do it were up to me. There were only these guidelines to follow:

1. The new firewall has to be a superset of existing one, meaning it should support everything the old one does plus new features
2. It has to support VPN connections
3. It doesn't have to support High Availability (HA), as the old one will remain in place as a backup firewall
4. Users must not be affected in any way by this change – they shouldn't be even aware something changed

Research

Before starting to look for the right product I set up some additional criteria, which were approved by senior management. They were a combination of technical requests and some financial aspects, regarding scalability and future growth. So I went to a quest for a holly firewall with these in mind:

- The current firewall is a **proxy or application level gateway**, which turned out to be very good solution against many different application level (ISO OSI Layer 7) attacks, as well as all spoofing and relaying tries. More info on this subject can be found SANS literature [1].
- **Split DNS** on the firewall proved to be a great concept for easily setup of both, internal and external DNS. At the same time it was insensible for many DNS attacks like DNS poisoning or spoofing and illegal zone transfers.
- Good **logging capabilities** are essential for discovering the cause of potential problems.

On the other hand I knew which functionality of the Altavista Firewall I don't want to have:

- **mail server should not be embedded**. This feature caused us serial troubles whenever input or output queue got filled up for whatever reason, be it an internal mail server (Exchange) problem or temporary Internet shortout. Mail should only be forwarded or denied according to the rules set.
- **ethernet interfaces should be direction independent**. Limitation to only 3 interfaces (external, internal and DMZ) is ridicoulus for contemporary firewall. Especially if proxy level control can not be set over DMZ interface because it supports only packet filtering rules, like the Altavista firewall does.
- The new firewall should be an **appliance type**. Software firewall build on a non-dedicated hardware has several weaknesses, among which the most important ones are lack of purpously build elements (like ASICs for VPN encryption acceleration), common operating systems with all their vulnerabilities and especially troubleshooting in case of non-typical behaviour, when it is impossible to tell whether problem origins from hardware, operating system or firewall software.
- Classical architecture of a traditional Firewall consists of a computer (server), loaded with a common operating system like WindowsNT or some flavor of Unix, hardened to some point. Upon this comes a proprietary firewall application or in case of open systems (Linux) some freeware, like IpChains or IpTables, combined with Squid. There has been a lot of discussion about this topic in the Internet community, but I join to the group of experts who think that **only SecureOS concept is safe enough** to be resistant to everyday's discovering of OS vulnerabilities and their patching. It means that the underlying OS has been altered by vendor in such manner, that no changes to its structure are possible by anyone, making it resistant to common exploits. Another feature of SecureOS is dividing of operational levels or domains to a common level (lower) and more secure one (higher). Higher domain is allowed to read downwards but cannot write anything in the lower one. On the other hand lower level can not read anything from higher but should write upwards if so requested by more secure layer. This is similar to a concept of a military organization with secret and top secret services.
- **Certificated** product is definitely a good one. Independent certifications assures that a firewall passed very strict testing methods and indirectly make it better positioned on the market, which leads to assumption that this product will continue to be supported in the future. Therefore it should hold at least the following certifications: Common Criteria EAL4 or higher, ICSA and VPNC for VPN conformity and interoperability
- Firewall should support **remote administration through SSH**
- **VPN** feature should support **IPSec** for both, site to site and client (road warrior) to gateway connections.
- **Patching, upgrading** and eventual rebuilding of the firewall should be **fast** and **simple**.
- **Licensing** should not be limited to the number of supported users or nodes as well as for the number of coexisting tunnels.
- **Hybrid architecture** is preferable. Hybrid firewalls act upon all 7 OSI layers according to the needs. They are in position to work either as pure packet filtering devices or as full application proxy level firewalls, as well as stateful packet filtering and circuit gateways.

All these points made the decision process pretty straightforward, as there aren't many products on the market to satisfy all of them. I have read all possible test and opinions, finding most useful those at SC Magazine[2, 3]. Different forums and firewall mailing lists had helped me narrowing my choices. However there is a firewall that seems to be an incarnation of all my demands. Cyberguard Premium Line Firewall/VPN Appliances are exactly what I was looking for. Considering all the expenses I went for the entry model, Cyberguard FS250 and compared it to all of its competitors step by step. For a reasonable price it offered hybrid proxy appliance firewall, built on proprietary SecureOS, with IPSec VPN support for unlimited users and number of tunnels, superior performance and scalability. All these come along with highest possible certifications on the market, like CC EAL4+ with ongoing assurance maintenance program, ICSA, ITSEC E3, CheckMark and VPNC. [4, 5, 6]

During

It wasn't hard to convince senior management for buying this firewall, much tougher though was to delay the final installation, as they wanted to have it working at once. All of a sudden there was an urging need for VPN connections with our partners, as well as high demands for home users access to company and higher overall speed for LAN users connecting the Internet.

