
User's Guide

Tiny Personal Firewall V.2

Contents

Getting Started Chapter 1

System Requirements	4
Installation	5
Personal Firewall Architecture	6

Administration Chapter 2

Overview	8
Run as Service	9
Personal Firewall Status	10
Remote Administration	12

Setting up Security Chapter 3

Firewall Description	15
Trustful Address Group	16
LAN users	17
MD5 Signature	18
Security Levels	19
Minimal Security	19
Medium Security	20
Maximum Security	24
Creating Filter Rules	25
Intro to TCP/IP	26
Time Intervals	27
Endpoints	28
Ordering and Precedence	29
Port Addressing	30

Contents

Logs and Packet Analysis

Chapter 4

Creating Useful Logs	32
Interpreting the Logs	33
Syslog.....	34

Index

35

CHAPTER 1**GETTING STARTED****In This Chapter**

System Requirements	4
Installation	5
Personal Firewall Architecture.....	6

System Requirements

586 Pentium class

16 MB of RAM

1 MB of free storage

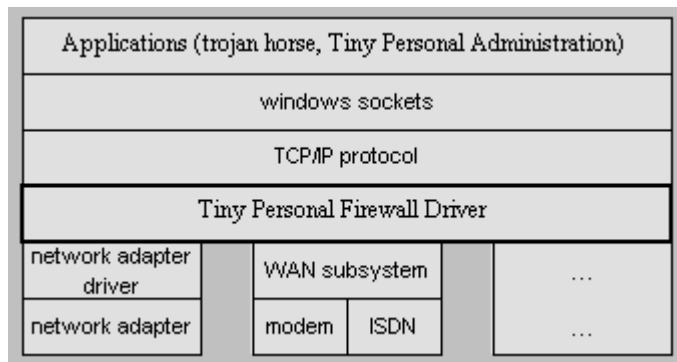
Windows 9x/2000/ME/NT

Installation

After double clicking on the Personal Firewall self-extracting utility you will be asked which folder you would like the program to be installed into. The default folder is the program files folder. By selecting the browse option you may change the folder that Personal Firewall will be installed into. You must restart the machine after installation.

Personal Firewall Architecture

Tiny Personal Firewall uses the administration utility as an interface for communicating with the engine, which communicates with the Personal Firewall driver. The driver resides on the lowest possible level of the operating system, just above the physical hardware drivers. This way, personal firewall is always the first and last line of defense.




ADMINISTRATION

In This Chapter

Overview	8
Run as Service	9
Personal Firewall Status	10
Remote Administration.....	12

Overview

Tiny Personal Firewall is divided into three utilities: the engine, the status monitor and the administration.



The **Engine** must be activated before you can access the administration utility. If you run Personal Firewall as a system service (recommended) the engine will automatically appear in the system tray during Windows startup. Otherwise, you must activate the engine from the program files folder. By right clicking on the engine icon you may access the administration window or exit the engine to disable the Personal Firewall driver.

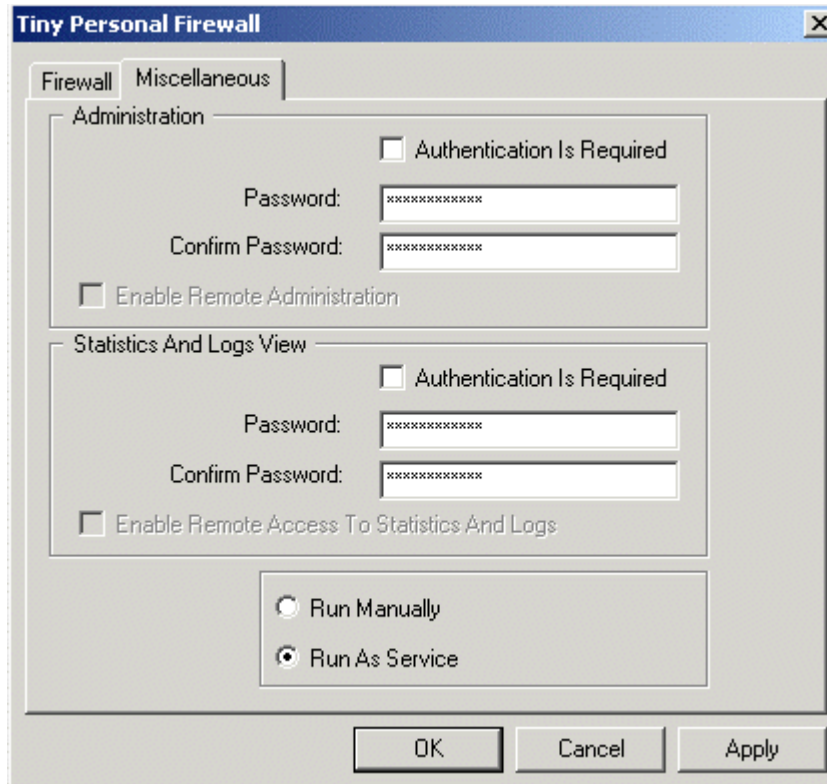
The **Administration** utility is the interface that allows you to modify your firewall settings and is discussed throughout the remainder of the manual.

