Global Information Assurance Certification Paper

INSTALLING, CONFIGURING AND ADMINISTRATING THE BORDERWARE FIREWALL SERVER AND ITS ROLE IN DEFENSE IN DEPTH

By

Graydon S. McKee IV

A practical paper submitted in partial fulfillment
of the requirements for the obtainment of the

GIAC Security Essentials Certification

Global Information Assurance Certification in
connection with the SANS Institute

Practical Assignment 1.4b – Option One

2003

INSTALING, CONFIGURING AND ADMINISTRATING THE BORDERWARE FIREWALL SERVER AND ITS ROLE IN DEFENSE IN DEPTH

By

Graydon S. McKee IV

ABSTRACT:

Network Security and the means to defend corporate systems have become increasingly hotter topics in today's corporate environment. Gone are the days of a single solution to protecting the corporate network. What has arisen from the experience of many is a multi layered approach to the concept of network security. This approach has been identified as Defense in Depth. A key role or layer in the Defense in Depth concept is the role of the corporate firewall. One such firewall is the Borderware Firewall Server from Borderware Technologies. The Firewall Server 6.5 is a Common Criteria Certified EAL4+ with EAL5 Vulnerability Analysis Application Proxy Firewall that is fully configurable. Proper configuration results in a strong framework from which the Defense in Depth Concept of Network Security can be applied.

This practical will demonstrate how the Firewall Server 6.5 can be configured, administered and tested to provide the framework for developing a secure network security strategy. Configuration of the Domain Name server, Mail Server, Proxies, the Proxy Server and Logging will be addressed as well as means to test the configuration to ensure the firewall acts as intended in a production environment.

Chapter 1

Types of Firewalls

Simply speaking there are three types of firewalls: packet-filtering firewalls, Stateful Firewalls and Proxy Firewalls. While these are the main types of firewalls, there are products on the market that are a mixture of some of these technologies. It is important to understand how each of these firewalls operates when designing your network security plan. Let us start with the simplest and fastest type of firewall, the packet filtering firewall.

Packet filtering firewalls are the most fundamental of firewalls and work primarily at layer 3 (Network) of the OSI Model. They are usually, but not exclusively routers, which examine each packet as it flows thru the device and filters them based upon a certain predefined rule set. These devices examine each packet header for the required information (i.e. Source or destination address and port requested) and drop all packets not meeting the requirements set forth in the devices rule set. Packet filtering firewalls are often very fast as they only examine the IP header before deciding to pass the packet along. Care must be taken when designing packet filtering firewall rule sets. A misconfiguration or forgotten rule may very well negate the entire rule set and violate the sites security policy. While packet-filtering firewalls offer some benefit when it comes to security, there are other technologies that address the issue of security in a more comprehensive manner.

Stateful firewalls are packet-filtering firewalls with OSI layer 4 (transport) awareness. It is for this reason they are sometimes referred to as dynamic packet filtering firewalls or Stateful Inspection Firewalls. They examine not only the origin and destination IP address but also keep track of open sessions or states. If a transmission comes to the firewall and this transmission appears to be a legitimate reply then the transmission is permitted to pass. One of the benefits of a Stateful Firewall is the low processor overhead. This allows more concurrent connections before degradation in performance is noticed. Stateful firewalls bridge the gap between speed and performance between packet filtering firewalls and Proxy firewalls.

The technology used in Proxy firewalls is different from that used in Packet Filtering and Stateful Firewalls. In each of the two previous examples, the device in question would inspect the packets flowing thru them and drop those who do not meet the criteria set forth in the rule set. What this means is that the client computer and the server computer have established a connection and are passing information back and forth between them. The packet filtering and Stateful firewalls simply examine that traffic in doing their jobs.

The Proxy firewall operates on a different premise. The client computer who desires to interact with the server computer does so via the proxy firewall. The client computer establishes a connection with the Proxy firewall, which acknowledges on behalf of the server computer. Once a connection has been established, the proxy firewall then establishes its own connection with the server computer. In essence, there are two separate connections going on here. Once this has taken place the proxy firewall then examines the packets it has received from one computer and applies its rule set to the packets before handing them off to the other computer.

Understanding these technologies is important in setting up your architecture in a manner that lends itself to security rather than works against it. In examining the principle of Defense in Depth, we can see that these technologies are not always exclusive of each other. They can be deployed or implemented on existing devices to complement each other. Envision a situation where a business is connected to the Internet thru its Internet Service Provider. The ISP typically terminates its connection at a router at both the client facility and internal to their own systems. This router can be set up to filter the packets that flow across it from the ISP into the business. Since this is the fastest means of filtering connections it is done with relatively little overhead and little degradation in performance. The connection is then routed directly to a proxy firewall that controls all connections between the trusted network inside the business and the untrusted world of the Internet. We now have two layers of security adding to the principle of defense in depth and contributing to the security of the network.

Chapter 2
Defense in Depth and the Borderware Firewall Server

Defense in Depth is a concept by which security is applied to a network in a layered fashion to create multiple layers of defense. This concept has replaced the old concept that dictated that network security be implemented in a simpler approach by securing the perimeter of the network while ignoring the center itself. The most often used analogy was a piece of candy which was soft and gooey on the inside but was surrounded by a hard crunchy coating. The problem with this approach is that once the hard crunchy coating was penetrated, nothing was left to protect the network.

The Defense in Depth concept developed to meet the deficiencies in the single barrier concept. Creating a layered approach to network security, an attacker must defeat multiple layers before he or she reaches the resources they wish to access (the gooey center of our candy.) While it is understood that no network security scheme is full proof, the concept revolves around the adage that the more difficult it is for an attacker to penetrate a networks security measures, the more likely the attacker will move on to find an easier target.

It has also become increasingly more apparent that the threat to network systems does not always come from outside the network. Probes and incursions from inside the network make the principle of Defense in Depth more relevant in today's networks. Joel Snyder in his June 2003 cover story for Information Security Magazine discusses the need to think inside out as well as outside in while designing a network security structure. While his article discusses the means by which multiple vendors designed a network security structure to focus on access control and authentication, it is useful to note that firewalls and firewall technology played a prominent role in most of the designs. They appeared not only on the perimeter but throughout the network in segmenting roles.

Mr. Snyder also identifies six roadblocks to implementation of Defense in Depth; some of which can be met by the deployment of a Borderware Firewall Server. These roadblocks are:

- Cost
- Performance
- Management
- Policy
- Authentication
- Binding

Several of these points can be overcome by use of Borderware's Firewall Server. Firstly, the cost of the Borderware's Firewall Server is substantial below some of the other products out there and is in numerous ways a superior product. The performance issue is minimized due to the Firewall Servers ability to achieve

throughput of near line speeds. Management is simple and straightforward with the Firewall Server. Its ability to handle up to six interfaces allows for a great degree of separation between network segments and a central point of administration. Authentication of users may still be an issue; however the Firewall Server allows authentication in the nature of Internet access thru an available Proxy Server running in non-transparent mode and authentication of Mail and FTP access. With the centralized administration the application of security policies is made simpler thus minimizing the impact of the Binding issue.

Today's Firewall systems must be configured and managed to handle security threats from all sides of the network. Segmentation of network resources is key to the Security Administrators ability to protect resources. The ability of the Borderware Firewall Server to accommodate up to six network cards provides Security Administrators with the ability and tools to segment their networks to their hearts content.

Firewall Systems provide the framework upon which the network should both be designed and defended. McClure, Scambray, and Kurtz, in their book: Hacking Exposed: Network Security Secrets and Solutions indicate that a strong firewall is key and that "most attackers make every effort to work around a strong firewall."

Chapter 3

Installation of the Borderware Firewall Server 6.5

Installation of the Borderware Firewall Server is a straightforward process that is easily accomplished. It can be purchased in two forms: an appliance device and as separate software and installed on an Intel based platform. The instructions contained here are aimed at the full installation of the software on your own hardware. For the most part these steps will parallel those used in the setup and configuration of the appliance device.

You will receive your hardware either on CD-Rom or in the form of an .ISO file that you will need to burn to a CD-Rom. You should also receive your license key at the same time. The Borderware Firewall Server will function for 30 days in evaluation mode without this license key however; it is recommended that it be obtained as soon as possible.

The Borderware Firewall server includes a fully functional secure operating system based on the FreeBSD flavor of UNIX. It is the only software that will be running on the hardware and it will overwrite any pre-existing data that exists on your server. While FreeBSD UNIX provides the framework, from which the operating system runs

The Borderware Firewall Server 6.5 is certified to EAL4+ with EAL5 Vulnerability Analysis. Borderware has extensively modified the operating system so that it is fundamental and structure is unique. Borderware has taken great pains to ensure that their S-Core Technology, as they refer to their operating system, is secure against attack. They have removed all features and processes that have been deemed insecure even to the point that no user accounts are maintained. Access to the system is only thru the management interfaces. This means that as the administrator you will be unable to access the command line or interact with the operating system directly. Access can only occur via the console or remotely thru Borderware's secure management interface for Windows: BWClient. Access can also be obtained thru Borderware Technical Support. Borderware has a means to access the firewall thru a function called Support Access. Borderware Technical Support makes a secure encrypted connection with the external interface of the firewall in order to get a clearer understanding of what is happening during a support call. This type of access is only open to Borderware Technical Support and is not available to the Administrator.

As the consumer has the option of just buying, the software and installing it on their own platforms there are some minimal requirements that the hardware must meet.

300 MHz Pentium, PII, PIII or Celeron Processor
64 Mbytes Memory

IDE Disk Controller
4 Gigabyte Disk Drive
PCI Network Cards
CD-Rom Drive

These are of course the minimum. You should make evaluate your network to make final decisions on the type of hardware you will deploy. You can choose from two to six interface cards, depending upon your needs. One hint here, Write down the MAC addresses of each of the cards you install in the device. You will need that information later when you configure each of these interfaces. Drives can be set up in a RAID configuration and there is some support for SCSI drives.

Once you have made your decisions about hardware, installation of the Firewall Server can begin. Make sure your BIOS settings allow booting from CD-Rom. Place the Disc in the CD-Rom Drive and boot the machine.

When the machine boots up, you will be prompted to make some choices. The first question you are asked is about the style of keyboard you are using. For most installations you should chose the US type of keyboard. You then have the opportunity to select between an automatic install and a custom install. Borderware suggests selecting an auto install and if you experience problems to go back and try a custom install. Experience has shown that the automatic install works well and a custom install should be tried only when special circumstances dictate the need. With either choice, you will be reminded that all data that currently exists on these drives will be erased. Described below is the installation procedure used if you were to choose the custom install option. The automatic installation will make several of these choices for you and if you do an automatic installation, you will not be asked to complete all of these steps.

Once you chose to do a custom install, you will be asked to confirm that you know that the information currently on your hard drives will be erased. Next, the installation program will detect your hard drive size. If you have more than one hard drive installed then you will be asked which hard drive you wish to install as the boot device. You are then asked if you wish to make a new software installation. Click ok and proceed to the next section where you are asked to confirm the loading of the software. Click Ok to proceed.

You will need to confirm the Disk Drive Geometry in the next section. This software is pretty accurate in its determination of the drive geometry however double check the figures for the number of cylinders, heads and sectors just to be sure. You can use the Tab key to move thru these screens. Next, you are asked once again to confirm that the software is being written to the boot disk.

Next, you have the option of which location you would like to run the install. Your choices are from the CD Rom or from the network. This document will detail the

CD-Rom install. A confirmation screen will open next asking you to verify the type of software you are installing. This is because Borderware occasionally ships several of their products on one CD-Rom. Chose the Firewall Server Option (It may be the only option showing.) Confirm the amount of system memory next. This occurs because of some problems Borderware found with a few Compaq servers. Just confirm the amount and select ok. The system will now reboot.

Upon reboot, you will be asked to partition your drives. You have the choice of either accepting the defaults or customizing them on your own. The screen should look something like this:

Disk	Partition	Size (Mbytes)	Usage
Ad0	Root	256	Operating System Area
Ad0	Swap	1024	Swap Area
Ad0	ftp+www	3562	ftp/www files
Ad0	Logs	3562	Log files
Ad0	Mail	3562	Mail Server Files
Ad0	Config	3562	Configuration Data Area
Ad0	Squid	3562	Squid Cache Area
Ad0	Free	2	Unallocated

Borderware recommends twice the amount of ram be allocated to the swap area. Experience has shown that the swap space should be two and a half times the amount of RAM. While this may be overkill for some network environments, you will avoid some problems as your network grows by setting the swap space this high. The above example is based on one 20 GB hard drive and the allotments for size are what the program defaults to for that size drive. You can also move partitions across different disks should your design require this. The screen used to reconfigure the sizes allows changes in 1, 10, and 100 Mb increments. This can be done by selecting the Edit button at the bottom of the initial screen. Once your reconfiguration is complete Click Ok to exit the reconfiguration screen; confirm the new settings and then click OK to move on to the next section. You are asked again to confirm that you want the disks configured as you specified on the previous screens and that all the data currently existing on these disks will be lost. Click yes to proceed. The installation now begins to set up the file system and a reboot occurs.

You are now asked for some system information such as region of the world, country, and date and time. Once providing these, you are allowed to log into the console. Note the default password is “Firewall” and as this system is based upon UNIX technology, you need to be aware of case as you type it. It is HIGHLY recommended that you change this password at your earliest opportunity. You can change the password by choosing that option from the Admin menu. Prudence suggests using a combination of letters, numbers, and special characters. One of the places that Borderware is lacking is that they limit you to eight characters. Once you have reset your password you must configure your interfaces. This must be done prior to licensing your firewall.

From the Admin menu chose the Configure Interfaces option. You will be given a total of three screens. Moving thru these screens is simple. Pressing enter progresses you thru the screens and if you make a mistake simply press ESC and you will leave this menu without saving any changes. You will have to start all over again.

As you move thru the screens, you will be asked to identify the cards. This identification is done by a combination of the interface number and the MAC address. For each interface, you will be asked for the IP address and the subnet of the network it will be connected. You also must specify the speed of the line they will be connected to. You can leave the selection to auto select or you can specify a specific speed. This is used to calculate the load when you are displaying the system activity and is helpful to have when you need to determine if you need more bandwidth. It is worth noting here that the Borderware Firewall Server is capable of processing 90 MBps so the speed of connection they will have to the outside and inside world may cause a bottleneck. An upcoming feature of the Firewall server will allow you to perform load balancing with multiple firewalls on a single connection. Depending upon your Network Interface Card, you may wish to allow the software to auto select the speed. If a problem occurs, you can always go back and reconfigure your interface cards at any time to specify a speed. You are also asked the speed of the slowest part of the link between the Firewall and the external network at this time. Again, this is for the calculation of bandwidth being used. You are then asked for the domain name of your organization as well as the host name you wish to have for this firewall. Once you have changed your default password and configured your interfaces you must reboot the firewall.

