

StarShield Handbook

StarShield Version 1.2

Table of Contents

1. Introduction

- 1.1. What StarShield is
- 1.2. What StarShield is not
- 1.3. History
- 1.4. Features
- 1.5. Software Copyright and Distribution (Licenses)
- 1.6. Contributors and Credits

4. Configuration

- 4.1. The Console Menu
- 4.2. The Web GUI
- 4.3. The System Screens
- 4.4. The Interfaces Screens
- 4.5. The Services Screens
- 4.6. The VPN Screens
- 4.7. The Status Screens
- 4.8. The Diagnostics Screens

5. The Firewall Screens

- 5.1. Rules
- 5.2. Inbound NAT
- 5.3. Server NAT
- 5.4. 1:1 NAT
- 5.5. Outbound NAT
- 5.6. Traffic Shaper
- 5.7. Aliases

6. Network Address Translation

- 6.1. NAT Primer
- 6.2. Inbound NAT
- 6.3. Server NAT
- 6.4. 1:1 NAT
- 6.5. Outbound NAT
- 6.6. Choosing the appropriate NAT for your network

7. Traffic Shaping

8. IPsec

- 8.1. Preface
- 8.2. Prerequisites
- 8.3. Configuring the VPN Tunnel
- 8.4. What if your StarShield isn't the main Internet Firewall?

9. PPTP

- 9.1. Preface

- 9.2. Audience
- 9.3. Assumptions
- 9.4. Subnetting and VLAN routing
- 9.5. Setup of StarShield software
- 9.6. PPTP User Setup
- 9.7. PPTP Firewall Rules
- 9.8. Setting up a PPTP Client on Windows XP™
- 9.9. Some things I have found not to work over the PPTP Connection

- 10. OpenVPN
- 11. Wireless
- 12. Captive Portal
- 13. Reference
 - 13.1. IP Basics
 - 13.2. IP Filtering
 - 13.3. NAT
 - 13.4. Traffic Shaping
 - 13.5. DNS
 - 13.6. Encryption (PPTP/IPsec)
 - 13.7. Logging (syslog)

- 14. Example Configurations
 - 14.1. Configuring a DMZ Interface Using NAT
 - 14.2. Locking Down DMZ Outbound Internet Access
 - 14.3. Configuring a filtered bridge

- 15. Example Site to Site VPN Configurations
 - 15.1. Cisco PIX Firewall
 - 15.2. Smoothwall
 - 15.3. FreeS/WAN
 - 15.4. Sonicwall
 - 15.5. Nortel

Chapter 1. Introduction

Table of Contents

- 1.1. What StarShield is
- 1.2. What StarShield is not
- 1.3. History
- 1.4. Features
 - 1.4.1. Components
 - 1.4.2. Specifications
- 1.5. Software Copyright and Distribution (Licenses)
 - 1.5.1. Other Software Packages
- 1.6. Contributors and Credits
 - 1.6.1. Code
 - 1.6.2. Documentation

1.1. What StarShield is

StarShield is a complete embedded firewall box based on a bare-bones version of FreeBSD, along with a web server (thttpd), PHP and a few other utilities. The entire system configuration is stored in one single XML text file to keep things transparent.

1.2. What StarShield is not

StarShield is a firewall, and the purpose of a firewall is to provide security. The more functionality is added, the greater the chance that a vulnerability in that additional functionality will compromise the security of the firewall.

We feel these services should be run on another server, and are intentionally not part of StarShield:

Intrusion Detection/Prevention System

Proxy Server

Packet inspection at any layers other than 3 and 4

A general purpose web server

An FTP server

A log file analyzer

StarShield provides many features, and some you won't find in any other firewalls, including:

web interface (supports SSL)

serial console interface for recovery

- o set LAN IP address
- o reset password
- o restore factory defaults

- o reboot system

wireless support (With optional wireless module)

stateful packet filtering

- o block/pass rules
- o logging

NAT/PAT (including 1:1)

DHCP client, PPPoE and PPTP support on the WAN interface

IPsec VPN tunnels (IKE; with support for hardware crypto cards and mobile clients)

PPTP VPN (with RADIUS server support)

static routes

DHCP server

caching DNS forwarder

DynDNS client

SNMP agent

traffic shaper

configuration backup/restore

host/network aliases

1.4.1. Components

StarShield contains the following software components:

FreeBSD components (kernel, user programs)

PF (the OpenBSD Packet Filter)

PHP (CGI version)

thttpd

MPD

ISC DHCP server

ez-ipupdate (for DynDNS updates)

Dnsmasq (for the caching DNS forwarder)

racoon (for IPsec IKE)

Chapter 4. Configuration

Table of Contents

4.1. The Console Menu

4.2. The Web GUI

4.3. The System Screens

4.3.1. General Setup

4.3.2. Static Routes

4.3.3. Firmware

4.3.4. Advanced

4.4. The Interfaces Screens

4.4.1. Assign Interfaces

4.4.2. LAN

4.4.3. WAN

4.4.4. Optional Interfaces

4.4.5. Wireless Interfaces

4.5. The Services Screens

4.5.1. DNS Forwarder

4.5.2. Dynamic DNS

4.5.3. DHCP

4.5.4. SNMP

4.5.5. Proxy ARP

4.5.6. Captive Portal

4.5.7. Wake on LAN

4.6. The VPN Screens

4.6.1. IPsec

4.6.2. PPTP

4.6.3. PPTP Users

4.7. The Status Screens

4.7.1. System

4.7.2. Interfaces

4.7.3. Traffic Graph

4.7.4. Wireless

4.8. The Diagnostics Screens

4.8.1. System Logs

4.8.2. DHCP Leases

4.8.3. IPsec

4.8.4. Ping

4.8.5. Reset State

4.8.6. Backup/Restore

4.8.7. Factory Defaults

4.8.8. Reboot System

The first time a StarShield system boots it uses a default configuration in which its IP address is set to 192.168.1.1, and it acts a DHCP server for the 192.168.1.X network. In many cases this default is sufficient to allow you to plug your LAN into StarShield's LAN port and then use a web browser on a LAN machine to connect to 192.168.1.1:80 (the web server on the StarShield box), after which you can do the remaining configuration using the webGUI interface as described below. Usually, however, you will have to use the console menu the first time StarShield boots in order to set up its network ports, after which you can use the webGUI for the remainder of the configuration. The network ports can also be assigned from the webGUI, so the console menu is only necessary to get you to the point where you can access the webGUI.

4.2. The Web GUI

To edit your StarShield configuration, point your web browser at your StarShield box. StarShield runs a web server on the standard web port (80) of its LAN connection. When you first connect to your StarShield web server, it will ask you for a user name and password. The username is admin and the default password is starsystem. To improve security, change the password in the General Setup screen.

The default StarShield configuration may be sufficient for you. If not, look through each of the screens, described below, to find the specific items you want to change. After you have made and saved your changes on the StarShield box, remember to download a backup copy of your configuration to another machine on your LAN.

When you first access the StarShield webGUI you will see the System Status screen. Along the left hand side of all screens is a menu to allow you to navigate to other screens. The items under the Interfaces menu heading may be different in your system, depending on how many network interfaces you have and how you have named them. The descriptions in the following sections are organized in the same way as the items in the navigation menu.

Note

Some of the screen shots in the following sections include blurred areas. When you view your StarShield screens, these will contain information specific to your system.

4.3. The System Screens

4.3.1. General Setup

The General Setup screen allows you to control some general parameters of your firewall.

Figure 4.1. The General Setup screen
The General Setup screen

The General Setup screen allows you to change the following parameters:

Table 4.1. General Setup parameters

Parameter	Description	Example	Reference
Hostname	The unqualified hostname of your firewall.	myfirewall	IP Basics
Domain	The domain name to qualify your firewall hostname.	mydomain.com	IP Basics
DNS Servers	The IP address of one or more DNS servers for use by the firewall.	10.0.0.123	DNS
Username	The username to use when connecting to the StarShield webGUI.	admin	
Password	The password to use when connecting to the StarShield webGUI. The current password is not displayed; this field is used only to change the password. You should change this when you first install StarShield.		
webGUI Protocol	The protocol for the StarShield webGUI to use. If you select HTTPS, you will need to access your webGUI using a URL that starts with "https:".		
webGUI Port	The port for the StarShield webGUI to use, if not the default.		
Time zone	The time zone of your firewall. This affects the value of times printed to logs.		
Logging			

Time update interval Logging	How often your firewall should contact the NTP server to update its time.
NTP time server Logging	The name of the NTP (Network Time Protocol) server for your firewall to use.

4.3.2. Static Routes

The Static Routes sub section allow the user to set up static routes in order to reach network that must use a gateway different from the default one. By pressing the + icon, the system allows the user to add new static routes.

The parameters to set up a new route are the following:

Interface: select the interface to which the route must be applied

Destination Network: select the network that have to be reached with Classless Inter-Domain Routing (CIDR) code for subnetting (see RFC1517, RFC1518, RFC1519, RFC1520 for more details)

Gateway: the gateway that the firewall must use in order to reach the Destination Network

Description: enter an optional description for the inserted route

4.4. The Interfaces Screens

4.4.2. LAN

In the LAN section, it is possible to change the IP address and the netmask (in CIDR notation) of the firewall internal interface. The system must be rebooted in order to apply the changes as suggested after pressing the "Save" button.

4.4.3. WAN

In the WAN sub section, it is possible to set up all the parameters for WAN interface. The WAN Interface can be a Static IP address, a DHCP address, a PPPoE interface or a PPTP connection, as detailed in the following. On the basis of the connection type selected, the related sub panel must be filled.

A detailed description of all the fields follows.

Note

You do not need to disable this option if you are using IPsec VPN tunnels with private IP addresses. When the VPN packets come into the WAN interface, they will be coming from source IP of the WAN interface of the remote VPN device, not from the private IP subnet on the remote side.

*

Type: the connection type that must be used

o

Static: A static IP address is assigned to the interface with the related netmask and gateway

o

DHCP: a dynamic address is assigned to the firewall WAN by a DHCP server on the WAN side

o

PPPoE: PPP over Ethernet, that is useful for ADSL connection

PPTP: allows to set up PPTP for the ADSL providers that requires this protocol for the connection

*

General Configuration Panel: allow to override default MAC address and MTU

MAC Address: some cable connections require the MAC spoofing. The MAC address must be in the format xx:xx:xx:xx:xx:xx

MTU: the value in this field allows to set up MSS clamping for TCP connections to the value entered above minus 40 (TCP/IP header size). If the field is left blank, an MTU of 1492 bytes for PPPoE and 1500 bytes for all other connection types will be assumed

*

Static IP Configuration: in this panel the static IP and gateway for WAN interface must be set:

IP Address: the static IP with related netmask is set in this field

Gateway: the default gateway for the firewall is set in this field

*

PPPoE Configuration: The Username and password for the ADSL connection should be set up there

Username: the username the provider assign to your connection

Password: the password the provider assign to your connection

*

PPTP Configuration: the parameters inserted in this sub panel allows the user to establish the tunnel required by the PPTP ADSL connection

Username: the username the provider assign to your connection

Password: the password the provider assign to your connection

Local IP Address: the local IP address the provider assign to your connection

Remote IP Address: the remote IP address the provider assign to your connection

* Block Private Networks - This option puts in rules to drop traffic coming in on the WAN from private IP subnets. If you configure your StarShield with the WAN interface on a private subnet of another LAN, for example, you need to disable this option. Also, some ISP's assign customers private IP's, in which case you'll also need to disable this option.