I had to stop this rush generated by big expectations. Everyone believed our entire Internet oriented troubles and many others would be solved over night. Anyway, I succeed to convince them I am not a magician with some kind of a magic stick, so I need some time to study the new firewall and try it in our test environment for some days, before final act. Furthermore I needed to test all new features and configure the firewall in the same manner the existing AltaVista was. This was crucial, as swapping of the firewalls should happen at the same time. I planned also a spare test time for rollback if unexpected behavior should appear. Beside this the implementation of the new firewall wasn't dependable on myself only, the cooperation of my colleagues, who are in charge of the network administration and mailing system, was inevitable. Because all of us were very busy with other more urgent tasks, I asked senior management for a one month implementation time frame, which was granted after all.

Testing

Testing period was crucial for successful implementation according to my planes. There are two ways of approaching – testing on a completely isolated network environment or testing in the actual, live environment, but in a role of a spare unit without any responsibilities assigned. Both methods have pros and cons, but testing in live network simplifies the process of creating necessary mail and DNS servers at both sides, external and internal. Use of actual running services can also prove correct behavior of the firewall, but special care should be taken not to corrupt existing records and entries, within the company and DNS worldwide.

Configuration

I configured basic settings first, like outside and inside interfaces, domain and host names, administrator's passwords and enabled SSH for remote console administering. This feature is a great relief, because AltaVista Firewall didn't have it, unless using a VPN tunnel. I was really fed up with all that fuzz and noise in the computer room, therefore working in the peaceful environment of my office and my desktop only was more than precious.

Next thing to do was setting up split DNS in role of a secondary name server for either side. Cyberguard firewalls have a nice feature of automatically adding appropriate rules to the rule set, so all necessary rules had been put in place by itself. Of course, the easiest way to test an Internet gateway is to enable web surfing. Rich set of embedded proxies calls for using them, so I set up a proxy to let all internal http and https traffic through (outbound through firewall) and none external into (no inbound through firewall). Because most of the clients are set to use a proxy for port 8080, I added this port to the set of listening ports for http proxy.

HTTP

I didn't bother with setting mail and other features at this time. All procedures described, from unpacking the firewall, connecting it to make it work, took me half an hour, just like Cyberguard says it should. Everything went pretty straightforward and smooth and my desktop was able to surf the web as soon as I changed default gateway setting from the old firewall address to the new one. I did have a problem though – Cyberguard refused me, if I tried to set my proxy to be the new firewall, either for standard port 80 or added 8080. The log file showed the same thing – web traffic from my PC is denied by firewall.
The reason was discovered soon after. Firewall works normally in a transparent proxy mode, which means that users are not even aware of something catching their packets, inspecting them and retransferring them (if they are allowed by rules, of course). Anyway, if a user sets the firewall to be his proxy, then this is non-transparent mode from the firewall point of view. Cyberguard calls this mode of traffic flow, regarding also a direction, outbound to firewall. There is a checkbox on the proxy setting for this mode, and a single click solved the problem. The result was automatic adding of two more rules to the firewall rule set, stating that there is a proxy web traffic allowed from internal to the firewall for port 80 and 8080 and a packet filtering rule which allows connections from firewall to external interface.

These basic steps enabled web access for my desktop and two more test workstations, for http and https traffic (ports 80, 8080, 443) as well as embedded ftp. Intensive testing, firewall logs and packet analyzer proved everything is working properly. It was time to move on.

SMTP

One of the most vital functionality for our company is email. Lots of unsolicited mail messages (SPAM) forced us to set up a means to control the smtp inflow. There was an Alpha based OpenVMS system available, so an AntiSpam mail server was set up with some help from public available SPAM lists (RBL - Realtime Blackhole Listing)[7]. This system sits in front of our firewall and is known worldwide as the final mail relay for my company's domain. All I had to do was changing the IP address of the

gateway to forward mail and restart smtp process. This operation didn't require any others assistance and wouldn't affect users, taking only a minute to complete. Of course, the smtp proxy for inbound through was set up first on Cyberguard. Unfortunately Mr. Murphy never sleeps – it took me whole evening to find out the reason for not working. OpenVMS mail relay couldn't send anything, his log starting to fill up. Cyberguard showed smtp connection from this system is up and working, but no packets had been received. Use of a packet analyzer revealed the cause. First thing was cache on the OpenVMS system – it had to be flushed by restarting the machine. The second thing was trivial. The firewall is next hop (forwarder) for the mail relay, therefore an inbound to option must be checked at proxy settings. Having it done so, everything returned to where it belongs.