Now you license your firewall. If you are currently evaluating the product, you can skip this step. You need two pieces of information here: the license key provided in the license pack envelope and the System ID. The System ID is tied to the external Network Interface Card. If you need to rebuild the firewall sometime in the future you will need to use the same Network Interface Card that is being used as the external interface or you will need to call Borderware and ask them to reset the System ID that is with your license key. This System ID is shown upon boot up on the initial Borderware Splash Screen.

From the Admin Menu select the Update Software License option. What you need here is an activation key. Take your license key and the system id and go to http://www.borderware.com/activate.html to obtain the activation key. Once this key is input into the console you will need to reboot the firewall. You have just completed the install and need to begin to configure your firewall. You do this with the secure BWClient provided for Windows Platforms.

Installation of the BWClient is very easy and straightforward. Simply download the client from Borderware's Web site and run the installation program. The install program will drop an icon on your windows desktop and a simple double click will activate the program. You will notice that nothing is filled in on the BWClient. You need to specify the internal interface of your firewall to connect and begin to administer to your firewall.

Chapter 4
Configuration of the Borderware Firewall Server 6.5

Administering to the Borderware Firewall Server via the BWClient interface is straightforward. Borderware provides an extensive help feature that assists the administrator in understanding the various interfaces. Several of the important features will be discussed and the security implications of each will be addressed.

In order to begin administration via the BWClient you must first set the firewall up to accept the connection. This is done from the console in two places. For the first option, select the Services option on the console and then select the Configure Servers entry on the menu. There are two choices available on this menu: Internal Servers and External Servers. Select the Internal Servers option from these choices. There are several internal servers listed here. The two of interest for this process are the GUI Config and the Secure GUI Config servers. A word of note here, if you elect to enable GUI Config you will access the firewall over an unencrypted channel connecting to the firewall via Port 441. If Secure GUI Config is enabled you will establish a secure encrypted SSL connection via Port 442. This port is for a UNIX daemon called cvc_hostd. When you enable the service, you have the option of logging the connections. The entries to this log are kept in the security log and can be viewed. What is listed is the internal IP of the machine that connected to the Firewall, not the user name used to connect. The second option is found under the Admin Menu in the Configure Remote Administration Option. Here you simply chose to enable servers and chose which interfaces you want to connect from for remote administration. A short note here about configuring remote administration from the external interface: If you enable remote administration from the external interface, you must use the secure SSL connection via the BWClient and your user must be configured to access the remote configuration via a Crypto-Card Authentication process.

A username also needs to be set up via the console prior to logging into the firewall for the first time via the BWClient. This can be accomplished by choosing the Admin Menu and selecting the Configure Remote Administration option. When this opens, you wish to add an administrative user. When adding a user you are prompted for their username and the type of authentication allowed. Two types of authentication are supported in the 6.5 version: Crypto-Card Authentication and password (or none) authentication. If you chose no authentication, you are warned by the software that this leaves only the secure nature of a complex password in keeping unauthorized users from accessing the firewall. Remember that Borderware only allows passwords of eight characters and no more. While this password can be complex or simple it is still just a password and thus easy to crack. It should also be noted that no matter what type of authentication you elect to use, all users will use the same password to access the firewall. Multi-tiered Administration is not yet available so care should

be taken with the number of accounts you install on the firewall and who has access to them.

If you chose that, you want the Administrative user to use Crypto-Card Authentication you will be prompted with screens showing you how to configure the Crypto-Card to authenticate to the Borderware Firewall server. Once the user has been added you will be able to utilize the BWClient to access the Firewall remotely for configuration and administrative purposes.

Figure 1: BWClient Initial Login Screen

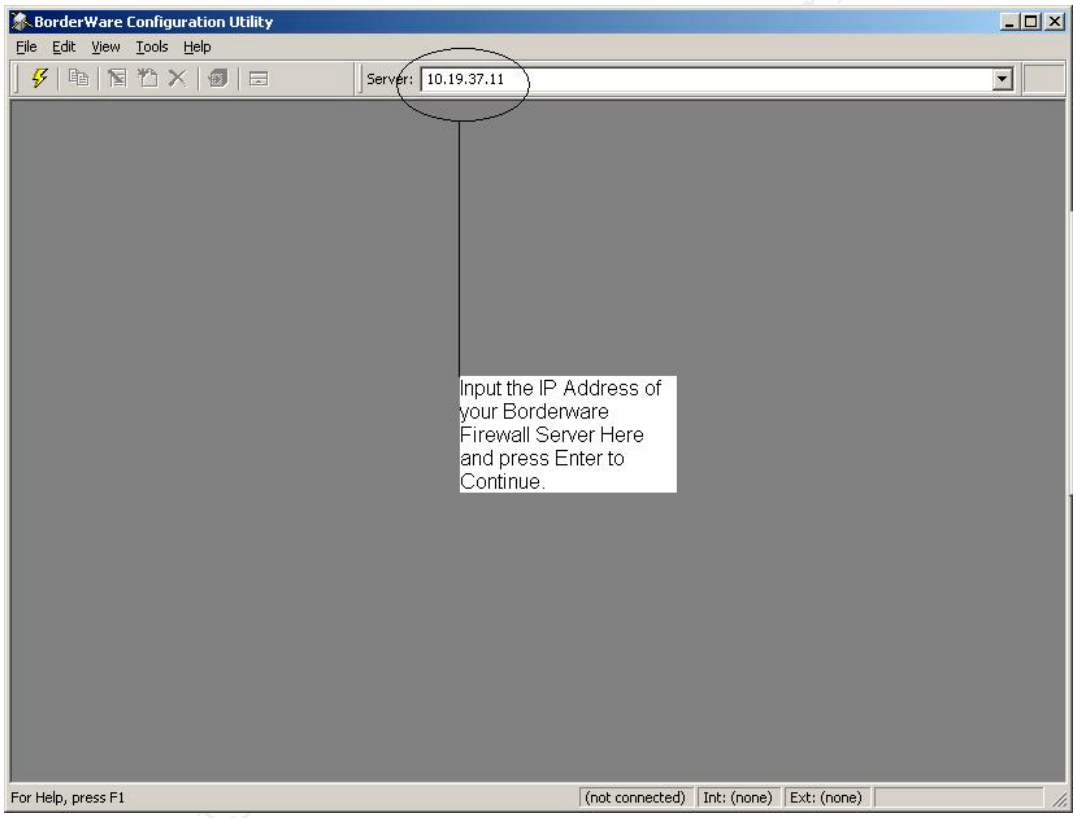

When presented with the initial screen of the BWClient you will need to specify the Internal IP address of your Firewall in the server field at the top of the screen. Please refer to Figure 1 for an illustration of this. (Note: All illustrations are taken from a firewall server built for this practical on an isolated network and thus all IP addresses shown have no relation to any production environment.) Once you press enter, you will be prompted for your username and password. If you have selected to have the administrative login authenticated via a Crypto-Card then you will see the screen shown in Figure 2.

Figure 2: Secure Login via the BWClient

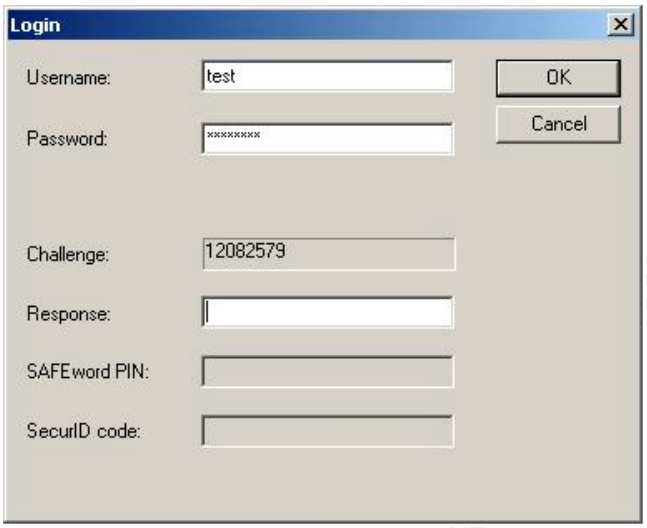

Once logged into the Borderware Firewall Server you are presented with many options. As this paper is not intended to be a feature-by-feature description of the Borderware Firewall Server, some areas will be glossed over rather quickly. The focus will be on those items needed to get the firewall up and running with minimal trouble. Each area has a help function that goes with it and very complete descriptions are given in the help files. These can be reached by clicking Help off the toolbar at the top of the screen or by clicking on the appropriate help button on each screen. The Administrative choices (Shown in Figure 3) have several options from which to choose. Some of the options will be dealt with in later sections of this paper while others are important to cover immediately.

The Administrative Folder

The screens available thru the Administrative Folder are System Settings, Password Settings, Multi-address Translation, Alarms, Secure Logins, Static Routes, Support Access, System Activity, Backup/Restore, Software Updates, Download Patch and Shutdown. While a basic description of each will be given, only Secure Logins, Static Routes, Support Access, Backup/Restore and Download Patch will be dealt with in any detail. Multi-address translation and Alarms will be dealt with in later sections.

The Systems Settings Menu allows you to Set the Remote Management type: Internal, Internal Secure and external. You can specify a SYSLOG server to send log messages to and you can change your time settings here.

The Password Settings Menu allows you to change the password for the configuration area and the FTP area. A separate password for each can be set so you can have some control over the two ways to access these areas.

The System Activity Menu allows you to view either a summary or a detailed view of the connections being permitted thru the firewall itself.

The Software Updates Menu shows you which patches have been installed on your firewall.

Figure 3: Administrative Choices

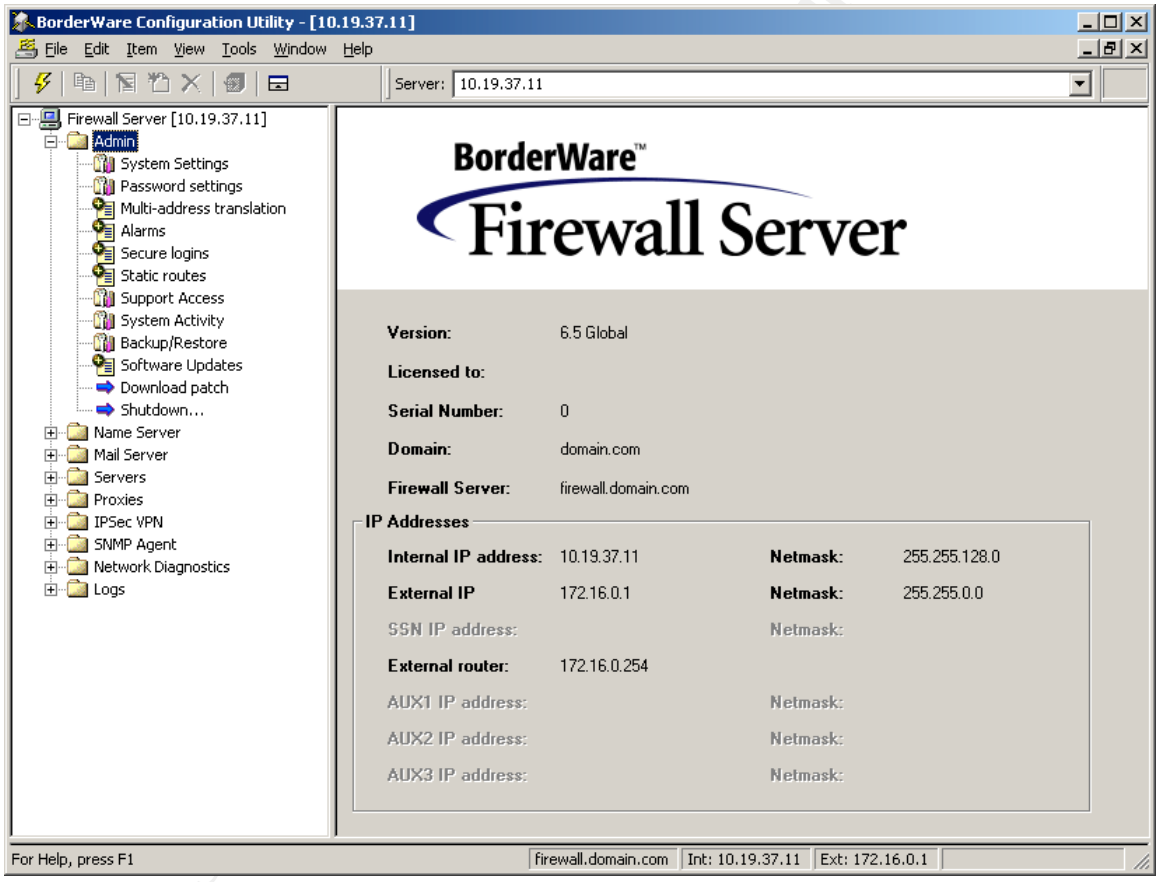

Shutdown allows you to either shutdown your firewall all together or reboots the firewall from the BWClient.

The Multi-address Translation screen is where you start to create separate external IP addresses that translate to servers located on either the Secure Server Network (Sometimes referred to as the DMZ) or the internal network. This section will be covered in later sections that deal with this process.

The Alarms Section deals with the various alarms that can be set to detect probes on various ports. Alarms can only be set for ports that are not in use. For instance, an Alarm cannot be set for Port 80, as this is the port that the firewall uses to pass HTTP traffic. If you specify a range of ports, the firewall will only monitor those ports that are not in use.

Setting up an Alarm is easy. Select the Alarm Option under the Admin Folder. Now click on Item from the toolbar and select Add New from the drop down menu. You could also click on the right hand side of the screen and then click on the blank piece of paper icon from the icon toolbar. Either way you will be lead to the Alarm Wizard Screen where you can begin to create your Alarm.