4.4.4. Optional Interfaces

Optional interfaces can be used for a variety of purposes. Generally they are used as second LAN interfaces or DMZ interfaces.

4.4.5. Wireless Interfaces

4.5. The Services Screens

4.5.1. DNS Forwarder

The DNS forwarder screen contains configuration options relevant to the DNS forwarding server on your StarShield.

Enabling the DNS Forwarder Check the first checkbox, "Enable DNS forwarder", to enable the service on the LAN interface. After enabling this, you will need to configure your client machines to use the LAN IP address of your StarShield as their DNS server.

Note

If the DNS forwarder is enabled, the DHCP service (if enabled) will automatically serve the LAN IP address as a DNS server to DHCP clients so they will use the forwarder. The DNS forwarder will use the DNS servers entered in System: General setup or those obtained via DHCP or PPP on WAN if the "Allow DNS server list to be overridden by DHCP/PPP on WAN" is checked. If you don't use that option (or if you use a static IP address on WAN), you must manually specify at least one DNS server on the System: General setup page.

DNS Host Name Registration

If your StarShield acts as the DHCP server for your LAN, and you need name resolution between hosts on the LAN, check the "Register DHCP leases in DNS forwarder" box. It will append the default domain in System:General setup. For example, if your machine name is my-pc and your default domain is example.com, it will register my-pc.example.com with the IP address assigned from DHCP, so the other hosts on your LAN can locate your machine by that name.

DNS Forwarder Overrides

If there are certain DNS host names you want to override for your internal DNS clients, add them under DNS overrides on this page. For example, if you want www.yourcompany.com to point to a different site internally than it does from the internet, enter an override for www.yourcompany.com with the appropriate IP address. This can also be used as a rudimentary (and easy to bypass) filter on web sites LAN clients can visit, by assigning the undesired host name to an invalid IP address. For example, to block www.example.com, put in an override to redirect it to an invalid IP address, such as 1.2.3.4. Note that using a different DNS server or editing the hosts file on the client machine gets around this restriction, but doing this is sufficient to block the site for the vast majority of users.

4.5.2. Dynamic DNS

Dynamic DNS allows you to have a permanent host name that can be used to access your network, generally used when your public IP address is assigned by DHCP and subject to change. This allows you to run your own web server, mail server, etc. using a DNS host name.

For links to providers of dynamic DNS services, visit the website of the dynamic DNS client used by StarShield, ez-ipupdate.

After you have signed up with one of the dynamic DNS providers listed, you can continue.

Configuring the Dynamic DNS Client

To start, first check the "Enable Dynamic DNS client" box at the top of the page.

In the "Service type" drop down box, select the service you signed up with above.

Some services support MX DNS records on dynamic DNS subdomains. This helps ensure you can get email to your host name. If your service supports this (dyndns.org is one that does, others do as well), fill in your mail server's host name in that field. If you do not need an MX record or if your provider does not support them, just leave the field blank.

Wildcards - If you want to enable wildcard on your dynamic DNS host name, check this box. This means all host names not specifically configured are redirected to your dynamic DNS name. So if your dynamic DNS is example.homeip.net, and you enable wildcards, www.example.homeip.net, mail.example.homeip.net, anything.example.homeip.net, etc. (i.e. *.example.homeip.net) will all resolve to example.homeip.net.

The next two boxes are for your username and password. Enter your account information from the dynamic DNS provider.

Click Save. Your dynamic DNS host name should immediately be updated with your WAN IP address. To verify this, ping your dynamic DNS host name. It should resolve to the IP address of the WAN interface of your StarShield. If not, check Diagnostics: System logs for information on why it failed.

4.5.3. DHCP

This screen allows you to enable the DHCP server on enabled Ethernet interfaces other than WAN.

Enabling the DHCP Server

To enable the DHCP server on a particular interface, click on the appropriate tab for the interface and check the "Enable DHCP server on interface" box.

Deny unknown clients

This option allows you to implement a more secure DHCP configuration. Many companies suffer from worm outbreaks and related security issues due to unauthorized machines being plugged into their network. This option will help ensure only authorized hosts can receive a lease from your DHCP server. With this option enabled, only hosts defined at the bottom of this page will receive a lease from DHCP.

The downside to this option is that it can be difficult to maintain when you have more than a handful of hosts on your network. Many will find the increased security worth the increase in maintenance. Note that this is only sufficient to stop the typical user that expects to be able to plug into your network and obtain a DHCP lease to get on the internet. Anyone with network and/or security expertise can easily bypass this.

Subnet, Subnet Mask, and Available range are filled in from the IP and subnet information from that particular interface.

Range

In the first box, enter the starting address of your DHCP range. In the second box, enter the ending address of the range. Note that you don't want to make this the same as the available range, as this includes the subnet address and broadcast address, which are unusable, as well as the address of your StarShield interface which also cannot be in the range.

WINS Servers

If you use an NT 4 domain, or have pre-Windows 2000 clients that need to access an Active Directory domain, you will need to fill in your WINS server IP addresses in these boxes. If you only have one WINS server, leave the second box blank.

Default and Maximum Lease Time

The default lease time is the length of the DHCP lease on any clients that do not request a specific expiration time on their DHCP lease. The default is 7200 seconds, or two hours. For the vast majority of network environments, this is too low. I would generally recommend setting this to a week, which is 604,800 seconds.

The maximum lease time must be more than the default lease time. Most networks will not use this value at all. In most instances, I set this to one second longer than the default lease time.

Click Save to save your changes, then click Apply to enable the DHCP server.

Static DHCP Mappings

Static DHCP mappings can be used to assign the same IP address every time to a particular host. This can be helpful if you define access rules on the firewall or on other hosts on your LAN based on IP address, but still want to use DHCP. Alternatively, you can keep the IP address box blank to assign an IP out of the available range, when you are using the "Deny unknown clients" option.

Click the + icon at the bottom of the DHCP configuration page to add a static DHCP mapping.

In the MAC address box, fill in the system's MAC address in the format xx:xx:xx:xx:xx:xx. For Windows NT/2000/XP clients, you can get determine the MAC address by opening up a command prompt and typing 'ipconfig'. For Windows 95/98/ME clients, go to Start, Run, winipcfg. For Unix clients, use ifconfig.

In the IP address box, fill in the IP address you want to be assigned to the client, or leave it blank to automatically assign one from the available DHCP range. If you put in a static IP address, it must not be within the range of the DHCP server.

It is recommended you fill in a description in the Description box to remind you what this entry is for, though this is an optional value.

Click Save when you are finished and the mapping will be added.

Note

The DNS servers entered in System: General setup (or the DNS forwarder, if enabled) will be assigned to clients by the DHCP server.

The DHCP lease table can be viewed on the Diagnostics: DHCP leases page.

4.5.4. SNMP

You can enable SNMP on your LAN interface on this screen. This is useful if you have a network management or monitoring system that takes advantage of it.

The System location and System contact boxes can be left blank, but can assist you in determining which device you are monitoring if you have several monitored hosts.

The Community is generally set to public, but if you have any regard for security at all, you should set this to something difficult to guess, containing numbers and letters. This community name is still passed over the network in clear text, so it could be intercepted, though the most anyone could get with that community name is information on the setup and utilization of your firewall. In most environments, this is likely of little to no concern, but is something to keep in mind.

After setting the values as you desire, click Save and your changes will be applied.

4.5.5. Proxy ARP

Proxy ARP can be used if you need StarShield to send ARP replies on the WAN interface for other IP addresses than its own WAN IP address (e.g. for 1:1, advanced outbound or server NAT). It is not necessary if you have a subnet routed to you or if you use PPPoE/PPTP, and it only works if the WAN interface is configured with a static IP address or DHCP.

If you enable 1:1, server, or advanced outbound NAT, you may need to enable proxy ARP for the IP address(es) being used by those translations. To do so, click the + on this page.

Enter either a single IP address, or subnet or range of addresses, optionally add a description to remind you why you made this entry, and click Save. Then click "Apply changes" for StarShield to enable proxy ARP.

For more information on when you do and do not need Proxy ARP, see this page.

4.5.6. Captive Portal

What is Captive Portal? from wikipedia.org

The captive portal technique forces a HTTP client on a network to see a special web page (usually for Authentication) before surfing the Internet normally. This is done by intercepting all HTTP traffic, regardless of address, until the user is allowed to exit the portal. You will see captive portals in use at most Wi-Fi hotspots. It can be used to control wired access (e.g. apartment houses, business centers, "open" Ethernet jacks) as well.

Check the "Enable captive portal" box to enable.

Interface - Select the interface on which you want to enable captive portal. It can only run on one interface at a time.

Idle timeout - Clients will be disconnected after this amount of inactivity. They may log in again immediately, though. Leave this field blank for no idle timeout.

Hard timeout - Clients will be disconnected after this amount of time, regardless of activity. They may log in again immediately, though. Leave this field blank for no hard timeout (not recommended unless an idle timeout is set).

Logout popup window - If enabled, a popup window will appear when clients are allowed through the captive portal. This allows clients to explicitly disconnect themselves before the idle or hard timeout occurs. When RADIUS accounting is enabled, this option is implied.

Note

Most any popup stopper will block this window. Worse, you cannot exclude a specific site, as this popup appears to come from whatever server the user tried to go to prior to authentication. If you have a popup blocker, you'll need to disable it prior to logging in, and then re-enable it after the logoff popup appears.

RADIUS server - Enter the IP address and port of the RADIUS server which users of the captive portal have to authenticate against. Leave blank to disable RADIUS authentication. Leave port number blank to use the default port (1812). Leave the RADIUS shared secret blank to not use a RADIUS shared secret. RADIUS accounting packets will also be sent to port 1813 of the RADIUS server if RADIUS accounting is enabled.

Portal page contents - Here you can upload an HTML file for the portal page (leave blank to keep the current one, or the default if you have not uploaded one previously).

Authentication error page contents - The contents of the HTML file that you upload here are displayed when a RADIUS authentication error occurs (generally because of an incorrect logon or password).

4.5.7. Wake on LAN

This service can be used to wake up (power on) computers by sending special "Magic Packets". The NIC in the computer that is to be woken up must support Wake on LAN and has to be configured properly (WOL cable, BIOS settings).

This might be useful, for instance, if you access your home or corporate network remotely via VPN, and need to access a machine that may not be powered on at all times. You can log into the StarShield device at that location and send a wake up packet.

To power on a machine, just choose the appropriate interface, put the MAC address of the machine into the MAC address box, and click "Send".

If you use this feature at all, you will probably want to create a list of the machines you want to remotely power on. If you click the + at the bottom of the screen, you can add a host to the list that is displayed. Once you have added the host to your list, you can simply click on the MAC address to power on the system.

4.6.1. IPsec

4.6.2. PPTP

4.6.3. PPTP Users

4.7. The Status Screens

4.7.1. System

Figure 4.3. The System Status screen

The System Status screen

4.7.2. Interfaces

4.7.3. Traffic Graph

Figure 4.4. The Traffic Graph screen

The Traffic Graph screen

The traffic screen allows you to select an interface, and view real time throughput graphs on that interface. This feature was introduced in version 1.1.

The Adobe SVG viewer is required to view the graphs. This page has a link to the installation for this viewer.

4.7.4. Wireless

4.8. The Diagnostics Screens

4.8.1. System Logs

4.8.2. DHCP Leases

This screen can be used to view your active and/or expired DHCP leases. Clicking the button on this screen will switch between showing only active leases and showing both active and expired leases.

Expired DHCP leases show up in gray text, while active ones are black. (this screenshot from a system with only expired leases)

4.8.3. IPsec

IPsec maintains two databases with connection details.