Different tunable parameters can be controlled through a proxy, so I checked some of them. I tested blocking of mails with specific types of attachments, as well as having specific words in the subject, sender or recipient line. Everything works fine, so I am considering having another means for fighting Spam running at the smtp proxy level.

This test turned out to be successful enough to leave all incoming mail traffic this way. Outgoing mail couldn't be tested at that time due to absence of mail server administrator, who is the only one authorized for touching the Exchange server. Anyway, no troubles were expected, as only one setting needed to be changed. The actual setting on the internal mail server used routing mode to point to the next mail relay (AltaVista Firewall) for outgoing mail, meanwhile Cyberguard requests this mode to be DNS, as it includes no mail server to accept or handle mails. This was left over to try on Day D.

VPN

This is the most wanted feature for our expanding business needs. Cyberguard supports only IPSEC tunnels, claiming other types of VPN don't meet security expectations (PPTP for example) or aren't compliant to standards (proprietary protocols), which are the only factor assuring connect ability to other vendor's products.
We need both types of VPN connection. Site to site tunnels are needed to connect our remote offices, client to gateway (road warriors) tunnels were to make our subcontractor services available, as well as letting senior management access company resources while being out of the office.

I must confess that I was most afraid of this part of testing. Surprisingly everything went smooth. The firewall was configured by the book, using plain password authentication and default parameters for encryption, like main mode key exchange with no PFS, 3DES and SHA1. Two types of VPNs were created. One was a site-to-site connection to remote office, better said to dislocated network. The other was endpoint for road warriors, sharing a pool of 100 virtual IP addresses for as many clients VPNs.
Having done that, I loaded Cyberguard VPN Client software on my laptop, setup up appropriate parameters and dialed to my private Internet account. A tunnel was up and running by a single click.

When I got to remote location I had some troubles with their ADSL router with embedded firewall and VPN server. Phase 1 of ISAKMP had been negotiated, while phase 2 failed. The problem was in different approach to setting IPSec parameters on the firewalls. As soon as I worked them out, the tunnel was up and both sites were connected. Because all protocols were intentionally allowed through, all connectivity tests succeeded.
Simple passwords are not considered to be very secure, so I tried out a two-factor authentication. Cyberguard firewall can be used to generate digital certificates, which can be used as additional means to secure both negotiating parties. Configuration was done following instructions in the Administrator's Manual and again no troubles appeared.

HTTPS

Several years ago appeared demands from senior management to deploy a cheap and quick solution for remote access to their email. Because Internet access was usually always available in such situations, solution had to support connectivity through web access. Since our email server is Microsoft Exchange, the most natural solution was to enable OWA (Outlook Web Access) and to redirect any https (port 443) traffic at the firewall. Being an elegant and useful solution, everyone uses it now. Last year a new web service began to run on one of the internal servers, using the very same protocol. Remote access to longer emails took more processing power from our outdated AltaVista Firewall, affecting these services. Therefore a decision was made to move OWA from port 443 to port 54321. A new proxy was created on the firewall to handle this port and all troubles had gone away when the clients changed their settings for remote access to the following syntax:
https://FW_IP_ADDRESS:54321/exchange

Unfortunately Cyberguard doesn't support port numbers above 32768; therefore port redirection couldn't be used when implementing the new firewall. The best solution I could think of was to define a new host name for OWA public access, publish it at our ISP's DNS and create Static NAT redirection to our mail server. After a couple of days, when name servers worldwide were updated, a test was run from various points in the world (thanks again, my friends everywhere!), just to show the decision was right. All other https traffic was not affected, as it ran over standard port 443, which was allowed through.

WWW

Finally, our web server had to be redefined. Up till now it ran on the external side with public address, therefore having no firewall protection at all. This was because nobody ever bothered to build in another network card and set up DMZ on the AltaVista in the first place. But when you start changing things, it is time to do them right!
Another interface was dedicated to be DMZ, all necessary DNS records were added to external part of the Split DNS on the firewall and web server was reconfigured and reconnected according to the new settings. Using the previous experience of redirecting IP address for OWA, a static NAT was set to show publicly known IP address to a newly defined DMZ. According rules were built using proxy, allowing

everyone web access only (only PUT command allowed within http stream), but to selected internal users all protocols were allowed for administration.
Testing again, checking logs, trying to penetrate to web server – 100% success!