Figure 4: Adding an Alarm

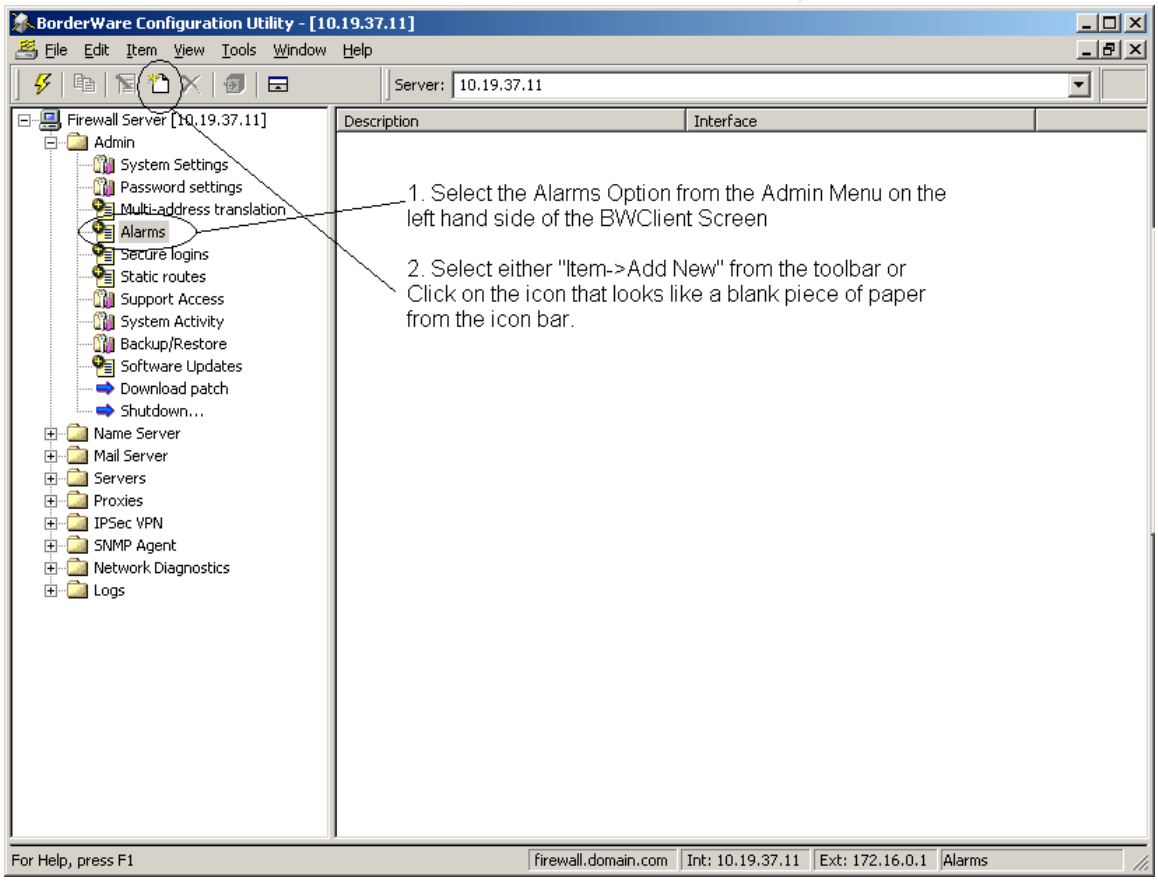

Specify a descriptive name for your alarm, which interface you would like this Alarm to examine, and whether you would like to look for TCP or UDP protocol scans. Clicking "Next" will move you to the next screen where you can specify if you would like to have the Alarm examine packets from all scanning machines or you may specify one specific source address. Moving forward thru the wizard allows you to specify the range of ports to monitor. If you wish to monitor a single port, you can limit the start range and end range to that single port to

accomplish that task. Next, you set the thresholds of the Alarm. This is done by the number of probes within a given time period. This value can be changed once the Alarm has been created to fine tune the Alarm. Lastly, you can specify the email address that should be notified once the alarm is activated. Once the Alarm has been created, it must be applied to the Firewall. This is done by both selecting File from the toolbar and then selecting the Apply Changes Option from the drop down menu or by clicking the icon on the icon toolbar that looks like a yellow box with a right aimed arrow moving across it. The Alarm will appear bolded until it has been applied to the Firewall and then it turns back to normal text.

The next section deals with setting up Secure Logins. This is the same process we covered when the initial user was created via the console. This time however, several more options are available. Creating a new user is done the same way you set up a new alarm. Now click on Item from the toolbar and select Add New from the drop down menu. You could also click on the right hand side of the screen and then click on the blank piece of paper icon from the icon toolbar.

Figure 5: Adding a Secure Login.

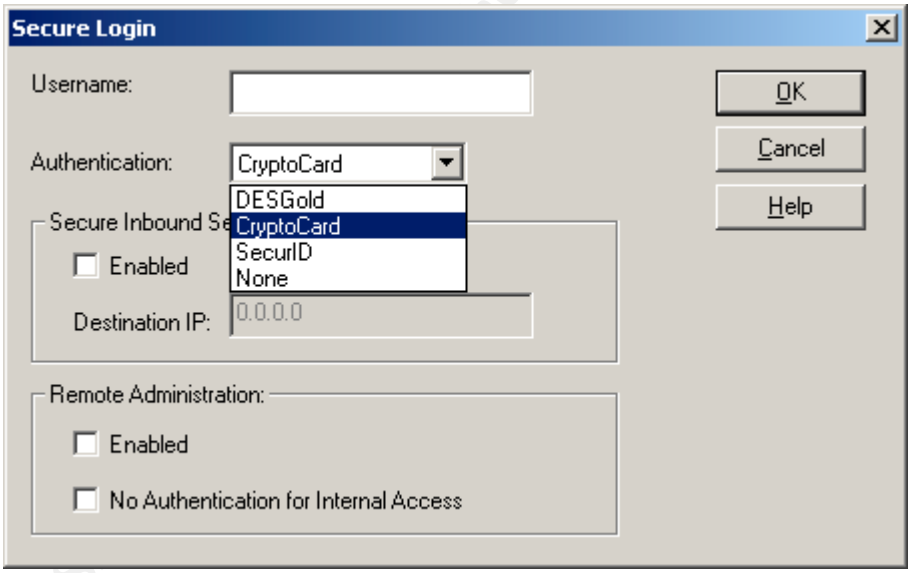

Figure 5 illustrates the screen you are shown when you add a new user. After adding the username, you are presented with more options for secure logins than you were shown via the console interface. The Borderware Firewall Server is designed to be configured and administered to via the BWClient so the options available thru the console interface are limited to only the basics. You can also grant users the right to access internal FTP and Telnet services on the internal network by enabling the Secure Inbound Services if you wish. You must specify a target IP Address for this service and access for these services will be limited to this machine. You must also enable this user to remotely administer the

Firewall. Remember all external access to the Firewall must be via SSL and must be done with authentication via a CryptoCard or another of the secure authentication options. External Administration cannot be accomplished by users whose authentication is set to none.

If you chose to have your users authenticate with DESGold, setup is very similar to the CryptoCard setup process. If however you wish to use the SecureID technology, you will need to have the appropriate technology running on your internal network to authenticate users with SecureID authentication. This is outside the scope of this document. Please refer to the documentation for SecureID about this process.

The Static Routes menu allows you to specify the gateway by which the firewall contacts those networks that are not directly connected to it. For instance, your business has a corporate WAN and an internal router that handles the traffic to and from these WAN links, you would simply list those networks and the interface of the router that will accept this traffic. If all of your networks are directly connected to the firewall via a Network Interface Card then you do not need to specify a static route. Unlike routers that can route traffic dynamically, if so configured, the Borderware Firewall Server does not support dynamic routing as a security feature.

The support access menu is the location you give Borderware Technical Support access to your firewall to assist you in troubleshooting. You are given the options of allowing internal access as well as inbound access. The option selected will depend upon the location of the firewall within your network. Most deployments of the firewall will need inbound access only. This feature can be turned on when needed and disabled when not needed. It is recommended that you leave this feature disabled until needed.

The Backup and Restore Menu allows you to backup the configuration of the Borderware Firewall Server. There are three options available. A text file can be downloaded to the FTP Area of the firewall, a text file can be sent via Email to a specified user and the file can be saved in XML format. From the BWClient you can only restore the configuration that was saved in XML format. You can create a backup of the configuration on a 5 ¼-inch floppy, Tape device or XML directly from the console as well. Text file configurations are made for archival purposes only and can be helpful for examining problems without switching between several screens.

Patches for the Borderware Firewall Server can be downloaded directly to the Firewall via the Download Patch Menu on the Borderware Client. Borderware also has several ways to save these patches both for archival purposes as well as for immediate use. Clicking on the Download Patch will bring up a Wizard that will assist you in this process.

Figure 6: Download Patch Wizard

You will be able to input your Support Contract Number when the Wizard first opens up. Your Serial number will also be displayed. This is the number that you provided when you registered your Borderware Firewall Server. An evaluation license will not be able to download patches directly in this manner. Once the valid support contract number has been inputted you can click next to progress to the next screen in the wizard. The next screen allows you to specify which server you would like the wizard to look for the updates on. It is recommended that for a vast majority of the installations that you leave this setting to the default, which is patchserver.borderware.com. If you wish to change this value, check with Borderware's Technical Support first.

The following screen asks you to verify the version level of your firewall and allows you to ask to be shown all patches. For the sake of brevity and ease of use, do not choose the show all patches option. You will be shown all the patches for this series of firewall and will have to know which patches are valid for your firewall. Figure 8 shows an example listing of available patches for the Borderware 6.5 series firewall.

Borderware uses a pretty straightforward nomenclature for the patches it deploys. An example would be the patch "fs65f0a." FS65 equates to Firewall Server 6.5. F0A equates to Feature Pack A. You will also see p01 and s01 as endings; they stand for service patch 1 and security patch 1 respectively. You can determine what your current software level is by opening the Software Updates menu and looking at the patches that are listed.

Once a software patch has been selected you have the option of saving that file to an area on the local machine the BWClient is running from, to the FTP Area of

the Firewall and to create a set of diskettes for the Patch to be installed from. It is recommended that you simply save the patch to the FTP Area of the Firewall. If you are going to save these files to your local machine so you can burn them to a CD-ROM and have them installed on the firewall from there, you simply place these patches in a folder on the CD-ROM called Patches. The firewall will look for the patches only in the patches folder if you are installing them from the CD-ROM.

Figure 7: Patch Listing

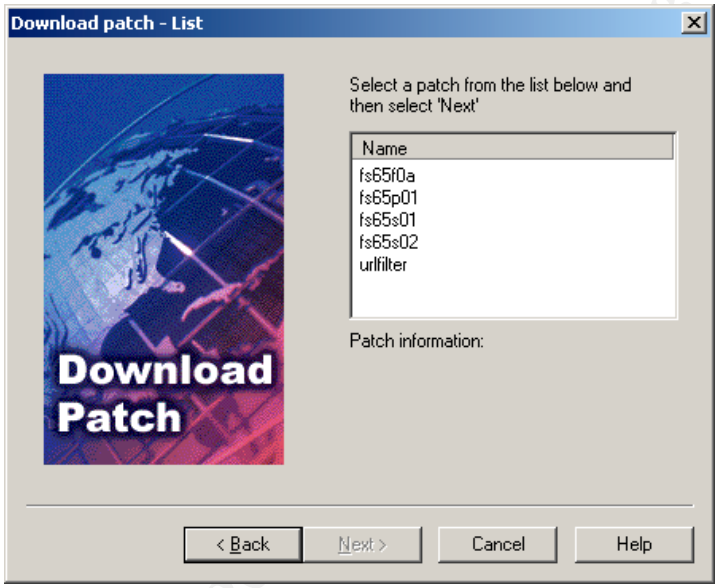

Patches can also be downloaded directly from the console however; the console does not have a status screen that shows you the current status of the download. As some of these patches can be rather large, it is often helpful to do download from the BWClient even if you intend to install them directly from the console. Once the patches are downloaded, you can install from the console or from the BWClient. If you elect to install from the client you must go to the Software Updates screen, right click on the update you wish to install and select install from the options available.

Patches being installed directly from the console can be done by choosing the Software Updates option under the Admin Folder. Chose to install patches and it will show you a list of patches that are currently in the FTP Area. If you chose to install from a CD-ROM it will move those patches from the CD over to the FTP Area and then you can continue to install. It is recommended that you always reinstall the Support Access Patch after every Patch install session. Patches should always be installed in the following order: Feature Packs, Service Packs and then Security Packs. Always read the release notes for each service patch, as they will enumerate the features affected by each patch.

The Name Server Folder

Setting up the Domain Name Server on the Firewall is rather simple and straightforward. Upon opening the Name Server Folder, several other screens become available. The first screen is the General Screen. It is on this screen that you set the Start of Authority settings for the domain(s) the firewall will host. It is suggested that you leave the default settings in place however they are configurable should your environment dictate other settings.

You should update the list of root servers when you first start the firewall but thereafter it should not be necessary. You can do this by clicking the box next to that option and by clicking the apply button at the top of the screen. The name server will restart once this is complete.

If this firewall will not be the primary DNS server then you will need to set the firewall, as a slave and it will forward the requests to the primary DNS server. This is normally done when the firewall is used to segment areas of an internal network. It is recommended that if the firewall is to be used as the boundary between the internet and the internal network that it is the primary DNS server. The list of forwarders can be found on the DNS Forwarders Menu under the Name Server Folder. It must be noted that the servers on the firewall (mail is a good example) cannot use any other DNS server other than the one on the firewall. If a request is made and the firewall is set to forward then the request will fail thus presenting problems with the delivery of mail. Adding a new entry is done in the standard Borderware fashion of right clicking on the white field and choosing “add new”; once the entry is entered you will need to apply the changes via the shortcut button on the toolbar.