Security Association Database

First is the Security Association Database (SAD). This database maintains a list of all current IPsec Security Associations (SA's).

Security Policy Database

Second is the Security Policy Database (SPD). This database maintains a list of all the IPsec policies on the system. You will have two SPD entries for each IPsec VPN connection you have configured, regardless of whether the connection is up. This database tells the system what traffic will pass over VPN, and specifically which tunnel it traverses.

Table 4.2. The two entries for each VPN connection are as follows:

Source	Destination	Direction	Protocol	Tunnel Endpoints
local IP subnet for VPN connection	remote IP subnet for VPN connection		protocol in use (ESP or AH)	Public IP address of local StarShield - Public IP address of remote endpoint
remote IP subnet for VPN connection	local IP subnet for VPN connection		protocol in use (ESP or AH)	Public IP address of remote endpoint - Public IP address of local StarShield

At this screen, you will see two entries for each IPsec connection that has been successfully negotiated. One from the local public IP to the remote endpoint's public IP, and one in the opposite direction. This indicates that IPsec negotiations were successful, and that traffic should now be passing your VPN connection if everything else is configured appropriately.

By clicking on the X, you can delete the SA. StarShield will attempt to recreate it after deleting it. If you have a VPN connection with duplicate SA's (more than one from same src to same dst) and the connection has gone down, delete all the SA's associated with the connection. It should renegotiate and come back up within a few seconds.

4.8.4. Ping

This screen gives you a GUI to ping (send ICMP echo request) from the StarShield. Fill in the IP address or hostname of the machine to ping, choose the number of pings in the count drop down, and click the Ping button.

Note

The StarShield ping screen cannot ping over VPN connections for the same reason SNMP does not work over VPN out of the box. See this FAQ entry for more information. So do not use this screen as an indicator of whether your VPN is working.

4.8.5. Reset State

This screen allows you to reset the state tables on your StarShield for the NAT and firewall state tables.

Just check the boxes for the table(s) you want to clear, and click the Reset button.

Resetting the state tables will remove all entries from the corresponding tables. This means that all open connections will be broken and will have to be re-established. This may be necessary after making substantial changes to the firewall and/or NAT rules, especially if there are IP protocol mappings (e.g. for PPTP or IPv6) with open connections.

The firewall will normally leave the state tables intact when changing rules.

NOTE: If you reset the firewall state table, the browser session may appear to be hung after clicking "Reset". Simply refresh the page to continue.

4.8.6. Backup/Restore

StarSystem IT s.r.l.

This screen allows you to backup your existing configuration, or restore a previous backup file.

To backup your StarShield, click the "Download configuration" button. This will download a file called (by default) config.xml.

If you ever need to restore a previous backup file, go to this page, and under the "Restore configuration" section, click Browse. Locate the config.xml file you backed up above.

4.8.7. Factory Defaults

Clicking Yes on this page will reset StarShield to the default out of the box configuration options and clear any configuration you have done on the device.

If all else fails when trying to configure something on your StarShield, sometimes it is easiest to start over from scratch on the entire configuration. In that instance, use this feature to reload the default settings.

4.8.8. Reboot System

Click Yes on this page to reboot the system.

As a general rule of thumb in StarShield and FreeBSD in general, rebooting probably isn't going to fix any problems you are having. But it is worth a shot in many circumstances.

Unlike so many systems, rebooting isn't a suggested maintenance procedure on StarShield. There is no need to reboot the system unless you have a specific reason for doing so.

Chapter 5. The Firewall Screens

Table of Contents

5.1. Rules

5.2. Inbound NAT

5.2.1. Interface

5.2.2. External address

5.2.3. Protocol

5.2.4. External port range

5.2.5. NAT IP

5.2.6. Local port

5.2.7. Description

5.2.8. Auto-add a firewall rule to permit traffic through this NAT rule

5.2.9. Editing Inbound NAT Firewall Rule

5.3. Server NAT

5.3.1. Adding a Server NAT entry

5.3.2. Using the Server NAT entry

5.3.3. Enable Proxy ARP if necessary

5.4. 1:1 NAT

5.4.1. Adding a 1:1 NAT entry

5.5. Outbound NAT

5.6. Traffic Shaper

5.7. Aliases

5.7.1. Adding an Alias

5.7.2. Using Aliases

5.1. Rules

5.2. Inbound NAT

Inbound NAT allows you to open ports on your public IP address(es) to hosts in your LAN or OPT networks. Click Firewall -> NAT, and then on the Inbound NAT tab to add an entry.

5.2.1. Interface

Interface is generally WAN because we want to permit traffic coming in from the Internet. You can also select any optional interfaces here.

Optional interfaces might be useful on a DMZ interface to allow access from the DMZ to a port on a host on your LAN. For example, if you want to use a LAN DNS server, you could put an Inbound NAT rule in on the DMZ interface opening UDP port 53 to your DNS server's LAN IP address, and use StarShield's DMZ interface IP address as your DNS server on DMZ hosts. There isn't really any advantage over doing this versus putting in a firewall rule to permit this traffic and using the LAN IP address of the DNS server, rather than NAT'ing it.

5.2.2. External address

External address is set to the WAN interface's IP address. If you have multiple public IP's, you can use other addresses here that you have previously defined on the Server NAT tab.

5.2.3. Protocol

StarSystem IT s.r.l.

Sede Legale ed Operativa: Località Fratte, 49 – 38057 Pergine Valsugana (TN)

Cod. Fisc. e P.IVA 00698480225 - Tel. 0461.185.10.40 – Fax 0461.019.979

Internet: www.starsystem.biz - e-mail: info@starsystem.biz

Choose which IP protocol the service you are using requires, either TCP, UDP or TCP and UDP.

5.2.4. External port range

Either select the desired protocol from the drop down box, or type in the port range in the text boxes. You can leave the "to" field empty if you only want to map a single port.

Note

When you want to open more than one port to a system, for example HTTP and HTTPS, do not use a port range from HTTP to HTTPS. This will work, but it also opens up 361 ports that you don't need opened between TCP 80 and 443. If you need to open two non-sequential ports to a system, you need to put in two Inbound NAT entries.

5.2.5. NAT IP

This is the internal IP address of the machine to which you are mapping the ports. In the given example, the LAN IP address of the web server is 192.168.1.25. This can also be a host on an optional network, and ideally it will be to a host on a DMZ. You should avoid opening ports to your LAN if possible.

5.2.6. Local port

This is the port on the NAT IP defined above to which we want to translate the connection. In this case it is the same as the external port, but it doesn't have to be.

5.2.7. Description

Optional as always, but we strongly recommend putting in a description so you remember the purpose of this entry, and to make your rules easier to read and comprehend.

5.2.8. Auto-add a firewall rule to permit traffic through this NAT rule

I recommend you check this box in all circumstances. If you need to tighten the default rule, you can do so later. If you don't let the webGUI create the rule automatically, it's more likely to be incorrect or problematic.

Click Save, then click Apply changes. You'll see your result, similar to the following.

5.2.9. Editing Inbound NAT Firewall Rule

After adding an Inbound NAT entry and allowing the system to automatically create the firewall rule permitting traffic through that NAT entry, you can go to the Firewall -> Rules page to edit the rule. You might want to do this if, for example, you don't want to allow the entire Internet to access the service you have opened.

You'll see the rule under your WAN interface, similar to the following.

Click the next to the rule to edit it. You'll see something similar to the following.

To restrict access to this service, change the Source from any to either a network or single host and enter the appropriate details. After confirming your changes, click Save, and Apply changes.

5.3. Server NAT

If you want to use a public IP address other than the WAN interface address with Inbound NAT, you need to define the address in Server NAT first.

5.3.1. Adding a Server NAT entry

Click Firewall -> NAT, and click the Server NAT tab. Click the to add a new entry.

After double checking your entry, click Save and Apply changes.

The first time you add a Server NAT entry, you may have to reboot for the change to take effect. If you are prompted to reboot, you must do so before you can use the Server NAT entry.

5.3.2. Using the Server NAT entry

Now if you go to the Inbound NAT tab and click the to add a new entry, and click in the External address box, you will see the Server NAT entry you entered above.

5.3.3. Enable Proxy ARP if necessary

Depending on the way your WAN connection is setup, you may also need Proxy ARP for Server NAT to function.

If any of the following applies to your setup, you should be fine without proxy ARP:

*

the additional IP addresses that you're trying to use are part of a subnet that is routed to you by your ISP (i.e. your ISP has a static route for that subnet with your StarShield's WAN IP address as the gateway)

*

you're using PPPoE or PPTP on WAN

Using proxy ARP under these conditions will not achieve anything. If however you use static IP addresses or DHCP on WAN and don't have a routed subnet, adding proxy ARP entries for the additional addresses/ranges/subnets in the webGUI will make sure that StarShield responds to ARP queries for these addresses on the WAN interface.

5.4. 1:1 NAT

1:1 NAT maps an internal IP to external IP, generally mapping a public IP address to a private IP address and vice versa. When you assign a 1:1 NAT mapping, any traffic coming from that host to the Internet will be NAT'ed to the defined external IP, and any traffic coming into the external IP will be NAT'ed and passed to the internal IP if firewall rules permit. (by default, the firewall rules do not allow any inbound traffic to 1:1 NAT mappings)

You can also map entire subnets with one entry.

You can also use this on optional networks, but that is not a common use of this functionality.

5.4.1. Adding a 1:1 NAT entry

Go to the Firewall -> NAT screen and click the 1:1 tab. Click the to add a new entry.

5.4.1.1. Interface

Interface will be WAN in most all cases.

5.4.1.2. External subnet

The external subnet will be set to the IP address or subnet you wish to map. Usually this will be a single IP address (and hence a /32 mask). If you have, for example, a full class C public subnet and your LAN or DMZ is a full class C subnet and you want to 1:1 NAT everything to its own public IP, you need to enter your entire public IP subnet here.

5.4.1.3. Internal subnet

In most cases this will be a single IP address on either your LAN or an optional interface like a DMZ. Or in the case of 1:1 NAT'ing an entire subnet, enter the subnet address here. The mask given in the External subnet is used, as they must be identical.

5.4.1.4. Description

Description is optional but recommended.

After verifying your entries, click Save and Apply changes.

Note

Depending on the way your WAN connection is setup, you may need Proxy ARP for 1:1 NAT to function. See the Proxy ARP section under Server NAT for more information.

5.5. Outbound NAT

5.6. Traffic Shaper

5.7. Aliases

You may have noticed throughout the webGUI there are some address boxes with a blue background. This blue background indicates you can use aliases in this field. The source and destination boxes on the Firewall Rules Edit screen are two examples of this.

Aliases act as placeholders for real IP addresses and can be used to minimize the number of changes that have to be made if a host or network address changes. You can enter the name of an alias instead of an IP address in all address fields that have a blue background. The alias will be resolved to its current address according to the defined alias list. If an alias cannot be resolved (e.g. because you deleted it), the corresponding element (e.g. filter/NAT/shaper rule) will be considered invalid and skipped.

5.7.1. Adding an Alias

Go to the Firewall -> Alias screen and click the to add an alias.

5.7.1.1. Name

The name of the alias - you'll use this in the blue boxes throughout the system.

5.7.1.2. Type

Either a reference to a single host, or a network.

5.7.1.3. Address

This is the IP address or subnet that this alias represents.

5.7.1.4. Description

As always, optional, but recommended.

After verifying your entries, click Save, and Apply changes.