The **Status** monitor allows you to see the ports that each application is bound to, as well as other useful information associated with each binding. This utility is described in more detail later in this chapter.

All three utilities can be activated from the Start menu -> Programs -> Tiny Personal Firewall.

Run as Service

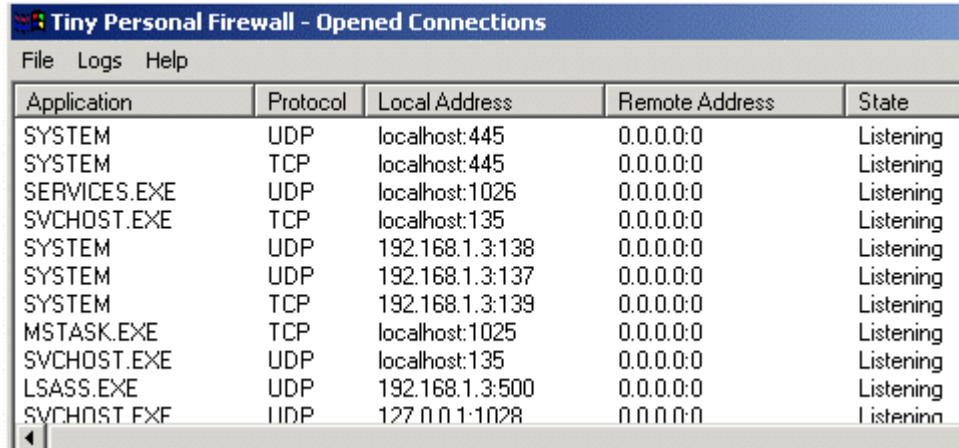
For optimal security, it is suggested that you run Tiny Personal Firewall as a service. This will allow the Personal Firewall driver to function the moment after communication is possible from your computer. All of your filtering rules will be active long before any application, such as a trojan horse, has the chance to send or receive data. To run Personal Firewall as a service you must enable the option from the Miscellaneous window of the Administration utility.



Personal Firewall Status

Applications and system services that communicate with other computers tell your operating system that all incoming packets with a particular destination port should be forwarded to the application or service listening on the specified port. The status monitor displays various information about all of these applications and services that are granted communication.

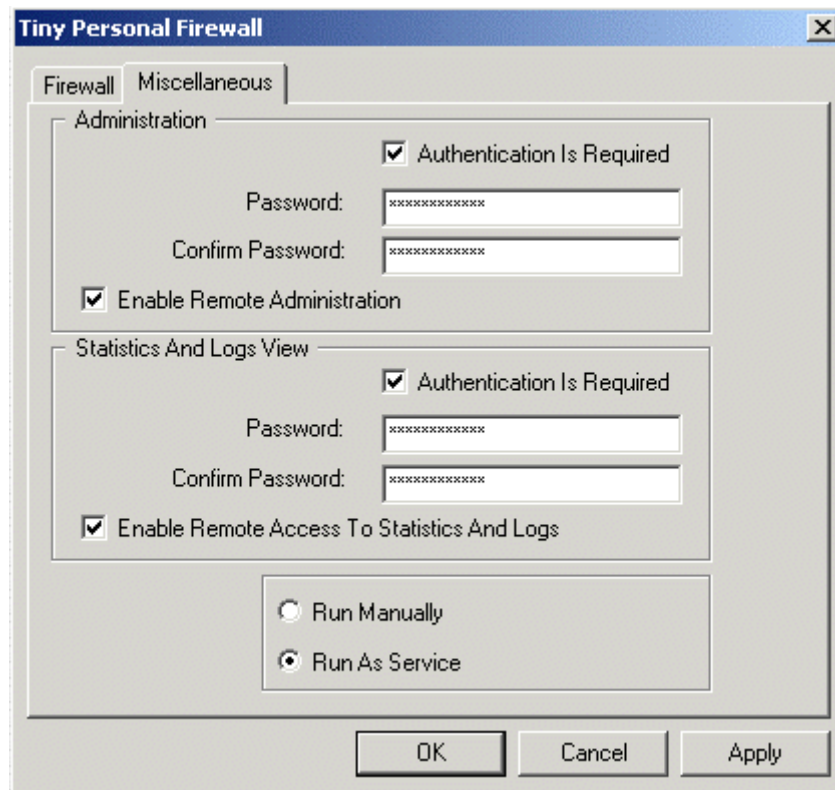
The **application** column simply states the application, such as `icq.exe`, or system, for service, that has a permitted open connection. The **protocol** will specify if the application or service is bound to either TCP or UDP. The **local address** refers to your computer, and includes the port number that the application or service is bound to. If there is active communication between the service or application, the **remote address** will display the IP address of the other party and the port that it is receiving connections from. Also, the **state** will display 'connection out.' If there is no established communication, the state will remain 'listening.' The **creation time** displays the date and time (military) that the most recent or current connection was established. The **Rx** and **Tx** refer to received data and transmitted data.