Figure 8: The Name Server General Menu

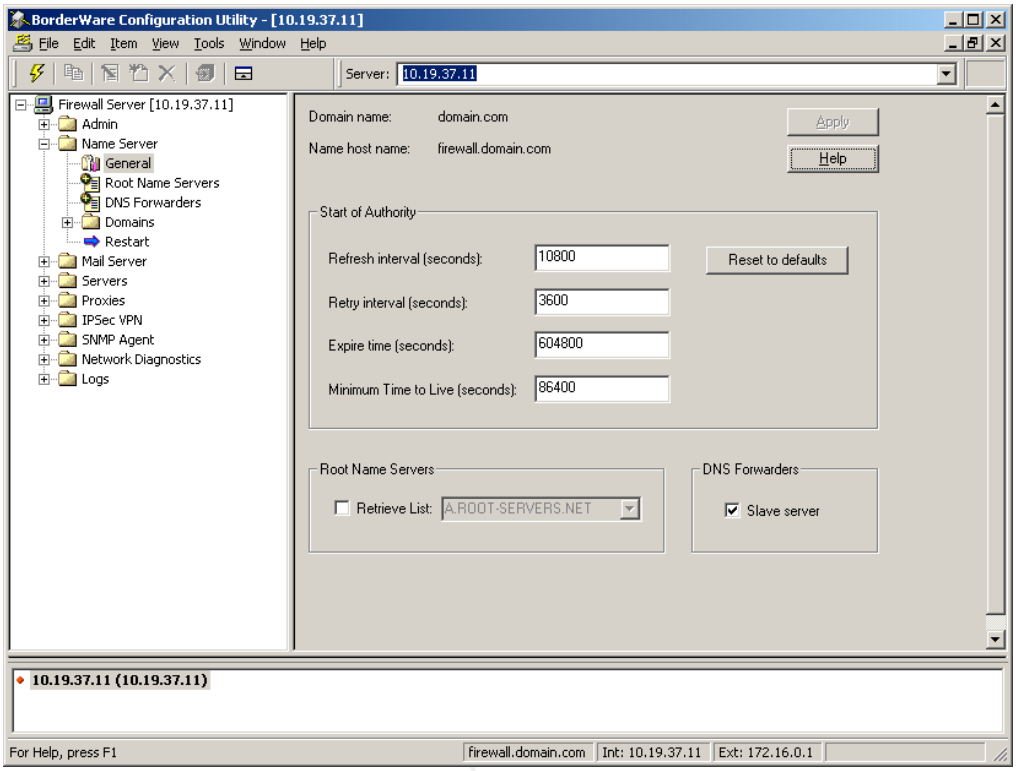

The Borderware Firewall Server has two separate Domain Name Servers: an Internal and an External. These are two separate processes. The Internal DNS can resolve DNS requests for requesters within your internal network. If the Internal DNS cannot resolve a request then that request is passed to the External DNS server to resolve the request. The reverse is not true however. If a request is made to the External DNS server and the External DNS server cannot resolve it then a response is sent back to the requester indicating such. No requests can be passed from the external server to the internal server for security reasons.

Adding a domain to either server is straightforward. Right clicking on the white area on the right side of the screen will allow the addition of a domain. Three choices are available when adding a new domain. The firewall can act as primary for the domain record, can indicate that it is secondary for the domain (the information about the primary is required when this is done) or it can delegate authority for that domain (again the information about the primary is required at this time.) Once an entry is added, the changes need to be applied by clicking on the shortcut button on the toolbar.

When a domain is added, it should be modified to reflect the proper name server. In most cases, this name server will be the firewall itself but not always. It is for this reason that a name server must be identified for each domain that is entered into the record. Secondary name servers can also be entered.

An MX record can be recorded at this time as well. It is important that this information be correct if the domain is going to receive mail. If the domain is not going to receive mail then this step can be eliminated. Preferences can be set for multiple mail servers to provide redundancy in mail delivery.

Both the Internal and External DNS servers are configured the same way. It is important to remember which server you are modifying in order to maintain security and to ensure that the entries work properly. Keeping an eye on the left hand menu bar will help to remind you as to which server you are modifying.

Entering reverse lookup records or PTR records are not necessary as this is done automatically for each domain that the firewall hosts when a forward record is created.

The Mail Server Folder

The Borderware Firewall Server comes complete with two mail servers. One is a SMTP Mail server that passes mail from external sources to an internal mail server and vice versa and the other is a POP Mail server that stores the messages at the firewall itself. The POP Mail server allows internal users to connect using the popular POP Mail clients to retrieve mail. If you wish to allow external users to access mail held by this server a VPN or other remote access solution needs to be implemented to allow access to the internal network first.

Borderware has also provided a proxy to allow relaying of mail from the internal network out to the external Internet in the event an administrator experiences problems configuring his mail server to deliver the mail itself thru this proxy. This however will disable features available in the built in SMTP Mail server.

It is important to understand how the Borderware Firewall Server processes mail before talking about the features that are available.

When a message is received on the external, interface the Borderware Firewall Server first checks to see if it has an Alias entry for the recipient. The Alias entry allows the delivery of the message to go to several recipients thru the entry of only a single email address. If an Alias entry is detected, the message is broken up into separate messages for each of the persons listed as being apart of the Alias and then processing continues. If an Alias is not detected the message is also passed along for further processing.

The processing system then checks to see if a local POP account is present for the recipient. If an account is detected the message is delivered to this account and processing stops; if an account is not detected the message continues to be processed.

The processing system then checks to see if a Mail Mapping exists. If a Mail Mapping does exist then the mapping is applied and the message continues to be processed; if a Mail Mapping does not exist then the message continues to be processed.

The processing system now checks for pre defined routes by which to deliver the mail. If a route is identified then the message is delivered via that route. If a route is not defined then an error message is sent back to the sender informing them the message could not be delivered.

It should be noted that the Borderware Firewall Server does not perform an MX record lookup at this last stage. This prevents the relay of mail by someone who sends mail directly to the firewall. Should someone send mail directly to the firewall and use an invalid sending address, the mail will sit in the queue for three days before its automatically deleted.

The Borderware Firewall Server does have a feature that will prevent the relaying of mail as described above. A simple checkbox on the Mail Server -> General menu allows the relaying of mail off the external interface of the firewall (Require FQDN.) With this option enabled, the firewall will only accept mail if mail domain is defined on the firewall. If there is not a mail domain defined, the firewall server will reject the message. It should be further pointed out that this option only effects mail received on the external interface. All mail received from the internal interface is unaffected.

Figure 9: Mail Processing Flow Chart.

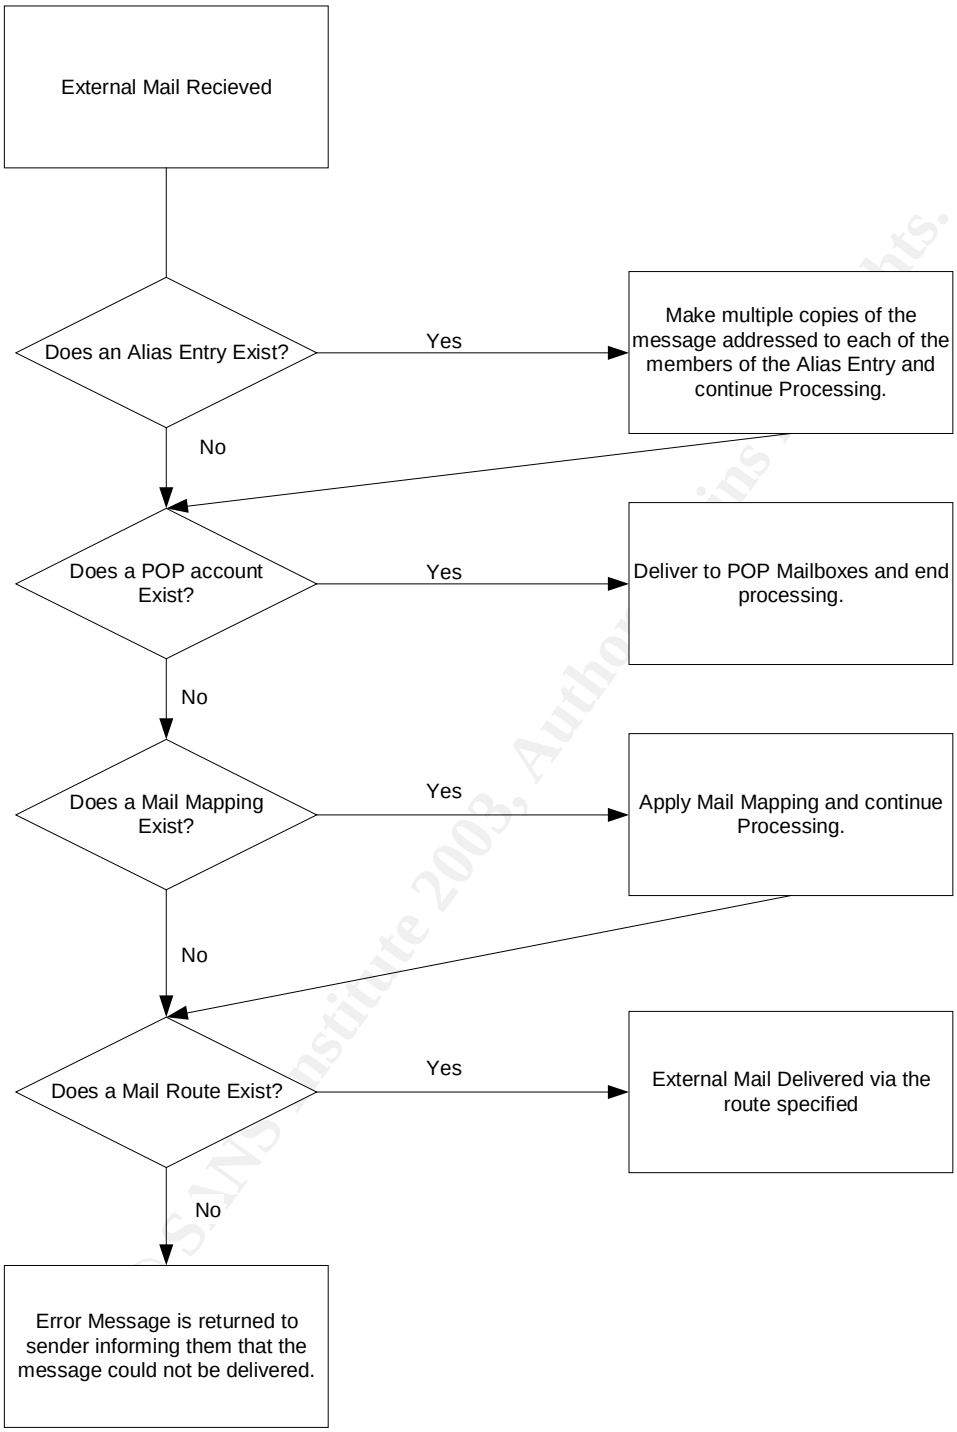

Figure 10: The General Menu for the Mail Server.

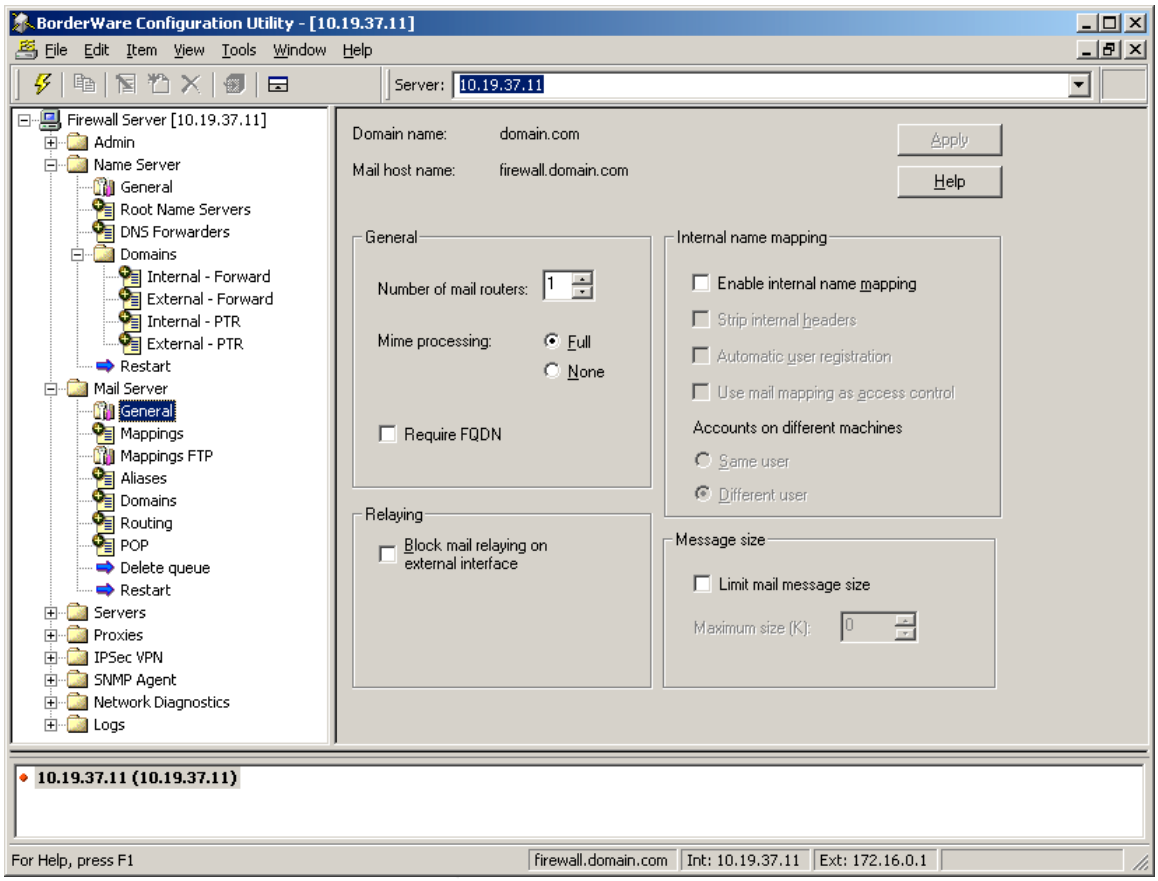

General Setup of the Mail Server is straightforward and simple. The number of mail routers used depends on the amount of mail traffic the firewall is expected to handle. Increasing this number allows separate processes to run in the delivery of mail but also affects the overall performance of the machine.

You can also choose to require a Fully Qualified Domain Name (FQDN) be present in all the mail that you receive. Choosing this option will reject any messages that do not report to have a FQDN in the “From:” field of the message. This can have the affect of cutting down on the amount of SPAM your domain receives.

An option to enable Internal Mail Mapping is available from this screen. This option allows you to apply the Mail Mappings you set in a following menu to be applied to both internal and outbound messages. Once this option is enabled, several other options become available. These include the ability to strip internal headers. These headers can pose a security risk if the information is relayed outside of the internal network. You can also automatically generate mail mappings. This is done when a user sends a message thru the firewall from

internal to external for the first time. The internal host is stripped to reflect user@domain.com instead of user@host1.domain.com. If the user sends another message from a different host, another mapping will be applied unless the account on different machines option is selected. With this option, you have the choice of allowing another mail mapping to be made by selecting the option that tells the firewall that the message from a different host is another user. By selecting the same user option, the firewall does not create another mapping because you have told it that this would be the same user.