5.7.2. Using Aliases

Now that you have entered an alias, you can use it in any of the boxes with blue backgrounds by selecting type "Single host or alias" and typing in the alias name in the "Address" box.

Chapter 6. Network Address Translation

Table of Contents

6.1. NAT Primer

- 6.1.1. Types of NAT
- 6.1.2. Other Resources

6.2. Inbound NAT

6.3. Server NAT

6.4. 1:1 NAT

6.5. Outbound NAT

6.6. Choosing the appropriate NAT for your network

6.1. NAT Primer

Network Address Translation (NAT) allows you to use RFC 1918 private IP addresses for addressing on your internal network, and allow all hosts on the internal networks to access the Internet using one public IP address.

Due to the typical expense of obtaining public IP addresses, most networks do not purchase one public IP address for each network host. NAT allows multiple machines to connect to the Internet using a single public IP address.

6.1.1. Types of NAT

There are two most commonly used and most familiar types of NAT, bidirectional or 1:1 (pronounced one to one), and Port Address Translation, or PAT.

6.1.1.1. 1:1 NAT Explained

1:1 NAT maps one public IP address to one private IP address, for both incoming traffic and outgoing traffic.

6.1.1.2. PAT Explained

6.1.2. Other Resources

RFC 1918 - Address Allocation for Private Internets - February 1996

RFC 1631 - The IP Network Address Translator (NAT) - May 1994

Network Address Translation at Wikipedia

6.2. Inbound NAT

Inbound NAT allows you to open up TCP and/or UDP ports or port ranges to hosts on networks protected by StarShield. You may need to open ports to allow certain NAT-unfriendly applications and protocols to function properly. Also if you run any services or applications that require inbound connections to a machine on your internal network, you will need inbound NAT.

6.3. Server NAT

Server NAT just gives you the ability to define extra IP addresses, other than the WAN IP, to be available for use for Inbound NAT.

6.4. 1:1 NAT

6.5. Outbound NAT

By default, StarShield automatically adds NAT rules to all interfaces to NAT your internal hosts to your WAN IP address for outbound traffic. The only exception is for any hosts for which you have configured 1:1 NAT entries. Therefore, if you are using public IP addresses on any of the interfaces behind your StarShield (with

the exception of bridged interfaces) you need to change StarShield's default NAT behavior by enabling advanced outbound NAT.

If you are using public IP addresses on all the interfaces behind your StarShield, check the "Enable advanced outbound NAT" box and click Save. Now nothing will be NAT'ed by StarShield.

If you have a public IP subnet off one of your interfaces behind StarShield and a private IP subnet behind another interface, you will need to enter your own NAT mappings on this screen. For example, if you have a LAN subnet of 192.168.1.0/24 and a DMZ subnet with public IP addresses, you will need to enable advanced outbound NAT, and click the plus at the bottom of this tab to add a NAT mapping for your LAN network. For this scenario, you will want to add a rule for interface WAN, source 192.168.1.0/24, destination any, target box blank, and enter a description of your choosing.

6.6. Choosing the appropriate NAT for your network

So by now you may be thinking "so what kind of NAT do I need?", to which the answer is "it depends."

For networks with one public IP, the only option is Inbound NAT, since that public IP will be assigned to StarShield's WAN interface.

For networks with multiple public IP addresses, the best choice is either 1:1 NAT, or Server and Inbound NAT, or a combination of both. If you have more servers than public IP addresses, you will need to use Server and Inbound NAT, or 1:1 NAT combined with Server and Inbound NAT. If you have sufficient public IP addresses for all of your servers, you should use 1:1 NAT for them all.

Inbound and Server NAT is most suitable when you have more servers than public IP addresses. For example, if you have three servers, one HTTP, one SMTP, and one FTP, and have only two public IP addresses, you must use Server and Inbound NAT. For small deployments, this isn't bad to deal with. As the number of hosts increases, things get far more complicated. You'll end up having to remember things like for public IP address 1.2.3.4, port 80 goes to server A, port 25 goes to server B, port 21 goes to server C, etc. If you can't clearly picture a network in your head while troubleshooting problems, things become much more difficult. With ports going all over the place like this, once you get a number of ports forwarded it's extremely difficult to picture the network in your head. Given the complexity introduced by such a configuration, we recommend having one public IP address per publicly-accessible host.

Chapter 7. Traffic Shaping

Chapter 8. IPsec

Table of Contents

8.1. Preface

8.1.1. Site to Site VPN Explained

8.1.2. Remote Access IPsec VPN

8.2. Prerequisites

8.3. Configuring the VPN Tunnel

8.4. What if your StarShield isn't the main Internet Firewall?

8.1. Preface

IPsec (IP security) is a standard for providing security to IP protocols via encryption and/or authentication, typically employing both. Its use in StarShield is for Virtual Private Networks (VPN's).

There are two types of IPsec VPN capabilities in StarShield, site to site and remote access.

8.1.1. Site to Site VPN Explained

Site to site VPN's connect two locations with static public IP addresses and allow traffic to be routed between the two networks. This is most commonly used to connect an organization's branch offices back to its main office, so branch users can access network resources in the main office. Prior to VPN's, much more expensive private Wide Area Network (WAN) links like frame relay, point to point T1 lines, etc. were commonly used for this functionality. Some organizations are moving towards VPN links between sites to take advantage of reduced costs.

Site to site VPN's can also be used to link your home network to a friend's home network, to provide access to each other's network resources without opening holes in your firewalls.

While site to site VPN's are a good solution in many cases, private WAN links also have their benefits. IPsec adds processing overhead, and the Internet has far greater latency than a private network, so VPN connections are typically slower (while maybe not throughput-wise, they at least have much higher latency). A point to point T1 typically has latency of around 4-8 ms, while a typical VPN connection will be 30-80+ ms depending on the number of hops on the Internet between the two VPN endpoints.

When deploying VPN's, you should stay with the same ISP for all sites if possible, or at a minimum, stay with ISP's that use the same backbone provider. Geographic proximity usually has no relation to Internet proximity. A server in the same city as you but on a different Internet-backbone provider could be as far away from you in Internet distance (hops) as a server on the other side of the continent. This difference in Internet proximity can make the difference between a VPN with 30 ms latency and one with 80+ ms latency.

8.1.2. Remote Access IPsec VPN

StarShield provides two means of remote access VPN, PPTP and IPsec (with OpenVPN available in beta versions only for now). StarShield's mobile IPsec functionality has some serious limitations that hinder its practicality for many deployments. StarShield does not support NAT-Traversal (NAT-T) for IPsec, which means if any of your client machines are behind NAT, IPsec VPN will not work. This alone eliminates it as a possibility for most environments, since remote users will almost always need access from behind NAT. Many home networks use a NAT router of some sort, as do most hot spot locations, hotel networks, etc.

One good use of the StarShield IPsec client VPN capabilities is to secure all traffic sent by hosts on a wireless network or other untrusted network. This will be described later in this chapter.

FIXME - A second limitation is the lack of any really good, free IPsec VPN clients for Windows. Most of your remote users will likely be Windows laptop users, so this is another major hindrance.

For most situations, PPTP is probably the best remote access VPN option in StarShield right now. See the PPTP chapter for more information.

This chapter will go over configuring a site to site VPN link between two StarShields, and will discuss how to configure site to site links with third party IPsec-compliant devices. The Example VPN Configurations chapter goes over, in detail, how to configure site to site IPsec links with some third party IPsec devices. If you have gotten StarShield working in a site to site IPsec configuration with some third party IPsec device, we would appreciate if you could put together a short write up of how you got it configured, preferably with screenshots where applicable.

8.2. Prerequisites

Before getting started, you need to take care of the following.

1.

Your StarShield must be setup and working properly for your network environment.

2.

Both locations must be using non-overlapping LAN IP subnets.

i.e. if both sites are using 192.168.1.0/24 on the LAN, no site to site VPN will work. This is not a limitation in StarShield, it's basic IP routing. When any host on either of your networks tries to communicate with 192.168.1.0/24, it will consider that host to be on its local LAN and the packets will never reach StarShield to be passed over the VPN connection. Similarly, if one site is using, for example, 192.168.0.0/16 and one using 192.168.1.0/24, these subnets are also overlapping and a site to site VPN will not work.

Keep in mind the more networks you link together the more important this basic fact becomes. Do not use unnecessarily large subnet masks. If you setup your LAN as 10.0.0.0/8, but only have 100 hosts on it, you're unnecessarily limiting your ability to add VPN networks anywhere in the 10.x.x.x space.

3.

If StarShield is not the default gateway on the LAN where it is installed, you must add static routes to whatever system is the default gateway, pointing the remote VPN subnet to the LAN IP of StarShield.

4.

You will need to either control or be in contact with the person who does control the other VPN concentrator. If it is another StarShield system, then share this document with the other administrator. If it isn't then have them consult the documentation that came with the IPsec device they are using.

5.

Host and application level security become more important when connecting multiple networks, how much depending on how much you trust the other network. The VPN tunnel will not respond to firewall rules at the time of this writing, so you will not be able to limit which hosts can be accessed by users across the VPN connection. If a worm would get into the network you are connected to via VPN, it could easily spread to your network. If a system on the remote network is compromised by an attacker, he could easily hop over the VPN to attack your systems without any firewall protection.

6.

Pay attention to what you are doing! If you have a VPN to your office, and a VPN to your friend's home network, your friend can now hop over to your company's network from your network. Or, if your friend gets infected with a worm, it could then infect your machines and continue to propagate over the VPN connection to your office. Most companies would probably fire you if your friend was caught on their network. Best bet

here is if you have a site to site VPN into your network at work, do not connect with friends, or use one network and firewall for accessing work and one for accessing your friend's network.

Ok now that we have the basics let's get started on the firewall settings.

8.3. Configuring the VPN Tunnel

Log into your StarShield and click IPsec , under VPN.

Ok now we need to add a VPN connection, to do this click on the icon.

You will be presented with a great form, I will be pasting screen shots of each section as we discuss it.

The first area is the one you use to establish what network ranges will use this IPSEC tunnel.

This is the first set of fields that we need to concentrate on. Later, when testing your tunnel, you can actually fail to establish level 2 connection if this data is incorrect. I will note what to pay particular attention to as we go along.

1.

Mode, this is a hard set option and frankly you don't need to change it (nor can you.)

2.

Disabled, this is a great "on / off" button if you need to disable the tunnel for what ever reason. Simply select the edit or from the main VPN: IPsec window and click this checkbox element, then select apply at the bottom of the page. When you need the tunnel again, reverse the process.

3.

Interface, this is how you determine which part of your network will be the termination point (end point) for the VPN Tunnel. If you are connecting to a remote server, then WAN is your option.

4.

Local subnet. This is where you can set which parts, hosts, or the entire LAN can be accessed from the other side of the VPN tunnel. The easiest thing to do is to set the LAN subnet as the option; this means your entire LAN will be accessible from the remote network. IMPORTANT: The other end of the tunnel has this same field, well it probably has 99% of these fields actually, make sure the other end is set exactly as you set this end. E.g. if you said "Single host" in this section and entered the IP address of that host, the other person would set that host in his "Remote Subnet" field. The same goes for you, and with that mentioned we move to the next field.

5.

Remote Subnet. This is more than just labeling which hosts and / or host you want to access on the other network, as mentioned in item 4 it is paramount that you set this exactly like the other end's "local subnet" section. If not, level 2 of the VPN connection will fail and traffic will not pass from one VPN segment to the other.

6.

Description: It is a good practice to always leave notes about why you are doing something. I suggest you enter something about what this VPN tunnel is used for, or about the remote end of the tunnel to remind yourself who/what it is.