Application	Protocol	Local Address	Remote Address	State
SYSTEM	UDP	localhost:445	0.0.0.0:0	Listening
SYSTEM	TCP	localhost:445	0.0.0.0:0	Listening
SERVICES.EXE	UDP	localhost:1026	0.0.0.0:0	Listening
SVCHOST.EXE	TCP	localhost:135	0.0.0.0:0	Listening
SYSTEM	UDP	192.168.1.3:138	0.0.0.0:0	Listening
SYSTEM	UDP	192.168.1.3:137	0.0.0.0:0	Listening
SYSTEM	TCP	192.168.1.3:139	0.0.0.0:0	Listening
MSTASK.EXE	TCP	localhost:1025	0.0.0.0:0	Listening
SVCHOST.EXE	UDP	localhost:135	0.0.0.0:0	Listening
LSASS.EXE	UDP	192.168.1.3:500	0.0.0.0:0	Listening
SVCHOST.EXE	UDP	127.0.0.1:1028	0.0.0.0:0	Listening

Remote Administration

To allow Personal Firewall to be remotely configured you must enable this feature from the miscellaneous tab in the main window. Additionally, you may allow your status and logs to be viewed remotely.



The Beta version of personal firewall includes the functionality of remote administration; however, the option of remote access is not provided within the administration interface. To remotely access the status monitor you must go to the start menu -> run. Select the browse option and locate the administration program. Append **-remote** to the end of the directory path to remotely access the status monitor. You will be prompted to enter the IP address and password for the remote client. To remotely access the administration, you will follow the same steps as above and append **-remote -cfg**. The full string for the picture below would look like this "C:\Program Files\Tiny Personal Firewall\PFWADMIN.exe" -remote -cfg.



SETTING UP SECURITY

In This Chapter

Firewall Description	15
Trustful Address Group	16
LAN users	17
MD5 Signature	18
Security Levels	19
Creating Filter Rules	25

Firewall Description

The Transmission Control Protocol/Internet Protocol (TCP/IP) is the only means of communication accross wide area networks such as the Internet, and is the most common form of communication accross local area networks. The Internet Protocol (IP) is responsible for the identification of each packet. For example, when you address a letter you must follow a standard format so that the post office is able to process your letter. Similarly, the Internet Protocol attaches a header to each data packet. Within this header is standardized information such as the packet's protocol, who the packet is addressed to, where it's coming from, what ports to use (if applicable), if it is a response or a request and so on. A true firewall, such as Tiny Personal Firewall, uses a driver that is capable of inspecting this header information and as a result can filter, translate, or record based on specified criteria.

The core functionality of Personal Firewall is its stateful inspection capability. Before packets leave your computer, Personal Firewall inspects the header and records who it's going to and what port it will return on. When the response packet comes back to your computer, Personal Firewall will compare the packet to its records to make sure that you requested the packet. If it is not recognized, the packet is simply dropped.