The danger with automatic registration is that mappings are automatically created without the administrator's knowledge. This feature can be helpful when you first set up the firewall. However, it may pose security problems later. Once you have a stable Mail Mappings Table, be it one created automatically thru the procedure just described or one created manually, you should enable the option that uses mail mappings as access control. This way only, those senders with a current mail mapping will be allowed to send messages outside of your network. This can be a method of control over temporary employees who need access to email for their job but do not need to communicate outside of the domain itself.

You also have the option of limiting the size of messages that pass thru the firewall in either direction. This setting is based upon the Kilobyte size of the message and is pretty straightforward.

Mail Mappings can be created manually rather simply by selecting the Mappings Menu under the Mail Server heading. An administrator should either right click on the white field on the right or select Item->Add New from the drop down menu at the top of the screen.

The Server Folder

The server processes are part of the Firewall Servers core system. They are specifically designed to protect against a client on one side of the firewall from knowing too much about a server on the other side of the firewall. The servers are also one of the first places an intruder will attempt action against. Borderware has designed servers to provide a variety of services on each of the main interfaces. There are three kinds of servers, External, Internal and SSN. Each provides services unique to their environments. The firewall protects against spoofing by denying requests from a network not associated with the interface the request was received upon. Thus a request received from the external network, which reports to be from the internal network, will be denied.

The Borderware Firewall Server comes with this list of available servers:

Type of Server	Description
DNS Queries	Responds to DNS requests. As there are only two DNS Servers, an internal and an external, the SSN Interface shares the Internal DNS Server.
DNS Zone Transfers	Allows transfer of DNS zone files.
Finger	Responds to Finger requests. The message returned is configurable.
FTP	From the external network, only Anonymous FTP is allowed. Admin and Anonymous are permitted from the internal network.
GUI Config	Allows connections from BWClient - only from the internal network
Secure GUI Config	Allows authenticated connections by BWClient from any interface
Ident	A protocol used by some security systems to check on the identity of a user requesting a service before that service is granted. The firewall responds to Ident request with a dummy "hidden-user" reply. It is recommended that you enable this on the external interface to satisfy Ident requests from external servers, Otherwise, connections to some sites may be slower, and your logs may fill with failed Ident request.
Log Rejected packets	Generally a debugging tool that is enabled for any particular server on any particular interface.
Ping	Allows the firewall to respond to ICMP echo requests (pings).
POP Mail	Allows users to access their POP mailboxes on the firewall (only from internal network).
SMTP Mail	Handles the receipt and delivery of mail
Traceroute response	Responds to traceroute requests
WWW	A set of simple static web pages can be placed on the firewall.

Servers are enabled in a straightforward manner. Select the interface you wish to enable the server on from the Server folder and double click. You must provide the following information to enable the server.

- Check the option to enable the server
- If you wish to log all packets for the specific service, you can click to Log Data Packets. This information is stored in the General Messages Log File.
- You can select the maximum number of sessions that can be concurrent.
- You can set up access rules for each service.
- If you chose the FTP or Anonymous FTP, the FTP Writable option appears. This would allow users to transfer files into the Incoming Directory of the FTP area of the Firewall.

Setting up and configuring the FTP area is described in the Help Files located in the BW Client Program. It is similar to all other FTP servers and thus will not be covered here.

NTP is another server that is available on the Borderware Firewall Server. It is a service that can only be activated from the console so you will not find any mention of it in the BW Client except an entry in the help files. The Borderware

Firewall Server can be configured as either a client or a server for NTP. From the console, select the Services Menu and then select the Configure NTP Option. If you want the firewall to act as a client, select the Add NTP Server and enter the IP address of the timeserver you wish to sue and its version. The comment area can store the email address of the contact person of the NTP Server. If the time of the firewall is off by more than a few minutes from the NTP Server, a reboot will be required for the new time to be set.

You can track NTP events in the General Messages Log file. Client entries begin with “ntp” and server entries begin with “xntp”

If you wish, the firewall to act as a NTP Server you must select Enable/Disable NTP Service and a reboot will be required. The version of NTP being used is 3.

The WWW server included with the Borderware Firewall Server is a simple web server with limited functionality. It is designed for a static environment. Use of CGI Script is not supported. Web pages can be added to the firewall by connecting to the firewall via FTP and changing directories to the /servers/WWW/data directory. Files can be placed here and subdirectories can be created to keep the site organized. Your home page must be called index.html. If this file does not exist, the server will present the requesting IP with a list of files in the /servers/WWW/data directory.

The best solution to providing a web page is to configure a dedicated web server, locate it on the SSN segment of your network, and configure the firewall to pass all web requests to this server.

If you wish to run the Finger Server and wish to configure the response sent, back to the requestor you can modify the file “Config” in the /servers/finger/data directory, which can be accessed via FTP.

Figure 11: Adding a Mail Mapping

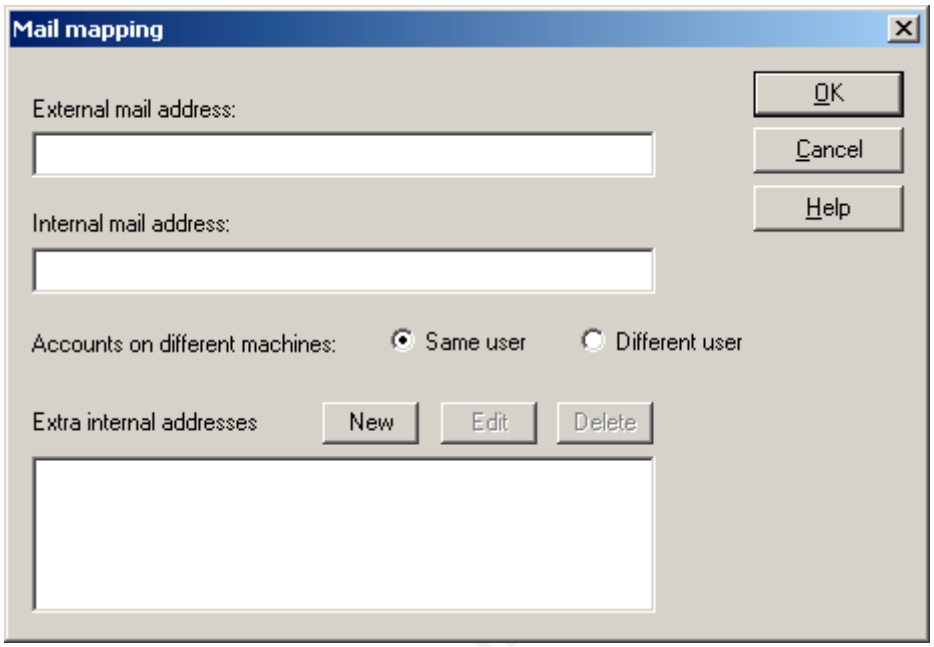

An Administrator can both add a Mail Mapping and edit a Mail Mapping on the same screen. An administrator must simply supply the external address, the internal address and define if a message received from different internal hosts should be considered the same user or a different user. Once the Mapping has been created, the administrator will need to apply the changes to the Mail Mappings Table by either the shortcut button on the toolbar (the yellow box with a right facing arrow on it) or by selecting Item->Apply Changes from the drop down menu at the top of the screen. The Mail Mappings Table can be downloaded and uploaded from the FTP area for offline administration. The file is named mailmaps_download and is located in the /load_db directory of the FTP Area. This is done from the Mappings FTP menu and needs the FTP password to be supplied before the options become available.

Adding an Alias is equally simple. Entering the Alias entry screen is done the same way a mail mapping is added and allows the delivery of the message to a specific address or to just be deleted. Borderware Firewall Server comes with several pre-defined Aliases. A pre-defined Alias cannot be deleted but some of them can be modified.
The Pre-defined Aliases that cannot be modified are:

- Devnull
- Bitbucket

The Pre-defined Aliases that can be modified are:

- Postmaster
- Postoffice
- Mailer-daemon
- Maildrop
- Root
- Usenet
- News

Pre-defined Aliases which can be modified, can be modified to only one destination however this can be circumvented by adding an additional alias should it be needed. An example of this would be

- Admin -> user1@host1
- Admin -> user2@host2

Adding a Mail Domain is done thru the same procedure as used with Mail Mappings and Aliases and is used to specify the different mail domains that the firewall will be accepting mail for. An example of this would be adding the mail domains of domain.net and domain.org for those organizations that hold the rights to both of those domains in addition to the domain.com which the firewall is automatically set up for.

Adding a route is done in the same fashion as all other additions we have discussed.

Figure 12: Adding a Mail Route

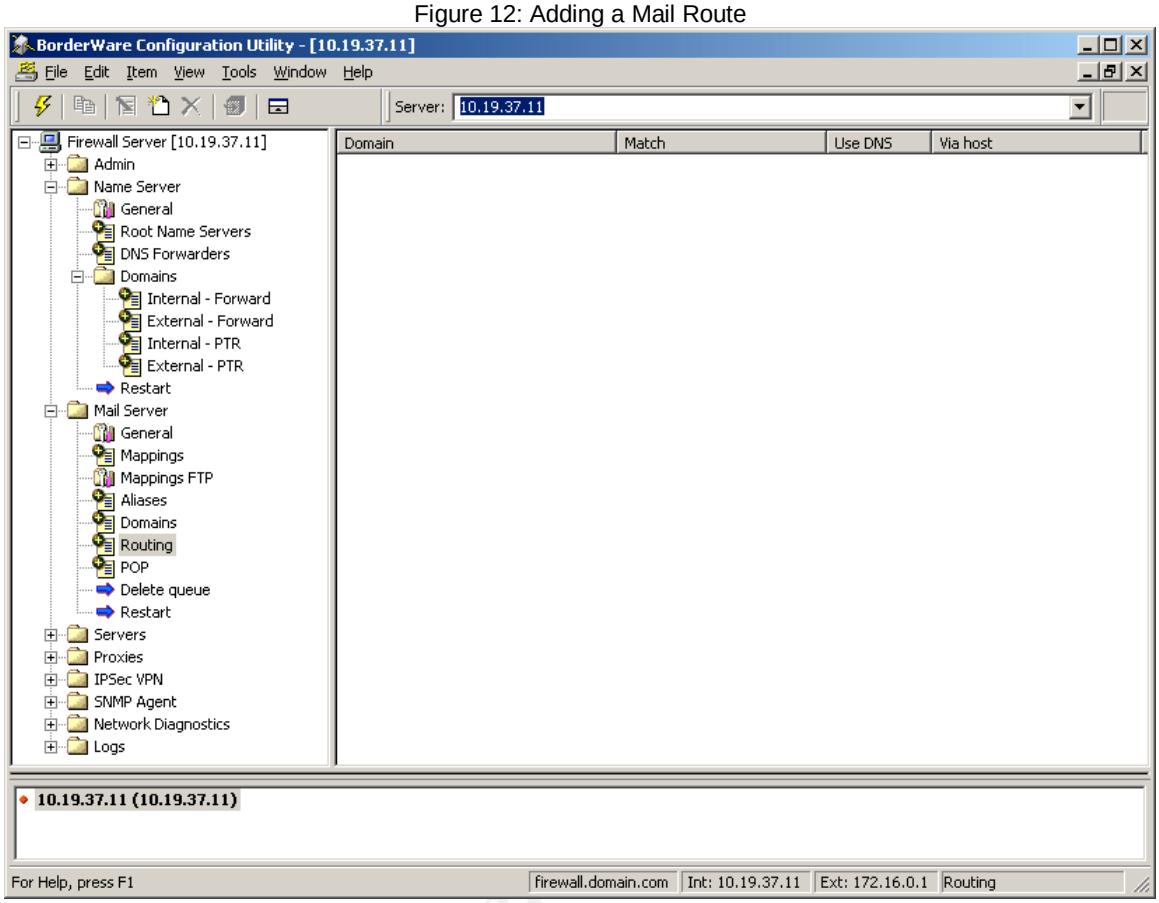

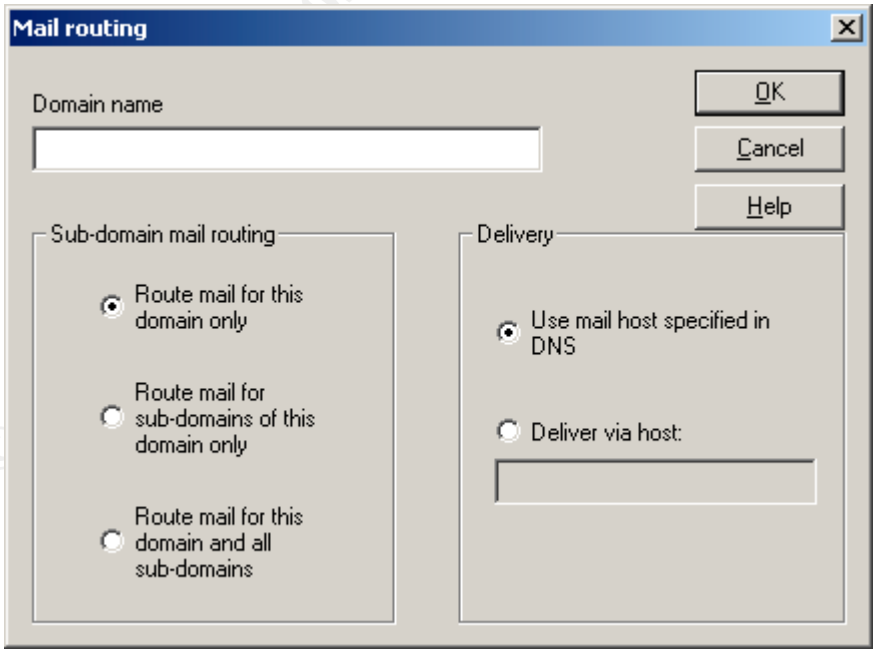

The administrator simply needs to supply the domain name for which the mail will be accepted, determine Sub-domain options as shown in the illustration and determine if mail should be delivered to one specific host or is a DNS lookup

should be preformed to facilitate the delivery. Clicking OK and then apply the changes made will activate these choices.

The last item to be concerned with the Mail Server is the POP component. Addition of a POP account is done exactly the same way as all the other additions we have made in each category.