Ok all the basic for the routing have been established. Now we move on to phase 1 of the VPN authentication process.

Okay the easy part of the VPN tunnel. The trick here, and even in phase 2, is to make sure that both VPN servers have EXACTLY THE SAME SETTINGS for all of these fields. Well okay, they will have different "My identifier" but make darn sure that they know each others names... more on that later.

1.

Negotiation mode: This is the type of authentication security that will be used. Unless you are under close watch by someone with paranormal like craziness, just leave this as aggressive. It is indeed far faster and will insure that your VPN tunnel will rebuild itself quickly and probably won't time out an application if the tunnel was down when the resource on the other end was requested. (more about that under Lifetime)

2.

My identifier: This is the key to probably 90% of the email on the list where people seem to not get the VPN tunnel up, or want to know how to do this with dynamic IP addresses, etc. Very simple, set your identifier to something that isn't going to change. So if you leave it as My IP address (* This will be the IP address of the "interface" you listed in the first section. *) then make sure that IP is static and persistent. If you use a DHCP assigned address then I would suggest using domain name instead This is because domain name can be completely your own even if you do not own the domain name. Make yours sexylovemonkey.com just for fun. ;)

3.

Encryption Algorithm: 3DES is the world de facto... if you are connecting to another StarShield, or a system that will support it, change this to Blowfish. It is a more secure and about twice as fast! Now of course, if you are trying to connect to a VPN device that only supports DES then you will need to downgrade and hope no one decrypts your key exchange. MAKE SURE BOTH VPN DEVICES ARE USING THE SAME ENCRYPTION ALGORITHM.

4.

Hash Algorithm: this is the hash used for checksum. MD5 is a good choice, SHA1 is the new up and comer and it is more reliable then MD5, but not all things support it. Again make sure you are using the same setting as the other end of the tunnel, and if you can use SHA1 go for it!

5.

DH Key Group: Most systems will support at least up to 1024 bit. This is a good place to stick to, going with more will eat up more resources and less makes your tunnel less-secure.

6.

Lifetime: This field is far more important then it appears. This lifetime, as opposed to the one in phase 2, is how long your end will wait for phase 1 to be completed. I suggest using 28800 in this field.

7.

Pre-Shared Key: Contrary to some suggestions this key must be exactly the same on both VPN routers. It is case sensitive, and it does support special characters. I suggest using both. E.x. f00m0nk3y@BubbaLand

Okay if you managed to coordinate and get both VPN systems set the same all should be good for phase 1. We really don't want to stop here, so let's go right into phase 2.

Phase 2 is what builds the actual tunnel, sets the protocol to use, and sets the length of time to keep the tunnel up when there is no traffic on it.

1.

Protocol: ESP is the de facto on what most VPN systems use as a transport protocol. I suggest leaving this as is. Note: The system should auto generate a firewall rule for you to allow ESP or AH to the endpoint of the VPN. We will check this later, if it does not you will need to make a firewall rule allowing ESP (or AH if

you changed this) traffic to the interface you established as your end point of the tunnel. I will outline that after figure 5.

2.

Encryption algorithms: Ok here is the deal on this. Like before in phase 1, make sure you are setting the algorithm exactly as it is set on the other VPN server. You can use several; when you do so everything you select is available for use. Honestly I like to keep things simple so I recommend only checking the one you are going to use. With StarShield to StarShield use Blowfish for speed and security over 3DES.

3.

Hash algorithms: again just as in phase 1 you want to make sure your selected hash matches the one on the other end. And like in step 2, don't add things you don't need. SHA1 is the suggestion if you can, but MD5 is always a good alternative.

4.

PFS key group: this works exactly like it does in phase 1. I suggest using 1024 bit, the default is off.

5.

Lifetime: This is the lifetime the negotiated keys will be valid for. Do not set this to too high of a number. E.g. more than about a day (86400) as doing so will give people more time to crack your key. Don't be over paranoid either; there is no need to set this to 20 minutes or something like that. Honestly, one day is probably good.

6.

Click Save

7.

Click Apply Changes

8.4. What if your StarShield isn't the main Internet Firewall?

FIXME - In some cases you have a firewall or router with layer 2 routing (protocol ACLs) sitting in front of your StarShield. If this is the case you will need to port forward ESP or AH (depending on which one you chose) to the StarShield. (NOTE: if you are running NAT on that firewall AH will not be an option.)

Figure 8.1. Example: StarShield behind a router

Example: StarShield behind a router

Chapter 9. PPTP

Table of Contents

- 9.1. Preface
- 9.2. Audience
- 9.3. Assumptions
- 9.4. Subnetting and VLAN routing
- 9.5. Setup of StarShield software
- 9.6. PPTP User Setup
- 9.7. PPTP Firewall Rules

- 9.7.1. Example of filtered PPTP Rules

- 9.8. Setting up a PPTP Client on Windows XP™

- 9.8.1. Testing our PPTP Connection in Windows™

- 9.9. Some things I have found not to work over the PPTP Connection

This chapter is based on Francisco Artes' StarShield-PPTP document, used with permission.

9.1. Preface

This chapter is intended to outline several different PPTP VPN type setups, it includes a how-to on setting up a Windows XP™ PPTP client to connect to the StarShield PPTP VPN server. Later versions of this document will include Linux and other clients.

All Trade Marks™ are represented in this document, and no intention is made that this document, StarShield, or the author are in any way related to any of the companies holding these Trade Marks. All Trade Marks are copy written by their respective companies.

The terms firewall and StarShield are used synonymously in this chapter. This is mostly because it is easier to say and type "firewall".

9.2. Audience

You need to have a basic understanding of TCP/IP and subnetting to understand this document. The author does make every effort to describe the items being discussed, but let's face it I can only go so far. (And I did include pictures, which apparently are each worth 1,000 words. So that makes this one HUGE document.)

If you have comments, questions, or suggestions in regard to this document please email <falc@netassassin.com>. I will try to get back to you as quickly as possible, but please do read this document thoroughly before writing. You may also want to check the StarShield website for email archives on frequently (or even one-time) questions.

9.3. Assumptions

Ok we are going to make several assumptions in this document, if you don't have these assumptions done already you will need to go get them done before PPTP will work correctly.

- 1.

Your firewall is already setup to do basic NAT and you have tested this, or at least it is doing what ever kind of routing you wanted it to do.

- 2.

You have configured at least one interface on the firewall so it is working and:

- 1.

StarSystem IT s.r.l.

The Client Machine(s) can route to (access) one of the interfaces of your firewall. Make sure of this. If it is an interface that you allow ICMP to access I suggest pinging it.

3.

You have a client machine running some form of VPN client that supports PPTP.

Ok now that we have the basics let's get started on the firewall settings.

9.4. Subnetting and VLAN routing

Ok so this isn't quite true VLAN routing, but we will (quite possibly) be working with a virtual network that doesn't exist until a PPTP connection is made. If you have a better term for this let me know and I will change it. We are however dealing with some virtual subnets, for instance the "Remote Address Range" will be a /28 and PPTP clients will receive a subnet of 255.255.255.255 (ff.ff.ff.ff for all you HEX people out there.) Just ignore that and trust in the magic of the PPTP Tunnel.

You can select (as you will see later) to set the "Server Address" and "Remote Address Range" to exist inside of the subnet that you defined for the LAN on the firewall. (e.g. IP Address and subnet bit you set for the LAN under Interfaces LAN on the StarShield menu.) Our example uses this setup. Pros and Cons? Well the major pro is that the firewall will allow traffic from this VLAN to route to the WAN (in most cases the Internet.) and it is nice and easy. Con's, it allows people to rout to the WAN if you don't want this then read the next paragraph.

You can also setup these two options to have an IP range that is outside of your LAN designation. E.g. LAN = 192.168.1.1/24 (really the 192.168.1.0/24 network) and the PPTP "Server Address" and "Remote Address Range" are set to 192.168.2.254 and 192.168.2.16/28 respectively. This will basically allow those using the PPTP connection to access the LAN, but the firewall will not route traffic for them to the WAN connection. Opt and WiFi networks will also be isolated depending on how you are routing to those networks and if they are in the same network segment (subnet) as the LAN.

Remember, that when you setup a PPTP connection (especially on Windows) all network traffic from that workstation is going to be sent via the PPTP tunnel.

9.5. Setup of StarShield software

Most people probably skipped right to this point. If you did, it should be easy enough with these examples if you do run into something go read the parts you skipped you may find the answers there you are looking for.

1.

The first thing we want to do is setup the PPTP server. To do this select PPTP from the VPN section of the StarShield interface. If you clicked the right thing you will have a screen that looks something like Figure 1.

2.

The next step is to enable the PPTP server. Click the "Enable PPTP server" radio button. (It only gets harder from here.)

3.

Now we have to type. (see harder) So enter the "Server Address" next. This can be an unused IP on your LAN, or another locally usable IP address in a separate subnet. It MUST be in the same networking class as the next entry.

4.

Remote Address range. This is going to be the range of 16 IP addresses that the server will issue to clients. Notice the /28, it is there to remind you there will be 16 hosts. Again, this MUST be in the same subnet class as the IP listed above. (Not in the same /28 though.... If you try to overlap the two the firewall will tell you that you made a mistake.)

In our example we used 192.168.1.254 for the "Server Address" and 192.168.1.192/28 as the "Remote address range." Think of the "Server Address" as the default route for the IPs you are going to be issuing to the clients. It is also the virtual interface for the PPTP server.

If you are confused here, or in step 3, please go back and read the section named "Subnetting and VLAN routing" as it covered this in more detail.

5.

If you have a RADIUS server of some sort feel free to fill in the next few boxes. I don't so they are blank on this example and frankly go outside of the scope of this document anyway.

6.

If you are really security conscious, and your client software supports it, check the box to require 128-bit encryption.

7.

Click "Save" We are all done setting up the server. Now let's setup some users.

9.6. PPTP User Setup

If you have a RADIUS server and you set it up in the previous section you can either choose to skip this one, or add users here that will be found and used before the PPTP Server sends a request to the RADIUS server.

For the rest of us, this stage is quite important as we need a user account to authenticate to the PPTP Server.

1.

Click on "users" under PPTP in the VPN section of the StarShield interface.

2.

Click the "+" icon and lets fill in some blanks!

3.

Enter a name in the "Username" box.

4.

Enter and then re-enter the password for this account. (You can't use special characters at the time of this document, just FYI.)

5.

Click "Save"

6.

When you get back to the next window you will need to click "Apply Settings" NOTE: This will disconnect any active PPTP connections. Being as we are just setting this up for the first time, and this is our first user, let's hope there aren't any to disconnect.

7.

If everything went well you should have a screen that looks something like Figure 2.

Now we need to setup a firewall rule so people using the PPTP connection can do something with it when they connect.

9.7. PPTP Firewall Rules

Yep you need to do this if you want the darn thing to work. But just like your LAN rule, you can make this as open or as restrictive as you want. Here you can limit the PPTP users to accessing only specific hosts on specific ports, or open it all up. We are going to assume you want full access for your PPTP users so we are going to setup a firewall rule that is exactly like the default LAN rule.

1.

Start by clicking "Rules" under the firewall section of the StarShield interface.

2.

Next click any of the "+" Icons on the screen so we can add a new rule.

As stated we are going to allow all our PPTP users to access all parts of the LAN, WAN, etc. If you wish to limit this access then you will need to modify things accordingly. I will present one example of such a rule after this default section.

3.

Simply go to the "Interface" section and select PPTP from the drop down. In the Description put something meaningful like "Default PPTP -> any."