As these tests were done on the firewall being connected to the network in parallel (coexistence) with the old one, the very last test would be swapping both devices. The time had come!

Implementation

Before the final step there was still some work to be done. All rules had to be generated to be a functional twin of the AltaVista firewall. All DNS records, proxy and packet filtering settings had to be checked. This was done in an hour. There was also one thing to change in our network environment.
Most of the servers are using NTP for setting up time. It makes sense to proclaim one internal machine to be NTP server for LAN, thus reducing the number of simultaneous connections to only one. This used to be one of the AltaVista firewall functions. However Cyberguard firewalls do not support this feature as it could lead to a possible exploit. Non-firewall server features are not allowed on the device, although it uses NTP for self-updating.
This meant we had some work to do before going live. One of our internal Unix machines was dedicated to be a new NTP server; all other servers using this feature were pointed to this computer. A rule was set up to allow NTP traffic updates through the firewall.
Finally IP addresses of external and internal interface were changed to the actual ones, as well as firewall host names and the device was restarted. At the same time AltaVista was shut down and Exchange mail server had its outbound way defined to use DNS.

We were very tense next few minutes after reboot. All kinds of tests were being run, just to show everything is normal. The whole crew spent next few hours thoroughly testing every possible feature. There wasn't a single mistake anywhere!

After

Next day users were informed that a new firewall is working and any problem should be immediately reported.
VPN connections to remote sites were configured and brought up, as well as a client connection to our partner, a contracted service provider, who was assigned a special tunnel with access rights only to a specific internal server. This connection turned out to be the only problem later on.
Partner uses an application to control our internal VOIP server. This application sends data to an external server, using FTP. The new firewall was blocking access to this server by some reason. Packet analyzer showed that specific commands within FTP stream triggered alarm at FTP proxy. Manual connection to the same server worked, as it should. Enabling keyword "RETR" in the list of allowed commands at the proxy resolved the problem.
After several days similar problem occurred when accessing identical FTP server on another location. This time we couldn't find the reason, however we found the

solution. Changing a rule from a proxy type to a stateful packet filter made all the problems disappear.

Conclusion

Once again a thoughtful plan and well tested network configuration proved to be irreplaceable. Only because all features of a new device were tested in advance could a swap over take place without users even noticing it. This approach had several benefits, as we all learned quite a lot during this process and the old firewall remained in place to be used as a cold standby. Should anything go wrong with the new appliance, there is always the old one there, waiting only to be switched on. Of course this is just a temporary backup solution, because it doesn't support VPN, but it's better than nothing and it costs nothing.

There is still a lot of work to be done regarding fine-tuning and regular administration of the firewall (tracing log files, backing them up, testing other features, improving responsiveness, etc.), but the basic task, I had been assigned, has been completed.

The most important thing to keep in mind is that the work is never done in security business. My advice to everyone: Keep up to date with all activity in this area and read as much as you can! [8,9]

References

[1] Cole, Eric., Fossen, Jason., Northcutt, Stephen & Pomeranz, Hal. SANS Security Essentials with CISSP CBK Version 2.1 Volume One. SANS PRESS. Chapter 14, page 668.

[2] The SC Magazine homepage:
http://www.scmagazine.com/home/index.cfm (July 3, 2004)

[3] Ido Dubrawsky. Firewall Evolution - Deep Packet Inspection.
July 29, 2003
http://www.securityfocus.com/infocus/1716

[4] A CyberGuard Corporation White Paper, September 2002.
CyberGuard Application Gateway Proxies: Beyond Packet Filtering
http://www.cyberguard.com/resource_center/WhitePapers/Application%20Proxies%20-%20Beyond%20Packet%20Filtering.pdf (July 3, 2004)

[5] Cyberguard Products Overview
http://www.cyberguard.com/solutions/product_fs.cfm (July 3, 2004)

[6] Paul Henry. Why CyberGuard? 26 reasons to choose CyberGuard Firewalls.
http://www.bluesky.com.au/Products/CyberGuard/Comparisons/CyberGuard26reasons.pdf (July 3, 2004)

[7] The MAPS RBL (Realtime Blackhole List) homepage:
http://www.mail-abuse.com/services/mds_rbl.html (July 3, 2004)

[8] R. W. Binnion. Network and Internet Security Issues and Solutions - an introduction.
http://www.itsecurity.com/papers/p34.htm (July 3, 2004)

[9] Symantec Internet Security Threat Report
Volume V
March 2004
http://enterprisesecurity.symantec.com/content.cfm?articleid=1539&PID=20336947&EID=504 (July 3, 2004)