Figure 13: Adding a POP Account

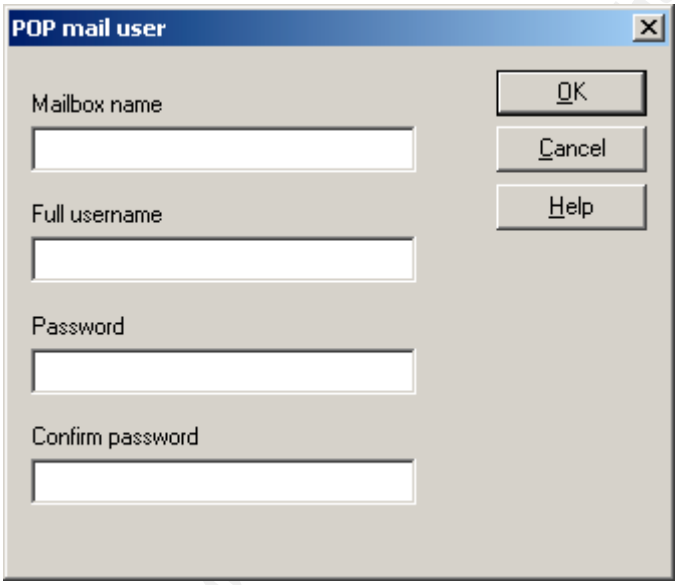

The administrator supplies the information required and applies the changes in the normal way. One word of caution here; setting up POP accounts on the Firewall server should be considered prior to the initial setup of the Firewall Server as partition space should be allocated for this use. Please note that size quotas are not available for each account. The size of each POP account can be obtained by viewing the View Pop Accounts option from the console's Mail Menu.

Mail Server Maintenance can be conducted from a different menu. It will be discussed in this section as it pertains directly to the current discussion.

The Administrator can view the mail in the queue from the BWClient interface by selecting the Mail Server option under the Network Diagnostics Folder. The option to delete the entire queue is available from the Mail Server Menu at the bottom. Should an Administrator wish to delete only selected messages from the queue, the administrator should do so directly from the console interface. This is done by selecting the Delete Mail from Queue option from the Mail Menu and supplying the message ID number of the message in question. This number can be obtained from the View Mail in Queue option of the Mail Menu on the console as well. Using the console, the administrator can only delete one message at a time where as the queue can be deleted in its entirety from the BWClient.

The Proxy Folder

The proxies are the heart of the Borderware Firewall Server. One point is key to the understanding of proxies. They work in connection with the servers to provide functionality to the firewall server and utilize access rules. It can be broken down as simply as this. A server is a process that can answer a request directly where as a proxy is a process that relays a request to a server on another machine. When working with the processes it is key to think of servers as resources that the firewall provides directly to the user where as a proxy is a conduit that transfers a request from one side of the firewall to the other. This function makes the Firewall Server the secure firewall that it is. In basic layman's terms, the proxy server functions just like a surrogate in the delivery of information. As was discussed in the first section, a client makes a request for information from a server on the other side of the firewall. The firewall acknowledges the connection and establishes a connection with the requesting client. At this point, the Firewall makes its own connection with the resource and repeats the request. Once the firewall obtains the requested information it passes this on to the requesting client. At no time do the requesting client and the resource talk directly.

One of the direct benefits of this type of system is that connections are controlled by the presence and absences of specific proxies. If a client requests access to information that does not have a corresponding proxy enabled on the firewall then no access can be granted to the requesting client. This helps to protect the internal resources should services be enabled accidentally or for internal use only.

Borderware allows full control over the proxies that you can enable on the Firewall Server. This is done by selecting the direction of the proxy from the point of origin to the destination point. Your selection is based upon the number of interfaces you have configured for your firewall. In order to simplify the explanation, the example used will be that only two interfaces are configured on the firewall but as Borderware supports up to six interfaces the example can be extrapolated out to that many interfaces.

Figure 14: Proxy Server Initial Screen

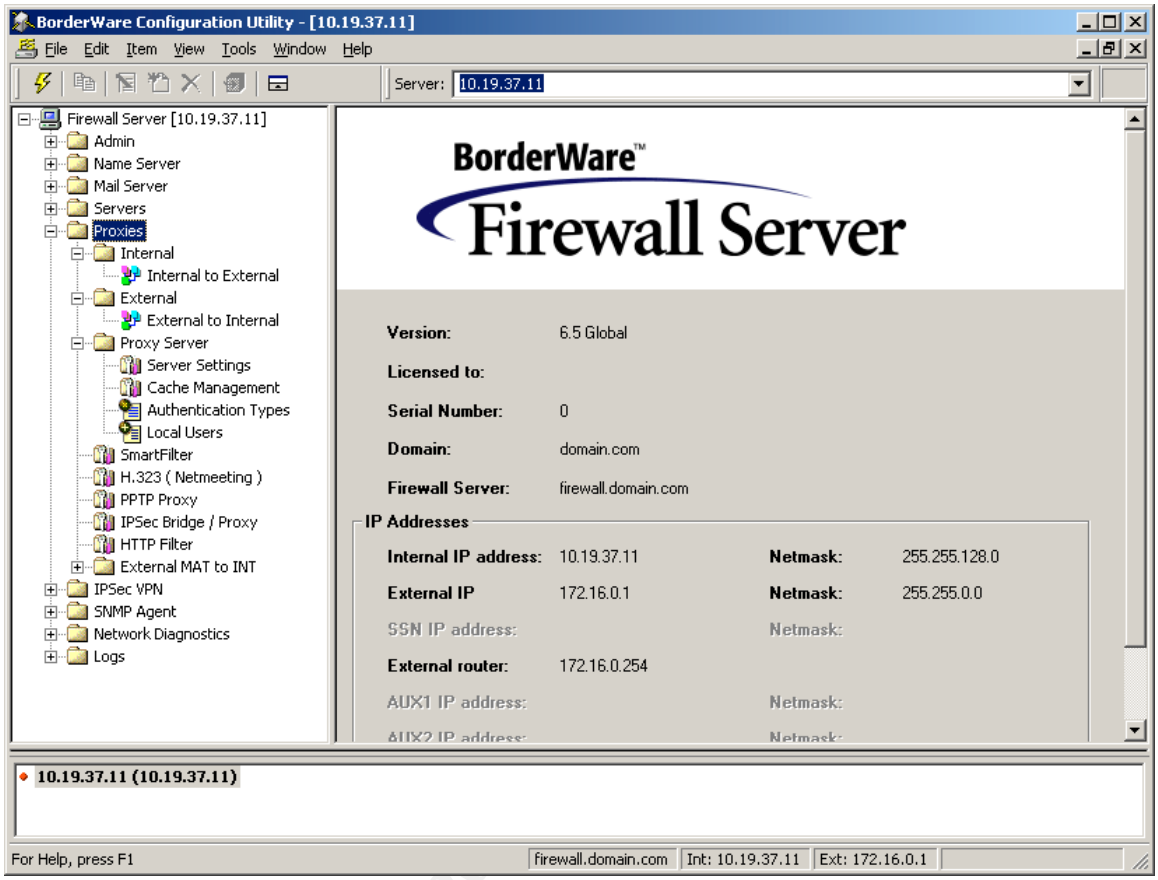

As seen in Figure 1, the administrator is given several choices for the direction of his proxies based upon the interface that originates the communication. While enabling a proxy is considered risky, it is the author's opinion that the amount of risk is associated not so much with the proxy itself but with the direction, the proxy is enabled in. An example of this is opening Port 21 from internal to external so that some internal users can access a FTP site that is located on the Internet. Opening this port presents a risk but the risk is mitigated by the fact that this proxy is only being allowed to function when the client on the internal network initiates the communication. While the communication occurring is bi-directional this is only true on a per communication basis. In our example, should the internal client end communications with the FTP server external to the network, communication can only be reestablished if initiated by the internal client again.

Opening a proxy is easy on the Borderware Firewall Server. First, the administrator must determine the direction communication will originate from. Borderware allows two basic types of proxies: Pre defined and generic proxies. Pre-defined proxies are those proxies that will log and examine packets being transmitted thru the firewall for all seven layers of the OSI Stack where as

generic proxies will only examine the packets thru layer three. It is recommended that pre-defined proxies be used whenever possible as this provides that extra level of security.

Enabling a proxy is straightforward. If the proxy is already defined and listed on the interface, right clicking on the proxy and choosing "Modify" from the drop down menu will allow you to configure it. Several options are available at this point.

Figure 15: Modifying a proxy.

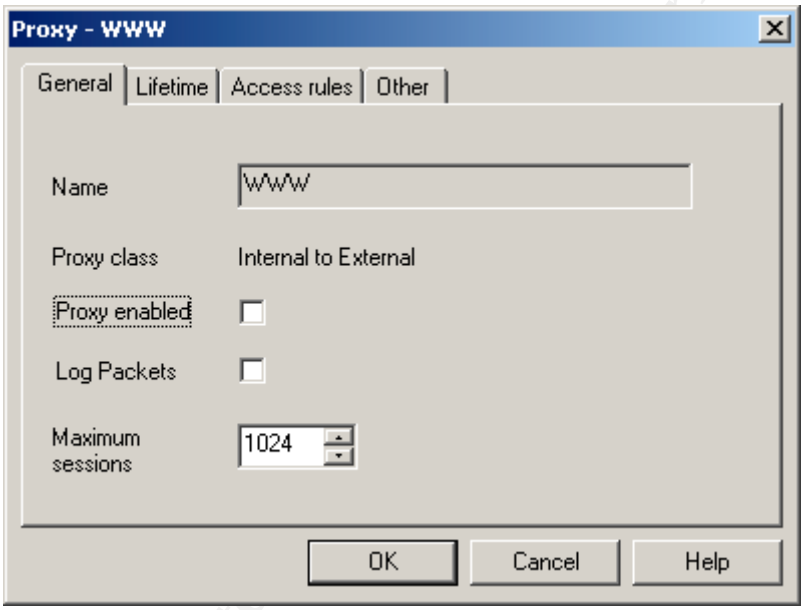

The administrator has several screens to choose from to effect changes on this particular proxy. Figure 2 shows the modification screen for a pre-defined WWW or Port 80 proxy. The general tab portion of this window allows the administrator to limit the number of sessions, establish logging on this port and, enable and disable the proxy.

Figure 16: The Lifetime Tab

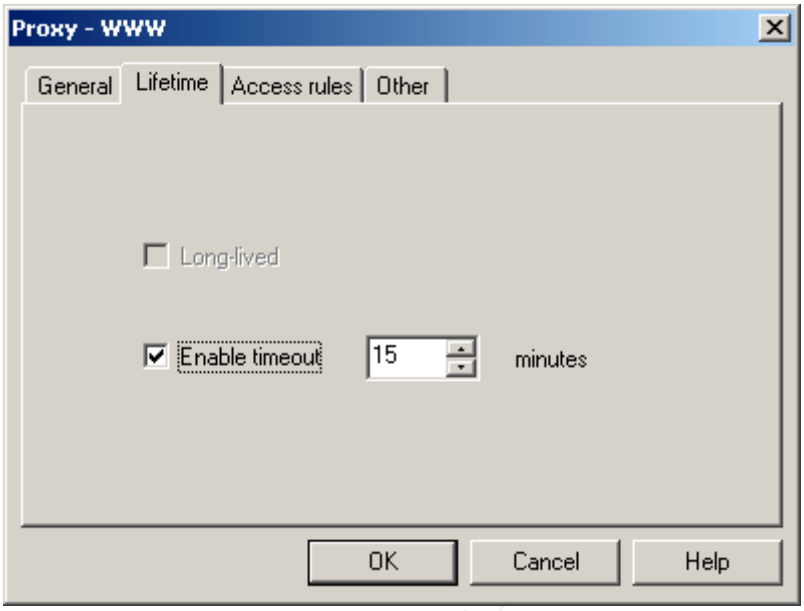

The Lifetime tab allows the administrator to set limits on the time a connection can remain open.

Figure 17: The Access Rules Tab

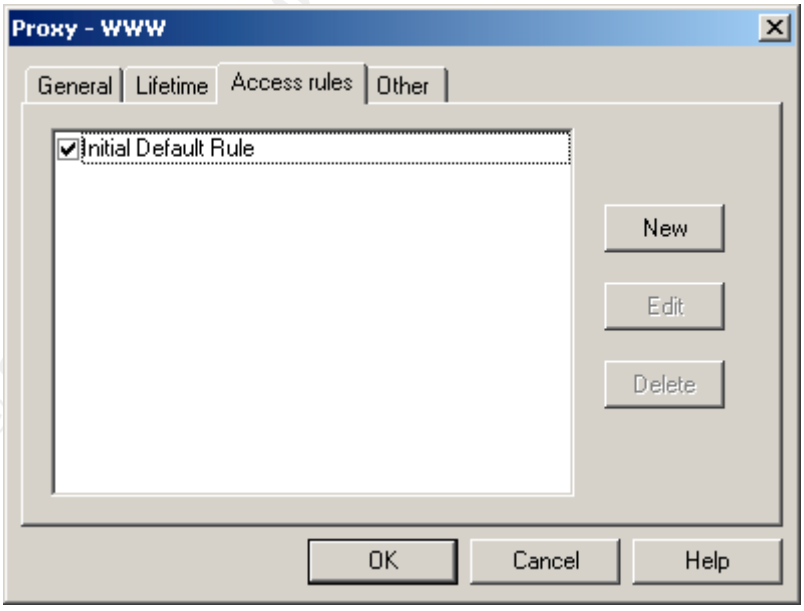

The Access Rules Tab is the heart of the proxy itself. It is here that the administrator determines who should have access to the resource the proxy is

intended to support. To enable an access rule, make sure the box next to it is checked and then click ok. To create a new access rule, click new and fill out the following information.

- **Rule Description:** Type a unique description that explains the rule. Be as specific as you can. Experience has shown that as time goes by the reason for certain rules can often be forgotten. The description, along with the status of enabled or disabled appears in all access rule lists.
- **Enable in the Rule:**
- **Service Class**: Select the service class to apply the rule to.
- **Max Sessions**: Specify the maximum number of concurrent sessions.

You can now specify the time of day you would like the rule to apply to. If a time of day is not specified, the firewall will apply the rule constantly.

You can specify the source and destination IP address the rule pertains to.

- **All:** Select this option if you want to allow or deny access to all IP addresses when this rule is applicable.
- **None**: Select this option if you want to allow or deny access to no IP addresses when this rule is applicable.
- **Only**: Select if you want only the addresses you have specified to have access.
- **IP Address:** field, enter the IP address to be allowed or denied.
- **Mask:** field, enter the netmask to apply to the IP address you are restricting.

An Administrator can also specify a range of IP addresses if so desired. This can be done simply by specifying the entire range in this manner. Specify the range 172.16.4.0 with a mask of 255.255.255.0. The access rule would apply to the entire range of 172.16.4.0 to 172.16.4.255. If a single IP address is desired, it should be input into the IP Address field with a mask of 255.255.255.255.