4.

Click Save

5.

You will have to Apply the changes on the next screen.

You are now done setting up the PPTP Server!

9.7.1. Example of filtered PPTP Rules

In some cases, most for those people who are granting PPTP access to others they do not fully trust, you will want to limit access (Specific Allow Rules) or mitigate specific access with Deny Rules. With specific allow users would be granted explicit permission to access hosts, and sometimes specific ports, and all other traffic is denied. The latter would be done if you wanted the PPTP clients to access the LAN & WAN but did not want them to access your SAMBA server for instance.

Our example is an allow rule granting permission for people on the PPTP network to use SSH on a LAN server with the IP address 192.168.1.151:

Save and Apply these rules as needed. Test them all to make sure they are working as designed. Most networks are compromised because no one checked the ACLs were activated or even working properly.

9.8. Setting up a PPTP Client on Windows XP™

This is super easy, and you only have to type one piece of information the entire time!

Start by accessing the Network Connections Panel. (do this however you like, I prefer to right click "Network Places" and select Properties.)

1.

Click "Create New Connection" in the left hand column of the "Network Connections" window.

2.

You are now presented with a Wizard. Click Next to continue.

3.

Select "Connect to the Network at my Workplace" from the menu.

4.

Select Virtual Private Network connection from the next panel.

5.

Name the connection.

6.

Now enter the IP or FQDN of the PPTP Server. (This can be any of the configured interfaces.)

7.

If you are the system admin you will be asked if you want this to be for your use only or for anyone's use. I suggest you limit it to your use only unless you want the VPN network to be made available to all user accounts on the workstation.

8.

Next you can either just finish or add a shortcut to the desktop. You are nearly done!

9.

When you launch the client for the first time (hopefully from the icon you asked it to create from the wizard, if not then you will need to access the "Network Connections" window again and double click your new connection.) you will be asked for a username and password. Click connect when you are done with this and if all goes well you will connect to the PPTP Server.

9.8.1. Testing our PPTP Connection in Windows [™]

1.

Start by opening a DOS window. (Command window)

2.

Run ipconfig and you should get something similar to the next figure:

As you hopefully will see you have the settings for your physical adapter (in my case I renamed it to ETH0)

You will also see the PPP Adapter with the name you gave the VPN Connection when performing the steps in the last section. It should have an IP address that is in the range you defined for the PPTP Server. It should also have the subnet of 255.255.255.255 and it will be using itself as the default gateway. Just live with it; it is how it works.

For the more advanced who wish to know if things are all working right, Figure 6, displays a full ipconfig on the virtual adapter.

3.

Now lets try doing something. If you followed the setup for this how-to you will have setup full access from the PPTP network to the LAN and WAN. If you setup selective rules you will have to test specifically what you setup. E.g. if you setup rules to only allow SMTP you will need to telnet to the host:25 that you designated in the firewall rule. Or write a new rule allowing ICMP to a host that will echo a reply back.

We will be sending a ICMP (Ping) to the firewall's internal interface to test the VPN connection.

4.

In my case the firewall is 192.168.1.1 (please use your internal address before writing to me to say pinging 192.168.1.1 didn't work on your 10.x.x.x network. Hehe) If done right (assuming your firewall isn't blocking internal ICMP packets) you are good for LAN access. (If you are blocking ICMP on the internal interface ping some other host on your home network.)

5.

Now lets test beyond the firewall. Ping isn't so good to use here as more and more people are blocking ICMP packets. So we will use tracert to check we are 1.) Routing via the PPTP tunnel and 2.) That we successful. Of course if you told the firewall to not allow WAN access then this step can be skipped.

As seen in the last figure, the first hop is the PPTP "Server Address" as this is the gateway/interface for the PPTP Network.

Now check things like HTTP, etc. If you have this much and followed the directions you should be able to do everything.

9.9. Some things I have found not to work over the PPTP Connection

These are more limits in PPTP than other VPN protocols.

*

NAT sometimes does not play nice with PPTP. Though StarShield seems to have this licked, and it works rather well.

*

Major "Gotcha!" If you are visiting a remote network where the network range is the same as the network range on the PPTP Network (your LAN network in most cases) then the PPTP tunnel will not work. E.g. You are using a WiFi connection in a local coffee shop and the network range it has put you in is 192.168.1.0/24. You try to connect to your home network via PPTP, but your home also uses 192.168.1.0/24. The tunnel/authentication to the PPTP server will happen, but no traffic will go across that tunnel due to the "confusion" in the TCP/IP stack on your workstation. To get around this use some odd network range at home. E.x. 192.168.88.0/24. Most people use 10.0.0.1 and 192.168.1.0 so try to set your home network differently. This will also help when you setup IPSEC tunnels between your house and say your friend's house.

*

Some ISP's use unreasonably short DHCP lease times, like one hour. If the PPTP client machine gets a short lease from DHCP, it will lose internet connectivity after the lease expires. This is because all network traffic, including your DHCP renewal requests, are going across the VPN. Since it can't hit the local DHCP server through the VPN, when the lease expires your machine will release its IP address. This causes the loss of all connectivity. You have to disconnect from the PPTP (if it doesn't disconnect itself), renew your IP address, and reconnect. This is common on Windows hosts, and likely other OS's as well. If this happens, contact the administrator of your DHCP server (likely the client machine's ISP) and get the lease time lengthened.

The author has seen this situation numerous times, and in every case, the ISP was willing to help and resolved the problem. Your mileage may vary.

*

UPnP packets from your LAN do not make it to the PPTP network. This is more than likely because the current version of StarShield does not support UPnP. (In English: those of use having dreams of accessing our ReplayTV™ or other media devices that use UPnP can dream of other things for now. It is actually more secure to not have UPnP on a firewall, but some people overlook that so they can use voice chat software and DVRs.)

*

Network Neighborhood in Windows does not work over PPTP connections because broadcasts are not forwarded across the PPTP connection.

I haven't really beaten the PPTP tunnel that much yet, so if you find more items that don't seem to work right let me know and I will add them here so people don't go crazy trying to figure out something that just won't work. ;)

Chapter 10. OpenVPN

OpenVPN is a new addition to StarShield in the 1.2 beta versions. Currently there is little documentation available.

Road warrior scenario - Peter Curran

Wireless network scenario - Peter Curran

For more information, see the OpenVPN project website.

Chapter 11. Wireless

Chapter 12. Captive Portal

Chapter 13. Reference

Table of Contents

13.1. IP Basics

13.2. IP Filtering

13.3. NAT

13.4. Traffic Shaping

13.5. DNS

13.6. Encryption (PPTP/IPsec)

13.7. Logging (syslog)

13.1. IP Basics

You can change the hostname and domain used by your firewall in the General Setup screen.

13.2. IP Filtering

13.3. NAT

NAT (Network Address Translation) permits you to use private IP address space on your LAN while still being able to access the internet.

There are two main types of NAT in StarShield, inbound, and 1:1.

13.4. Traffic Shaping

13.5. DNS

You can change the DNS servers used by your firewall in the General Setup screen.

13.6. Encryption (PPTP/IPsec)

13.7. Logging (syslog)

Log messages include a timestamp of when the event occurred. The system time on the firewall is synchronized to an NTP (Network Time Protocol) server. You can change the NTP server and related parameters in the General Setup screen.

It is recommended that you log your StarShield to a remote syslog server for diagnostics and forensic purposes. There are a number of free tools that do this for you on Windows, Mac, and Unix based systems.

Unix-based tools

The syslog daemon built into virtually every Unix-like system can be configured to accept log messages from remote hosts. Check documentation specific to your OS on how to configure syslogd to accept messages from remote hosts.

Other Unix Tools

syslog-ng

nsyslog

Windows-based tools

There are several free and commercial tools available on Windows to enable your system to accept syslog messages from hosts on your network.

Kiwi Syslog

One of my favorites on Windows is Kiwi Syslog. There is a version with "basic" features that is free, and a more advanced version with \$49 registration. Even if you are just looking for a free tool, the basic version has as many if not more features than any other free package on this list.

<http://www.kiwi-enterprises.com/>

3Com offers a couple of free utilities on this page. 3CSyslog is a GUI tool best used on a temporary or as-needed basis only. To collect logs using a service that will be running at all times, whether or not anyone is logged into the machine, try wsyslogd.

Several more for Windows and a couple for Mac listed on this site.

Chapter 14. Example Configurations

Table of Contents

14.1. Configuring a DMZ Interface Using NAT

- 14.1.1. Network Diagram
- 14.1.2. Adding the Optional Interface
- 14.1.3. Configuring the Optional Interface
- 14.1.4. Configuring the DMZ Interface Firewall Rules
- 14.1.5. Permitting select services from DMZ into the LAN
- 14.1.6. Configuring NAT

14.2. Locking Down DMZ Outbound Internet Access

14.3. Configuring a filtered bridge

- 14.3.1. General Configuration
- 14.3.2. WAN Configuration
- 14.3.3. OPT Interface Configuration
- 14.3.4. Enable Filtering Bridge
- 14.3.5. Configure Firewall Rules
- 14.3.6. Completing the Configuration

14.1. Configuring a DMZ Interface Using NAT

This section will explain how to add a DMZ interface to the two interface (LAN/WAN) base configuration from the Quick Start Guide.

You must have a functioning two interface setup before starting on configuring your DMZ interface.

The 1:1 NAT DMZ setup is most appropriate where you have multiple public IP's and wish to assign a single public IP to each DMZ host.

14.1.1. Network Diagram

This depicts the network layout we will have after configuring our DMZ interface.

14.1.2. Adding the Optional Interface

Log into your StarShield's webGUI, and click "(assign)" next to Interfaces.

Click the on this page to add your third interface.

Now restart your StarShield for the changes to take affect.

14.1.3. Configuring the Optional Interface

After your StarShield restarts, log back into the webGUI. Under Interfaces, you will see OPT1. Click on it.

Check the box at the top to enable the interface, give it a more descriptive name (I'll call it "DMZ"), and set up the desired IP configuration. The IP subnet must be different from the LAN subnet.

14.1.4. Configuring the DMZ Interface Firewall Rules

The main purpose of a DMZ is to protect the LAN from the publicly-accessible Internet hosts on your network. This way if one of them were to be compromised, your LAN still has protection from the attacker. So if we don't block traffic from the DMZ to the LAN, the DMZ is basically useless.

First we will put in a firewall rule on the DMZ interface denying all traffic to the LAN while still permitting all traffic to the WAN. Click Firewall -> Rules, and click the at the bottom of the page.

Filling out this screen as shown below will permit all traffic out the DMZ interface to the internet, but prohibit all DMZ traffic from entering the LAN. It also only permits outbound traffic from the DMZ's IP subnet since only traffic from a source IP within your DMZ should come in on the DMZ interface (unless you have a routed DMZ, which would be strange). This prevents spoofed packets from leaving your DMZ.

Click Save after verifying your selections. Then click Apply Changes.

14.1.5. Permitting select services from DMZ into the LAN

You probably have some services on your LAN that your DMZ hosts will need to access. In our sample network, we need to be able to reach DNS on the two LAN DNS servers, cvsup protocol to our LAN cvsup-mirror server, and NTP for time synchronization to the time server that resides on the cvsup-mirror server.

Always use specific protocols, ports, and hosts when permitting traffic from your DMZ to your LAN. Make sure nothing that isn't required can get through.

Note

Don't forget that source ports (TCP and UDP) are randomly selected high ports, and not the same as the destination port. You'll need to use "any" for source port.

My DMZ interface firewall rules now look like the following after permitting the required services from DMZ to LAN.

Note that I added a rule to deny any traffic coming in on the DMZ interface destined for the LAN. This was not required because of the way we configured the allow rule, however I like to put it in there to make it very clear where the traffic from DMZ to LAN is getting dropped.

When entering your rules, remember they are processed in top down order, and rule processing stops at the first match. So if you had left the rule we added above as the top rule, it would drop packets from DMZ to LAN without getting to the permit rules you added. I recommend you design your rules similar to how I have, with drop DMZ to LAN as the second last line, and permit DMZ to any except LAN as the last line.

14.1.6. Configuring NAT

Now you need to determine whether you'll use inbound or 1:1 NAT. If you have multiple public IP's, use 1:1 NAT. If you have only a single public IP, you'll need to use inbound NAT. If you have multiple public IP's, but more DMZ hosts than public IP's, you can use inbound NAT, or a combination of 1:1 and inbound.

14.1.6.1. Using 1:1 NAT

For this scenario, we'll say we have a /27 public IP subnet. We'll say it's 2.0.0.0/27. StarShield's WAN interface has been assigned with IP 2.0.0.2. I will use 1:1 NAT to assign the public IP 2.0.0.3 to the DMZ mail server and 2.0.0.4 to the DMZ web server.

Go to the Firewall -> NAT screen and click the 1:1 tab. Click the . I will add two entries, one each for the mail server and web server.

After adding the rules, click Apply changes. You'll now see something like the following.

14.1.6.2. Testing the 1:1 NAT Configuration

You can test the 1:1 NAT we just configured by going to whatismyip.com on the machine configured for 1:1. If you don't have a GUI, lynx will work, or you can fetch or wget the URL and cat the resulting file. (fetch `http://whatismyip.com && cat whatismyip.com | grep "IP is"`).

You should see the IP is the one you just configured in 1:1 NAT. If you get an IP other than the one you configured in 1:1, there is a problem with your configuration.

14.1.6.3. Using Inbound NAT

If you have only one public IP, or more need more publicly-accessible servers than you have public IP addresses, you'll need to use inbound NAT. Go to the NAT screen, and on the Inbound tab, click .

For this example, we will assume you have only one public IP, and it is the interface address of the WAN interface.

First, anything to the WAN IP to port 25 (SMTP) will go to the mail server in our DMZ.

Click Save, and click to add the inbound NAT rule for the HTTP server.

Click "Apply changes" and your configuration will be working. It should look like the following.

14.2. Locking Down DMZ Outbound Internet Access

We've limited DMZ hosts' accessibility to the LAN, but we can lock it down a step further using egress filtering. Many DMZ hosts don't need to be able to talk out to the Internet at all, or possibly only while you are running updates or doing maintenance or need to download software.

If we can keep our DMZ hosts from accessing the Internet, we can make an attacker's job much more difficult. Many exploits rely on the target being able to pull files from a machine the attacker controls, or in the case of a worm, from the infected host. I'll use Code Red and Nimda as an example. Infected hosts exploited the vulnerability, and the remote host pulled the infected admin.dll via TFTP from the already infected host. If you were running vulnerable web servers, but did not allow TFTP traffic outbound from your webservers, you could not have been infected. (reference)

Attackers most always try to pull in a tool kit or root kit of some sort onto machines they exploit. There are ways around this, but it just makes it that much more difficult. This will merely slow down a knowledgeable attacker (who'll find a way to get in one way or another), but it could stop a script kiddie dead in their tracks and keep some worms from infecting your network.

This is not a replacement for proper patching and other security measures, it's just good practice in a defense-in-depth strategy.

How does this work? You might be wondering how your servers will be able to serve content while not being able to talk out to the Internet. I'll use web servers as an example. When packets come in on the WAN interface through firewall rules you have entered to permit HTTP traffic, there is a state entry that permits any return traffic from that connection to traverse the firewall. Remember this only affects the ability to initiate connections outbound, not the ability to respond to incoming traffic requests.

Recommended configuration. As with all firewall rules, limit the accessibility as much as possible. Mail servers that must send outbound mail will need to initiate connections to destination TCP port 25 to any host. If the DNS servers your DMZ hosts use reside outside of the DMZ, you'll need to allow UDP port 53 to the DNS servers being used. I typically put in rules for upgrade purposes to permit outbound traffic to the ports required. For FreeBSD, TCP 5999 (cvsup) and TCP 80 (HTTP) will generally suffice. When I'm not upgrading the system, I use the "disable" checkbox to disable the rule, but leave it in place to easily enable it when needed. Just always remember to disable it when you're done updating the system.

14.3. Configuring a filtered bridge

A filtered bridge is a common way of configuring a DMZ segment. This can be used as a typical DMZ where you have hosts on the LAN interface, but is probably more frequently used to protect servers at a colocation facility where there are no LAN hosts.

Note

Remember you cannot access hosts on a bridged interface from a NAT'ed interface, so if you do have a LAN interface set up, you won't be able to access the hosts on the bridged interface from the LAN.

Network Diagram for this Configuration. The following diagram depicts the example configuration described in this section. The colocation facility has assigned you with the subnet 111.111.111.8/29, which includes usable IP's .9-.14. One of those is required for the colo's router, so you end up with 5 usable IP's.

14.3.1. General Configuration

After you have your network set up as shown, and the interfaces and LAN IP assigned appropriately, log into the webGUI to begin the initial configuration.

First go to System -> General setup, and configure the hostname, domain, DNS servers, change the password, switch the webGUI to HTTPS, and set your time zone. Click Save, and reboot StarShield for the changes to take affect.

14.3.2. WAN Configuration

Log back into the webGUI and go to the Interfaces -> WAN page. For the example network, we'll assign the static IP 111.111.111.10/29, default gateway 111.111.111.9. Unless your WAN network is private IP's, check the "Block private networks" box. Click Save.

14.3.3. OPT Interface Configuration

Click Interfaces -> OPT. Name the interface to your liking (for the example, we'll use Servers for the name). In the "Bridge with" box, select WAN. Click Save.

14.3.4. Enable Filtering Bridge

Go to the System -> Advanced page and check the "Enable filtering bridge" box. Click Save.

14.3.5. Configure Firewall Rules

Go to the Firewall -> Rules screen.

Note

Chances are for any configuration, especially if you're restricting outbound connections, you'll need a much more involved ruleset than is depicted here. Open what you know you need open, and watch for dropped traffic in your logs to see what else you might need to open. It takes some effort to get your firewall locked down as tightly as it can possibly be, but the long term effect of increased security is well worth the time spent.

14.3.5.1. OPT Interface Rules

Initially, you may want to configure a rule on the OPT interface permitting traffic to anywhere, then after things are working, tightening that rules as desired. For this example, we'll go ahead and implement locked down rules from the get go.

The mail server on our bridged interface needs to send mail to any host on the Internet. Both servers need to get to DNS servers at 111.111.110.2 and 111.111.109.2. We'll add disabled maintenance rules for HTTP and cvsup.

14.3.5.2. WAN Interface Rules

Since this example portrays a firewall at a colocation facility, we need a remote administration rule to allow traffic from our trusted location's static IP access to administration functions of the servers, as well as the StarShield webGUI. For this example, we'll permit all traffic from the trusted location (IP 11.12.13.30). You may want to tighten this rule. If you don't have anything on the LAN segment, remember to allow remote administration from somewhere so you can get into the webGUI without being on site.

We also need to add rules to permit SMTP traffic to the mail server and HTTP and HTTPS traffic to the web server.

14.3.5.3. LAN Interface Rules

You can leave or remove the default LAN to any rule if you don't have hosts on the LAN interface. In the example, the LAN interface will be unplugged once the onsite configuration is completed.

14.3.5.4. Firewall Rules Completed

14.3.6. Completing the Configuration

Everything should be working as desired now, as long as the servers are configured appropriately. Test that the configuration works as desired, including all inbound and outbound rules. Once you're satisfied with the testing results, your setup is complete.

Chapter 15. Example Site to Site VPN Configurations

Table of Contents

15.1. Cisco PIX Firewall

- 15.1.1. PIX Configuration
- 15.1.2. StarShield Configuration

15.2. Smoothwall

15.3. FreeS/WAN

15.4. Sonicwall

- 15.4.1. Sonicwall Configuration
- 15.4.2. StarShield Configuration

15.5. Nortel

StarShield can connect to any third party VPN device that supports standard IPsec site to site VPN's, which includes most any VPN device and firewall with IPsec VPN support.

This chapter will provide instructions on connecting StarShield with a number of third party IPsec devices.

Have you configured a VPN between StarShield and a device not listed here? Please document how you accomplished this. There is a section of the wiki dedicated to configurations for this chapter.

15.1. Cisco PIX Firewall

The following describes how to configure a site to site IPsec VPN tunnel between a PIX Firewall and StarShield.

15.1.1. PIX Configuration

First we need to make sure the PIX has 3DES enabled.

```
pixfirewall# sh ver

Cisco PIX Firewall Version 6.3(3)
Cisco PIX Device Manager Version 2.0(2)

Compiled on Wed 13-Aug-03 13:55 by morlee

pixfirewall up 157 days 5 hours

Hardware: PIX-515E, 32 MB RAM, CPU Pentium II 433 MHz
Flash E28F128J3 @ 0x300, 16MB
BIOS Flash AM29F400B @ 0xffffd8000, 32KB

0: ethernet0: address is 000b.4605.d319, irq 10
1: ethernet1: address is 000b.4605.d31a, irq 11
2: ethernet2: address is 0002.b3b3.2e54, irq 11
Licensed Features:
Failover: Disabled
VPN-DES: Enabled
VPN-3DES-AES: Enabled
```

If the "VPN-3DES-AES" line above does not show "Enabled", you need to install the PIX 3DES key. This is now available free from Cisco here for all PIX firewalls (click 3DES/AES Encryption License). Do NOT use DES

for a VPN if you want it to be cryptographically secure. DES is only slightly better than transmitting in clear text.

Next we'll see if any VPN configurations are in place on the PIX.

```
pixfirewall# sh isakmp policy
```

```
Default protection suite
encryption algorithm: DES - Data Encryption Standard (56 bit keys).
hash algorithm: Secure Hash Standard
authentication method: Rivest-Shamir-Adleman Signature
Diffie-Hellman group: #1 (768 bit)
lifetime: 86400 seconds, no volume limit
```

If you only see the default policy, there are no VPN's configured. This document cannot be followed verbatim if you have current VPN's (though you should be able to figure it out, just be careful not to break your existing VPN's with any duplicate names).

Allow IPSec connections to the PIX

```
pixfirewall(config)# sysopt connection permit-ipsec
```

Enable ISAKMP on the outside interface (where "outside" is the name of the internet-facing interface)

```
pixfirewall(config)# isakmp enable outside
```

isakmp policy command on PIX

```
pixfirewall(config)# isakmp policy ?
Usage: isakmp policy %lt;priority> authen %lt;pre-share|rsa-sig>
isakmp policy %lt;priority> encrypt %lt;aes|aes-192|aes-256|des|3des>
isakmp policy %lt;priority> hash %lt;md5|sha>
isakmp policy %lt;priority> group %lt;1|2|5>
isakmp policy %lt;priority> lifetime %lt;seconds>
```

Now we need to configure the ISAKMP policy on the PIX. Enter the following commands in configure mode:

```
isakmp policy 10 authen pre-share
isakmp policy 10 encrypt 3des
isakmp policy 10 hash md5
isakmp policy 10 group 2
isakmp policy 10 lifetime 86400
```

This policy uses pre-shared keys as authenticator, 3DES encryption, md5 hashing, group 2, and 86400 second lifetime.