Access rules are applied in this manner:

When you have decided which servers and proxies you are going to use, you should consider applying access rules. When a request for a connection arrives, the firewall takes the following action:

- The firewall receives a packet and checks the port and destination address to see if they are consistent with a currently enabled server or proxy. If they are not the packet is rejected. If they are processing continues

- The number of current sessions is checked against the maximum set for that service. If the maximum is exceeded, the connection is denied. If they are not the processing continues.
- The access rules are examined. The time is checked against access restrictions specified in the access rule and the source and/or destination address is checked. If the packet conforms to the restrictions it is allowed; if it does not it is denied.

If no access rules are specified then no access is allowed. The Initial Default Rule allows all access at all times of day with no session restriction. If you specify other rules, the Initial Default Rule should be disabled.

Access rules can be applied to the following classes of servers and proxies:

- Internal Servers
- External Servers
- SSN Servers
- Internal to External Proxies
- External to Internal Proxies
- Internal to SSN Proxies
- External to SSN Proxies
- SSN to Internal Proxies
- SSN to External Proxies

The Proxy Server Folder

In addition to the standard proxies, the Borderware Firewall Server contains an embedded Proxy Server. This Proxy Server enhances the capabilities found in standard Proxies. It provides improved reporting of traffic (Web, FTP, Gopher or WAIS) through the firewall, user authentication, caching of web pages and enhanced performance. It is based upon the Squid Web Proxy (www.squid-cache.org.) The added functionality of being able to filter Java, JavaScript and Active-X is also available thru the Proxy Server.

The Proxy Server can be run in one of two modes. The first mode, Transparent Mode, is invisible to the end user and requires no modification of the end users web browser. The second mode, Non-Transparent Mode, requires modification of the end users web browser to access the internet. Communication is conducted via Port 3128 and requires more effort to implement because each

desktop must be touched in the process. It does provide authentication for all outbound traffic which cannot be done in Transparent Mode.

Figure 18: The Proxy Server Settings Menu

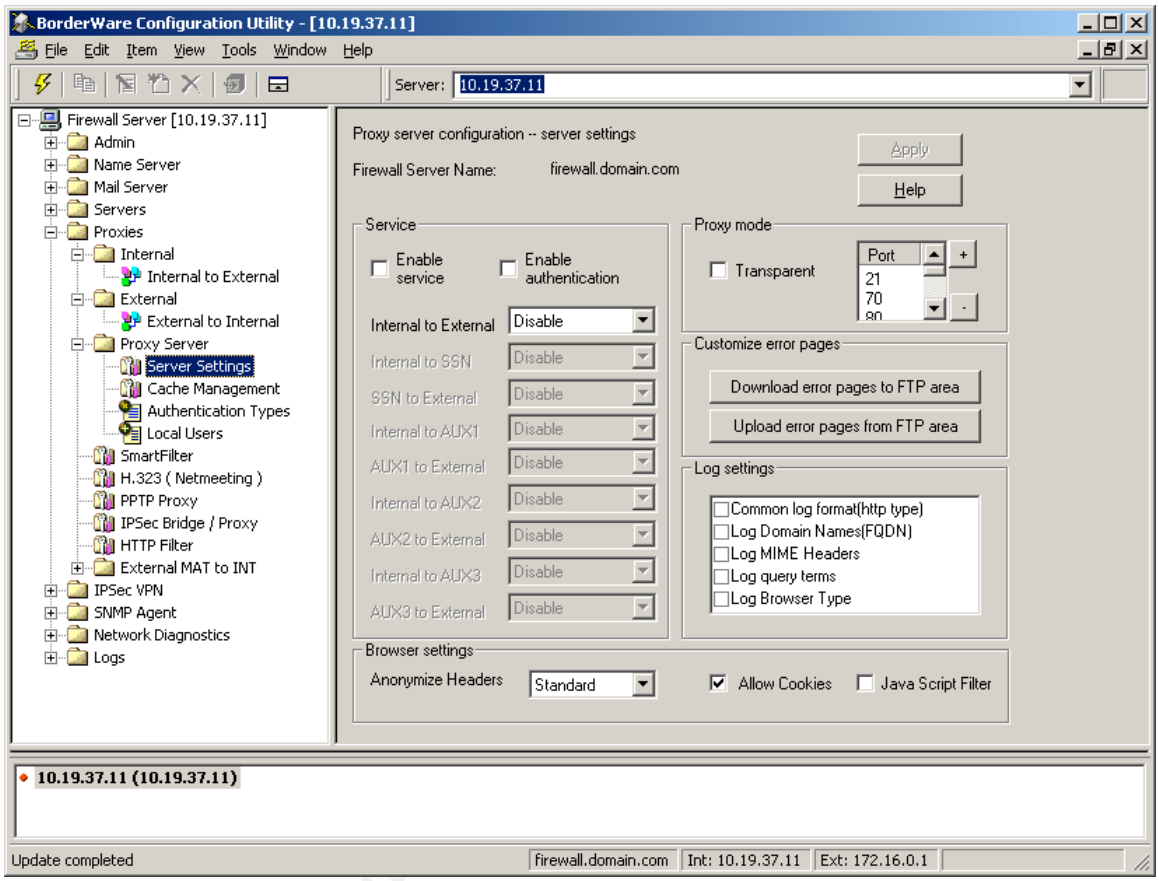

Let us examine both ways to implement the Proxy Server starting with Transparent Mode. First you will notice on the Server Settings Menu of the Proxy Server Folder several areas. They are the Service area, Proxy mode area, Customized error pages area, Log settings area and the Browser settings area. In the Service area you can enable and disable the service and then specify which options you would like to implement for each interface on the firewall. The service can be disabled, enabled without cache, or enabled with cache on each interface. The choice made depends upon the environment the firewall will be deployed in. Within the Proxy mode, area the box must be checked for transparent mode. Within the Customize error pages area you can upload or download the error pages that will be displayed to the user from the FTP area. This allows for the custom creation of error pages specific to the environment with pertinent contact information. The files will be located in the /load_db/errors/ area and can be modified on the administrator's workstation. They can be loaded back into this area and uploaded into the Proxy Server. These pages are

standard HTML pages and their file names must be preserved. The Proxy Server must be restarted before these pages become active.

You have the choice of enhanced logging with the Proxy Server running and several options are available to add to the logs.

The Browser settings area allows for several customizable settings. The first is to allow or disallow cookies. This function works exactly like it would on a stand alone browser except this setting affects all browsers using the proxy server to make web requests. If you set the Proxy Server to not allow cookies then some web sites may not work properly. A decision to do this should be based upon a need specific to the administrator's environment.

The Browser settings area also allows for decision to be made about header information being sent to the web sites being accessed. It is possible that a malicious web site may abuse the information being sent to it and launch an attack on known weaknesses based upon this header information. The Proxy Server can block this information from being transmitted. This is done thru a three level setting:

Off. All information will be sent to the requesting web site. No blocking will be preformed.

Standard. All headers except the following will be allowed:

- "From"
- "Referrer"
- "Server"
- "User-Agent"
- "WWW-Authenticate"
- "Link"

Paranoid. All headers except the following will be blocked:

- "Allow"
- "Authorization"
- "Cache-Control"
- "Content-Encoding"
- "Content-Length"
- "Content-Type"
- "Date"
- "Expires"
- "Host"
- "If-Modified-Since"
- "Last-Modified"
- "Location"
- "Pragma"

- "Accept"
- "Accept-Encoding"
- "Accept-Language"
- "Content-Language"
- "Mime-Version"
- "Retry-After"
- "Title"
- "Connection"
- "Proxy-Connection"

While in Transparent Mode, The Proxy Server handles requests made on Port 80 only. Because of this conflicts can arise with other proxies that are enabled on Port 80 in the same direction. Therefore the WWW proxy must be disabled in order for the Proxy Server to be enabled.

In contrast, Non-Transparent Mode can handle traffic on many ports and is configurable. In the Proxy mode area, you can add or subtract ports with the “+” and “-“keys.
User Authentication is another feature of the Non-Transparent Mode. There are three methods of authentication that can be used: a local user list, a RADIUS server and/or an LDAP server. These methods can be used separately or can be combined as the administrator and the local security policies dictate. If multiple methods of authentication are configured, the Proxy Server will attempt each method until a successful authentication has occurred. The sequence is determined by the priority parameter set for each specific method and methods with a lower priority are attempted first.

Figure 19: The Proxy Server Authentication Types Menu

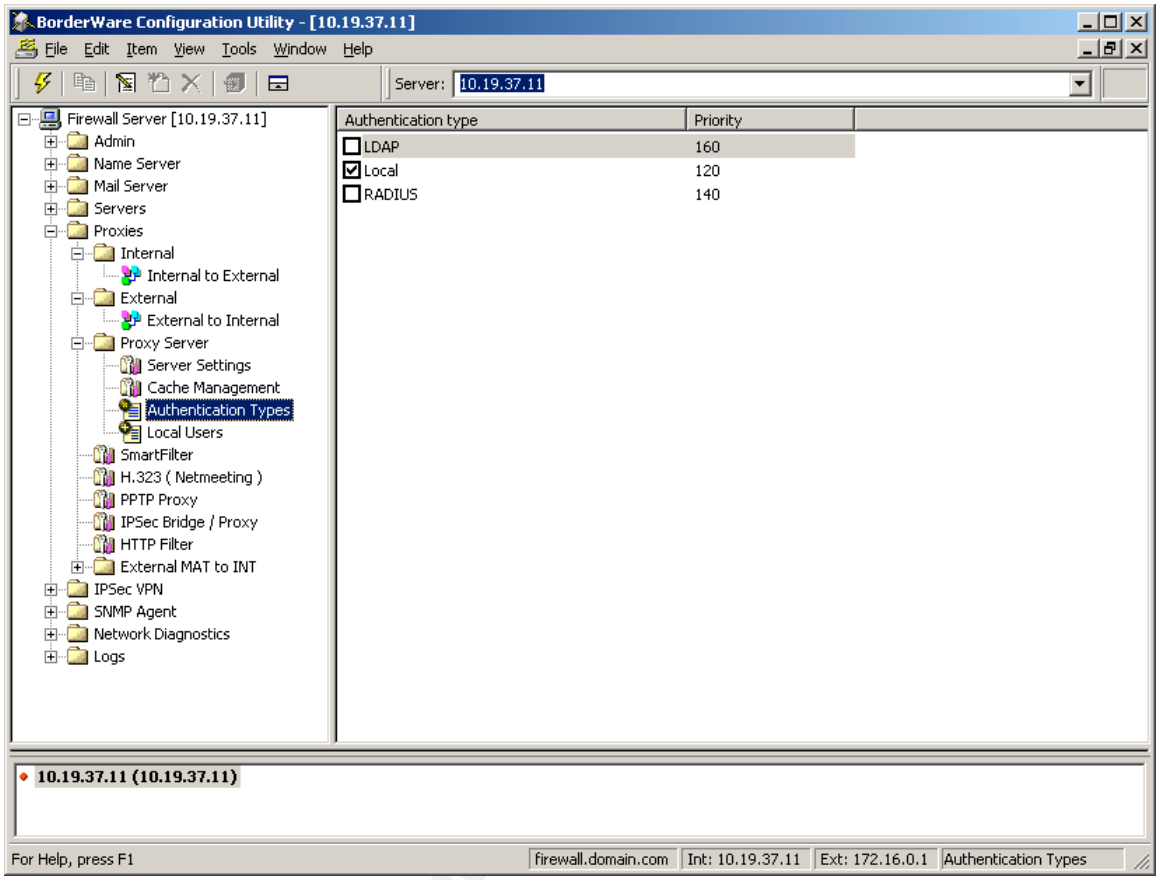

Each authentication type is set from the Authentication Types menu under the Proxy Server Folder. To modify each selection, double click on the entry and priorities can be set accordingly. It is also helpful to note any backup servers here as well in case the primary authentication server is not reachable.

Setting up Local Users is straightforward. In order to do this you must enable the local user authentication type as previously described and make sure its priority value is set. To add a new user, go to the Local Users Menu under the Proxy Server Folder and right click on the right hand side of the screen. Select Add New. You must provide a user name, full name and password for each user before clicking OK to continue. You will need to apply the changes for them to go into effect. You can modify a user in much the same fashion. Just right click on the user specified and chose Modify. Once a users password is set it cannot be recovered but can be changed if forgotten.

Setting up the other Authentication Servers is also straightforward. Again each must be enabled with a priority set before they can be configured. When adding a RADIUS server the following information must be provided.

- Host – the fully qualified domain name or the IP address is required.
- Priority – If multiple RADIUS servers are specified then this dictates the order they are queried in. This setting is unique to the RADIUS Servers and does not have anything to do with the priority set for the methods of authentication.
- Secret. This value must match that entered on the RADIUS server exactly
- Timeout. The amount of time the Firewall will wait for a response before moving on.
- Retries. The number of times each server will be retried before the next is asked.

Setting up an LDAP server is not much different. Only two entries are required when an LDAP server is entered as an authentication type thought.

- Host – the fully qualified domain name or the IP address is required.
- Priority - If multiple LDAP servers are specified then this dictates the order they are queried in. This setting is unique to the LDAP Servers and does not have anything to do with the priority set for the methods of authentication.

Several parameters are required for LDAP Authentication. These tell the Proxy Server the criteria that are to be used in order to locate the information it needs. These values are global and apply to each and every LDAP server listed. Depending on the LDAP Server you are employing, not all these entries may be needed.

- Base DN – this is the LDAP distinguished name (DN) that will be used to authenticate the user. Once the user has been found in the directory, the firewall will attempt to bind to the LDAP server as this DN.
- Search Filter – Search Filter is a string representation of the filter to apply in the search.
- Bind DN – this is the DN that will be used if the LDAP server does not allow anonymous bind requests.
- Bind PW – The password used for the above Bind DN request
- Size Limit – the maximum number of entries to be returned.
- Time Limit – the maximum amount of time allowed for a search.
- Scope – this value specifies how many levels deep in the directory structure the firewall should follow the search
- Deref – this value defines how the LDAP search behaves upon receiving a reference to an alias.
- Referrals – this value defines how the LDAP search behaves upon receiving referrals to another server or set of servers.

The Network Diagnostic Folder

The Network Diagnostics Screen is the location where you can troubleshoot various aspects of the Firewall Server. Several different screens are available for you to choose from. The choices are Interfaces, System, Connectivity, Name Server, Mail Server and Advanced. With each menu, you have the choice to not only display the results but to have then emailed to an address that you specify. This is useful if you need to keep documented results of the tests you are performing.

The Interfaces Menu allows you to test the interfaces to ensure that they are up and active and to reset them if you suspect a problem. Resetting the interfaces will break all connections occurring at the time so thought should be given to the effects this will have on your network prior to choosing this. Simply clicking an option and clicking on the Run Test button at the top of the screen will perform the requested Action.

Figure 20: The Interfaces Diagnostic Menu

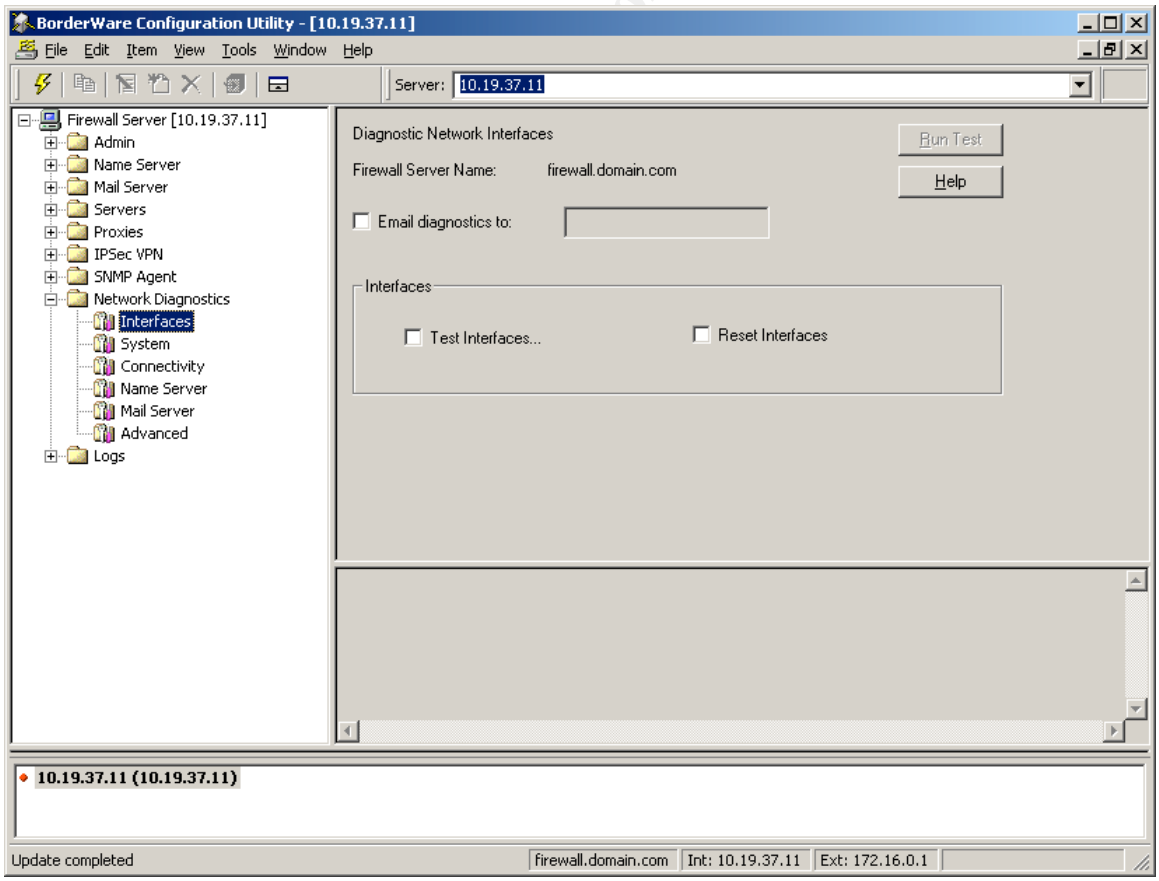

The System Menu allows you to see the static routes and some statistics on the routes as well as Disk usage and partitions if applicable.

The Connectivity Menu is helpful for troubleshooting connectivity within and without the network. You can ping each interface (like pinging the loopback address on a single network card computer), the external router, and a specific host on any interface. The specific host can either be the fully qualified domain name or the IP address of the host itself. You also have the ability to trace the route between the firewall and any host specified.

The Name Server Diagnostic allows you to query your own DNS servers on the firewall for required information. This is useful in troubleshooting DNS problems. You have the choices of Looking up a host by its name, address or performing a "Dig" on another DNS server using either your internal or external DNS server itself.

Figure 21: The Name Server Diagnostic Menu

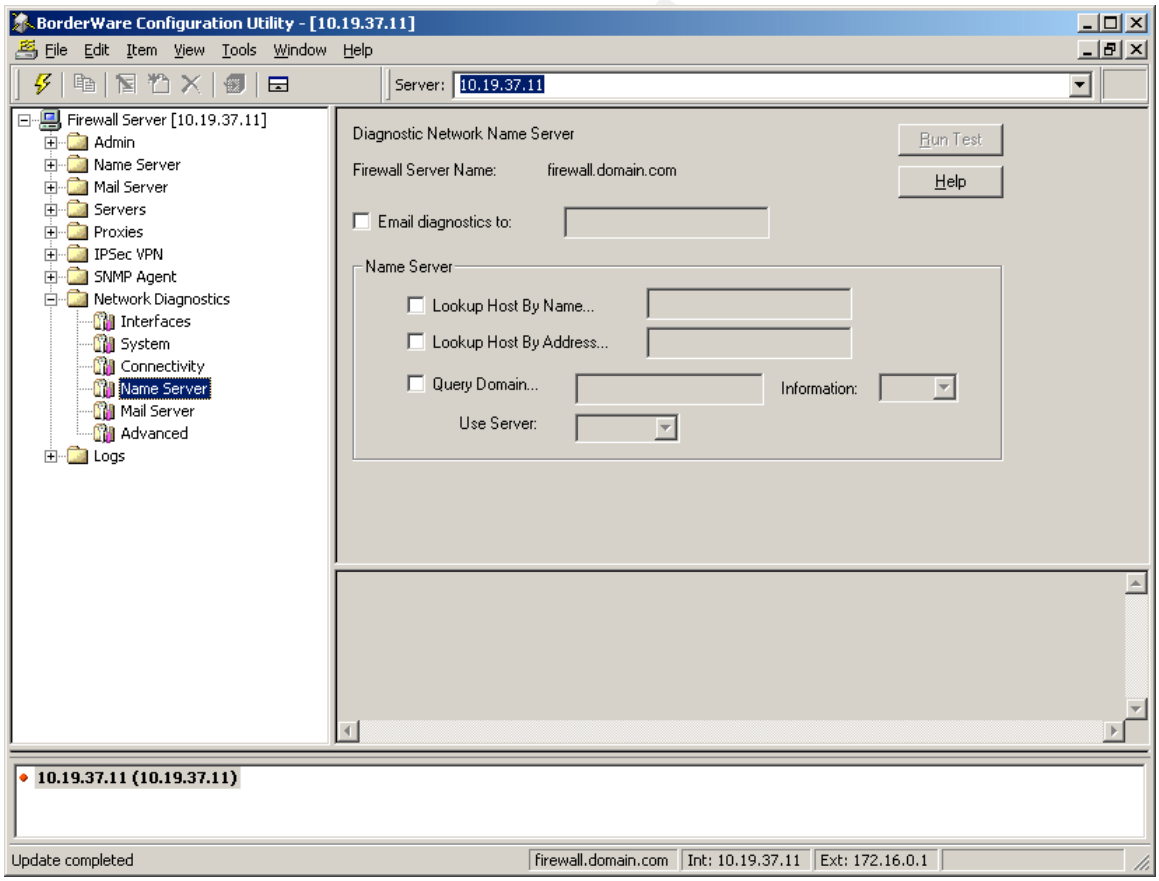

The Mail Server Diagnostic Menu allows you to send a test email and view the mail that is currently in the queue.

Advanced diagnostics provide detailed information about the state of the firewall. Several options are available for the administrator.

- Ifconfig – shows the configuration of each selected interface.
- Netstat – shows statistics according to the following available options
    - "rn" – show the network routing table
    - "an" – current open connections and ports being listened on
    - "m" – memory usage
    - "ind" – link status
    - "s" – statistics
- ARP – the address resolution protocol shows the table of IP address to MAC address translation.
- PSTAT – shows the files in use and swap space being used by the system
- FSTAT – shows all files being considered open by a process.
- VMSTAT – shows kernel statistics on items such as virtual memory, process and disk trap and CPU activity. The "m" option shows a snapshot of the current status where as the "512" option provides a continuous display of activity.

The Logs folder

The Logs folder allows you to access the FTP area where the Logs are kept via a web browser. The Admin username and FTP password are required to access this area. These logs can be downloaded via a FTP Client to save on another system or for archival purposes. The logs are located in /log directory.

Chapter 5

Auditing your Firewall

Once you have the Borderware Firewall Server installed and configured, the job is not over. Even the most careful administrator will forget something or misunderstand what they have configured and this could have devastating consequences for a system that is relied upon to provide security for your network. It is important to test what has just been configured to ensure the Firewall is working as expected. Lance Spitzner has published a good white paper on this process and it is the basis for this short section. Mr. Spitzner's information is generic in nature but extremely helpful in identifying those areas which need to be given attention.

Mr. Spitzner outlays a three step plan for performing an audit on any firewall (we are going to focus on the Borderware Firewall Server but the steps here are applicable to any firewall.) The three steps include one defining what to expect, two auditing the firewall itself, and three auditing the rule set. Mr. Spitzner goes into detail over the types of tools available to perform the audit however they will not be covered here. There are plenty of tools available from both the commercial and open source community to perform these tasks. Mr. Spitzner's white paper was written a few years ago and as technology changes a description of the tools he recommends may not be up to date.

The first step in performing an audit is pretty straightforward. Determine what you expect to see. This provides a measuring stick or guide by which you can gauge the performance if the firewall. Most of these should be currently defined in an existing security policy. A careful review of these policies will provide the needed guidance by which you can proceed.

The second step is auditing the firewall itself. Three items are covered in this step. The physical security of the device itself, the Operating System of the device and the vulnerability of the device. Is the physical device physically secured from attack? If an attacker can get physical access to the device then the firewall is no good or as Mr. Spitzner puts it: "game over."

Is the operating system secure? With the Borderware Firewall Server the operating system is hardened and inaccessible. Are all the patches current? This check should be done as part of a regular maintenance cycle anyway.

Now a scan of the firewall itself is needed. Make sure to scan the firewall from every interface. The most basic setup for a firewall is with three interfaces: the external, internal, and Secure Server Network (or DMZ.) Scan the firewall from each interface and look for open ports and running services. Review each to determine if these are needed and if not then shut them down. If the ports are

open and services need to be running then set up access rules to limit their access to only those who need to access it. This can be done by IP Address.

Once this information is obtained and corrections made, a test of the rule set needs to be preformed. It is at this stage that an audit of the firewall can turn into an audit of the network itself. Care needs to be taken to keep this audit on scope and any problems not associated with the firewall noted and addressed on another project.

In order to test the rule set of the firewall you will need to "scan every network segment from every other network segment." You will need to determine what traffic actually passes thru the firewall and why. Once suggestion Mr. Spitzner has is to scan one system on another network segment and compare that scan to one done on the same network segment. This comparison should yield results in what types of things can be seen from one segment to another. Use of a packet sniffer can be very useful during this stage to collect data and analyze results.

A test of authentication is also important. Test and verify that the systems you wish to require authentication for actually do require authentication and cannot be circumvented. If encryption is required, examine the packets related to this and verify that the data is indeed being transmitted in an encrypted form. Make note of any possible vulnerabilities and follow thru on those. This is a key element in Defense in Depth. Remember the firewall is a framework to Defense in Depth and not the intended to be the final layer of defense.

The last thing that needs to be done is verify the logging capability of the firewall. In all actuality this should be examined throughout the other steps. Verify that the firewall is logging the items you wish it to log and notating the tests you're performing. This is useful to provide a "fingerprint" of the types of information gathering that hackers use.

As was stated earlier, there are various tools to perform these functions. These tools are provided by commercial as well as open source sources and each has their own pros and con's.

Chapter 6

Summary

The Borderware Firewall Server offers a large degree of configurability to fit almost any environment from a small office setting to a large corporate enterprise. Proper configuration can result in a strong layer in the Defense in Depth concept of layered network security. In most applications, the Borderware Firewall Server will serve as a framework upon which the rest of the corporate network security strategy can be based.

The firewall server provides a fully functional and configurable name server with the ability to host multiple domains.

The firewall server supplies a mail server with the ability to deliver SMTP mail traffic and host POP mail accounts directly from the server. This server can prevent the relaying of mail from its external interface and can require FQDN's for mail being accepted to help cut down on the amount of SPAM received.

The firewall server has a fully configurable server and proxy interface complete with packet filtering up to layer 7 for pre-defined proxies and layer 3 for user-defined proxies. Setting up access rules is straightforward and intuitive.

The firewall server comes with a full diagnostic suite to provide valuable feedback in the event of connectivity and configuration problems. Detailed Logs are kept for review and analysis in an easy to locate area and can be imported into several popular report generating utilities currently on the market.

Bibliography

Snyder, Joel, “Turning the Network Inside Out,” Information Security, June 2003 URL: http://www.infosecuritymag.com/2003/jun/cover.shtml

Romanofski, Ernest, “A Comparison of Packet Filtering Vs Application Level Firewall Technology,” URL: http://www.group1ifw.com/whitepapers/a_comparison.htm

Unknown, “White Paper: Internet Security for Small Businesses - Firewall Protection for the Small Business Network,” URL: http://www.cisco.com/en/US/products/hw/routers/ps380/products_white_paper09186a008008872f.shtml

Cole, Eric, et al, SANS Security Essentials with CISSP CBK Version 2.1, 2003

McClure, Stuart, Scambray, Joel, Kurtz, George, Hacking Exposed: Network Security Secrets & Solutions Fourth Edition, 2003

Eagle, Liam, “Enabling the Defense in Depth Security Strategy,” URL: http://thewhir.com/features/depth-security.cfm

Paul, Brooke, “Building an In-Depth Defense,” July 9, 2001, URL: http://www.networkcomputing.com/1214/1214ws1.html

Spitzner, Lance, “Auditing your Firewall Setup,” 12 December 2000, URL: http://www.spitzner.net/audit.html

Borderware Help Files contained within the BWClient v1.9