Now we need to define the pre-shared key for this connection. (1.1.1.1 = public IP address of StarShield, qwertyuiop is the shared key, randomly generate something to use for your configuration)

```
isakmp key qwertyuiop address 1.1.1.1 netmask 255.255.255.255
```

Now we need to create an access list defining what traffic can cross this tunnel.

```
access-list monovpn permit ip 10.0.0.0 255.255.255.0 10.0.1.0 255.255.255.0
access-list monovpn permit ip 10.0.0.0 255.255.255.0 10.0.1.0 255.255.255.0
```

Define transform set for this connection called "monovpnset"

```
crypto ipsec transform-set monovpnset esp-3des esp-md5-hmac
```

Define security association lifetime

```
crypto ipsec security-association lifetime seconds 86400 kilobytes 50000
```

Now to set up the actual connection, the crypto map "monovpnmap". (where 1.1.1.1 is the public IP address of the StarShield device)

```
crypto map monovpnmap 10 ipsec-isakmp
crypto map monovpnmap 10 set peer 1.1.1.1
crypto map monovpnmap 10 set transform-set monovpnset
crypto map monovpnmap 10 match address monovpn
```

These lines specify type of VPN (ipsec-isakmp), peer IP address (1.1.1.1), transform set to be used (monovpnset, defined above), and that packets matching the access list "monovpn" created above should traverse this VPN connection.

Last step is to tell the PIX to not use NAT on the packets using this VPN connection and route them instead.

First we'll see if anything is currently routed.

```
pixfirewall# sh nat
nat (inside) 0 access-list no-nat
```

Look for "nat (interface) 0 ..." commands. The above means any traffic matching access list "no-nat" will be routed, not translated. In this instance, we are adding to a current access list (if you use a DMZ, you likely have something similar to this set up).

```
access-list no-nat permit ip 10.0.0.1 255.255.255.0 10.0.1.0 255.255.255.0
access-list no-nat permit ip 10.0.1.0 255.255.255.0 10.0.0.0 255.255.255.0
```

If you do not have a "nat (interface) 0 ..." command in your "sh nat" output, you can use the above two lines to create a "no-nat" access list. You then have to apply it with the "nat (interface-name) 0 access-list no-nat" command (replacing "interface-name" with the name of your LAN interface).

15.1.2. StarShield Configuration

Log into the StarShield web GUI, and under VPN, click IPsec.

If the "Enable IPsec" box is not checked, check it and click Save.

Click the + button to add a VPN tunnel. On the "Edit tunnel" screen, fill in as follows:

Leave "Disable this tunnel" box unchecked.

Interface "WAN"

Local subnet: Type: "LAN subnet"

Remote subnet: 10.0.0.0 /24 (fill in the subnet of the network behind the PIX here, rather than the made-up 10.0.0.0/24)

Remote gateway: public IP address of PIX

Description: add one to describe the connection (e.g. "PIX VPN")

Phase 1

Negotiation mode: Aggressive

My identifier: "My IP Address"

Encryption algorithm: 3DES

Hash algorithm: MD5

DH key group: 2

Lifetime: 86400

Pre-shared key: qwertyuiop (enter exactly what you defined as your pre-shared key on the PIX earlier)

Phase 2

Protocol: ESP

Encryption algorithms: only 3DES checked

Hash algorithms: only MD5 checked

PFS key group: 2

Lifetime: 86400

Note

In StarShield 1.2 beta versions, you may experience the connection dropping frequently with this configuration. If this happens, set the PFS key group in phase 2 to "off".

Note

If you don't specify a key lifetime in the StarShield config, the tunnel will work, but appear to go insane after a while. Supposedly Cisco's will negotiate a key lifetime, but I have not seen this work in my experience. This is also true of a Cisco VPN Concentrator. (anonymous wiki contribution)

15.2. Smoothwall

Rev. Tig posted the following information on connecting Smoothwall and StarShield via IPsec VPN in a post on the mailing list on September 30, 2004.

I could not find a working solution in the mailing list archives but here is how I have managed to create a VPN between Smoothwall Corporate with Smoothtunnel and StarShield and I thought I would share it here to same people going through the same headbashing experience I did :) This will be far to much of a teaching granny to suck eggs for most people on the list but it might help someone get up and running quickly.

Variety is the spice of life and just to confuse matters the StarShield box was stuck behind NAT :) The office I was linking to was in a serviced building and hence the connection was a shared one with a private IP and public one port forwarded to it.

I had never done this before so corrections are welcome :) I am not saying these are the best settings all I know is my VPN is up and running and it seems to be happy :)

What I have created is a VPN between one subnet at one site running Smoothwall Corporate Server 3.0 with Smoothtunnel and a StarShield v1 box sitting behind NAT with a private IP at the other site. Any other versions of the software may need slightly different settings but hopefully this should put you in the right ballpark.

StarSystem IT s.r.l.

Sede Legale ed Operativa: Località Fratte, 49 – 38057 Pergine Valsugana (TN)

Cod. Fisc. e P.IVA 00698480225 - Tel. 0461.185.10.40 – Fax 0461.019.979

Internet: www.starsystem.biz - e-mail: info@starsystem.biz

First off IPSEC over NAT, if at all possible don't :) If you have to or for some perverse reason you fancy a crack at this then read on, if you are just here for the Smoothwall bit scroll down :)

IPSEC over NAT does work but it can be a case of sacrificing the odd network card to the deity of your choice, what I did in the end was ask their network guy to just send everything and I will let m0n0 do the firewalling, this is what I would recommend as then you don't have to hassle them every time you want a port opening, but from what I have gathered is that all you need are port 500 forwarding and IP protocols 50 and 51 to be routed but the firewall. Apparently your IPSEC traffic goes through port 500 but IP protocols 50 and 51 are needed for phase 1 (authentication) and phase 2 (key exchange). If I am wrong (this is quite possible there will be a load of mails below correcting me :) If m0n0 is behind NAT and you are certain the other end is right but there appears to be no attempts to authenticate then check here first.

Now onto Smoothwall Corporate, now I know Rich Morrell posts on here so I have to be careful about what I say about the interface but that is just a personal taste thing :)

Right here are the Smoothwall settings :

Local IP : your RED IP address (if you are using Smoothhost then put the IP of your firewall in)

Local ID type: Local IP

Remote IP : the external IP of your NATted StarShield box.

Remote ID type : Remote IP

Authenticate by : Preshared Key

Preshared Key : put your shared key here

Use Compression : Off

Enabled : On

Local network : in this case it was 192.168.0.0/255.255.255.0

Local ID value : same as your Local IP

Remote network: in this case it was 192.168.1.0/255.255.255.0

Remote ID value : the same as your Remote IP

Initiate the connection : Yes

I will use these networks in this example as it shows you a little gotcha in StarShield that threw me because I was not thinking :)

Next block :

Local Certificate : (your local certificate)

Perfect Forward Secrecy : Yes

Authentication type: ESP (it has to be AH will NOT work over NAT)

Phase 1 crypto algo: 3DES

Phase 1 hash algo : MD5

Key life : 480 (mins)

Key tries : 0 (never give up)

Right now the StarShield settings :

Phase 1:

Mode : tunnel (well you can't change it and why would you want to :)

Interface : WAN

Local Subnet : 192.168.1.0 / 24 (don't do what I did and select LAN :)
Remote Subnet : 192.168.0.0 / 24
Remote IP : The RED IP of your Smoothwall box
Negotiation Mode : Main
My Identifier : IP Address : Your public IP (non NATed) for your
StarShield box
Encryption Algo: 3DES
Hash Algo : MD5
DH Key Group : 5
Lifetime : (blank)
Preshared Key : put your shared key here.

Phase 2:
Protocol : ESP
Encryption Algo: 3DES (only! untick the others)
Hash Algo: MD5 (again only)
PFS Key Group : 5
Lifetime : (blank)

That is it, you can now bring the link up from Smoothwall by going
into the VPN control tab and clicking UP!

15.3. FreeS/WAN

Josh McAllister provided the following sample ipsec.conf, which can be used to connect StarShield with
FreeS/WAN in a site to site IPsec configuration.

```
# /etc/ipsec.conf - FreeS/WAN IPsec configuration file

version 2.0 # conforms to second version of ipsec.conf specification

config setup
    interfaces=%defaultroute
    klipsdebug=none
    plutodebug=none
    uniqueids=yes

# defaults for subsequent connection descriptions

conn %default
    # How persistent to be in (re)keying negotiations (0 means
very).
    keyingtries=0
    #compress=yes

conn block
    auto=ignore

conn private
    auto=ignore

conn private-or-clear
    auto=ignore

conn clear-or-private
    auto=ignore
```

```
conn clear
  auto=ignore
```

```
conn packetdefault
  auto=ignore
```

```
conn josh
  type=tunnel
  left=ip.add.of.m0n0
  leftsubnet=m0n0.side.subnet/24
  leftnexthop=%defaultroute
  right=ip.add.of.freeswan
  rightsubnet=freeswan.side.subnet/24
  rightnexthop=%defaultroute
  authby=secret
  auth=esp
  esp=3des-md5-96
  pfs=no
  auto=start
```

```
m0n0-side:
Phase1
Neg. mode = main
Enc. Alg = 3DES
Hash Alg = MD5
DH key grp = 5
```

```
Phase2
Protocol = ESP
Uncheck all Enc. Alg. Except 3des
Hash alg = md5
PFS key group = off
```

15.4. Sonicwall

Contributed by Dino Bijedic < dino.bijedic (at) eracom-tech (dot) com >

The following describes how to configure a site to site IPSec VPN tunnel between a Sonicwall (PRO 300) and StarShield.

Editor's note: I would suggest using Main mode rather than Aggressive.

Figure 15.1. Network diagram
Network diagram
15.4.1. Sonicwall Configuration

Log in to Sonicwall

Click VPN -> Configure

Add/Modify IPSec Security Association

In Configure, select Security Association -> Add New SA

Name: Name of connection (Monowall test)

IPSec Gateway Name or Address: Type IP address of your StarShield (203.49.X.117)

Security Policy

Exchange: Aggressive Mode

Phase 1 DH Group: Group2

SA Life time (secs): 28800

Phase 1 Encryption/Authentication: 3DES & MD5

Phase 2 Encryption/Authentication: Strong Encryption and Authentication (ESP 3DES HMAC MD5)

Share Secret: type your share secret (novitest)

Destination Networks

Select "Specify destination network below".

The following screenshot shows what this screen will look like.

Click Add New Network

You will get: Edit VPN Destination Network (Note: This is Popup window – enable Popup in your browser)

Network: type your destination network (192.168.200.0)

Subnet mask: Type destination subnet mask (255.255.255.0)

Click Update

Figure 15.2. Example of Sonicwall configuration

Example of Sonicwall configuration

15.4.2. StarShield Configuration

Configure StarShield IPsec Edit Tunnel screen as follows.

Interface: WAN

Local subnet: LAN subnet

Remote subnet: 192.168.2.0/24

Remote gateway: 61.95.x.99

Description: Sonicwall

Negotiation mode: Aggressive

My identifier: My IP address

Encryption algorithm: 3DES

Hash algorithm: MD5

DH key group: 2

Lifetime: 28800

Pre-shared key: novitest

Protocol: ESP

Encryption algorithms: 3DES

Hash algorithms: MD5

PFS key group: off

Lifetime: 28800

Click Save at the bottom of the page to complete the VPN configuration.