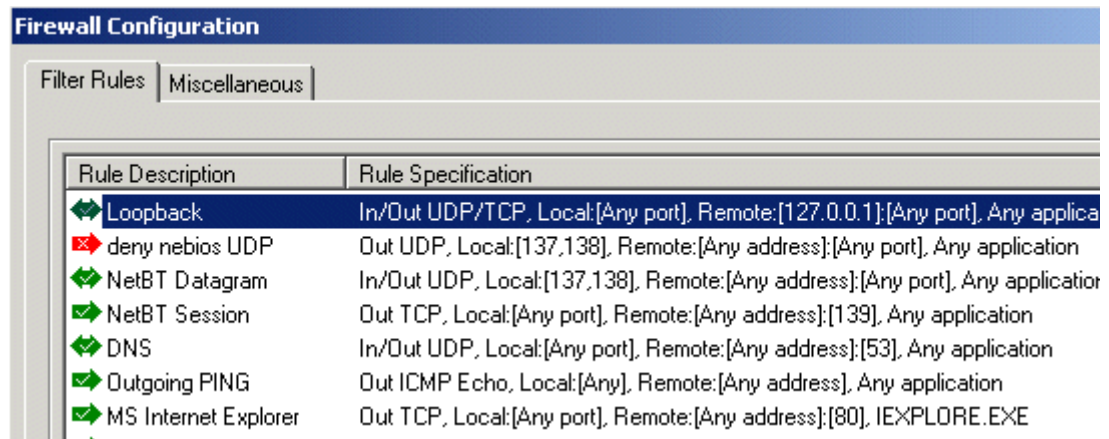
Port addressing is an important element in data communication. Applications and services bind to and listen on particular ports creating a high level of vulnerability. A firewall can ensure that these applications and services don't allow unwanted intruders to access your machine through an open port.

Trustful Address Group

In the advanced settings -> miscellaneous tab you will find the Trustful Address Group. This option is for simplification purposes. By themselves, any entries made to the address group will have no effect. The values inputted into the address group must be referenced by filter rules. For example, if you are on a Local Area Network and you want to allow shares to the LAN, then you would specify the IP address range or subnet of your LAN in the trustful address group. Then you must create a filter rule that allows incoming TCP/UDP packets with local endpoint of ports 137-139 (NetBios) and a remote endpoint from the trustful address group on any port. If you wanted your shares to be available to additional parties, say out over the Internet, you would specify those IP addresses in the trustful address group without having to create an additional filter rule. Note that the trustful address group can also be referenced by 'deny' rules.

LAN users

If you are in a local area network, your computer will broadcast UDP packets using ports 137 and 138. These packets tell other LAN users your name and workgroup. Since these packets are frequently broadcasted, we provided a filter rule to permit all datagrams on ports 137 and 138 so that the user will not encounter several prompt screens from the wizard. This means that your computer name and workgroup name are available to anybody. If you do not wish to disclose such information you must add the following filter rule: For **outgoing UDP** packets specify the **local endpoint** as a **range from 137 to 138**. Leave all other information as **"any."** **Deny** all packets that fit this description. Make sure you give the rule a description like "Deny UDP netbios," and that the rule is at the top of the list.




Filter Rules	
Miscellaneous	
Rule Description	Rule Specification
↔ Loopback	In/Out UDP/TCP, Local:[Any port], Remote:[127.0.0.1]:[Any port], Any applica
✖ deny netbios UDP	Out UDP, Local:[137,138], Remote:[Any address]:[Any port], Any application
↔ NetBT Datagram	In/Out UDP, Local:[137,138], Remote:[Any address]:[Any port], Any application
↕ NetBT Session	Out TCP, Local:[Any port], Remote:[Any address]:[139], Any application
↕ DNS	In/Out UDP, Local:[Any port], Remote:[Any address]:[53], Any application
↕ Outgoing PING	Out ICMP Echo, Local:[Any], Remote:[Any address], Any application
↕ MS Internet Explorer	Out TCP, Local:[Any port], Remote:[Any address]:[80], IEXPLORE.EXE

MD5 Signature

When creating permitting rules for applications, the user is informing Tiny Personal Firewall that an application named `someapp.exe` located in `c:\program files\someapp.exe` is allowed to bind to whatever ports the user says the application can use. This ensures, for the most part, that trojan horse applications do not have the right to bind to any ports for communication. It is possible, however, for a trojan horse to spoof its identification so the user may see `outlook.exe`, for example, and permit it without realizing that `outlook.exe` may be a trojan horse application in disguise. To prevent this application spoofing, Tiny Personal Firewall includes MD5 signature authentication support. MD5 is a hash algorithm that takes a 128 bit fingerprint of an application. Each time the application requests to bind to a particular port, Tiny Personal Firewall can take a fingerprint of the application and compare it to the original fingerprint. It is virtually impossible to duplicate a fingerprint, so trojan horse applications don't stand a chance.

Within the advanced filtering rules option the user may choose to check for a digital signature (MD5) for a particular application.




Security Levels

Minimal Security

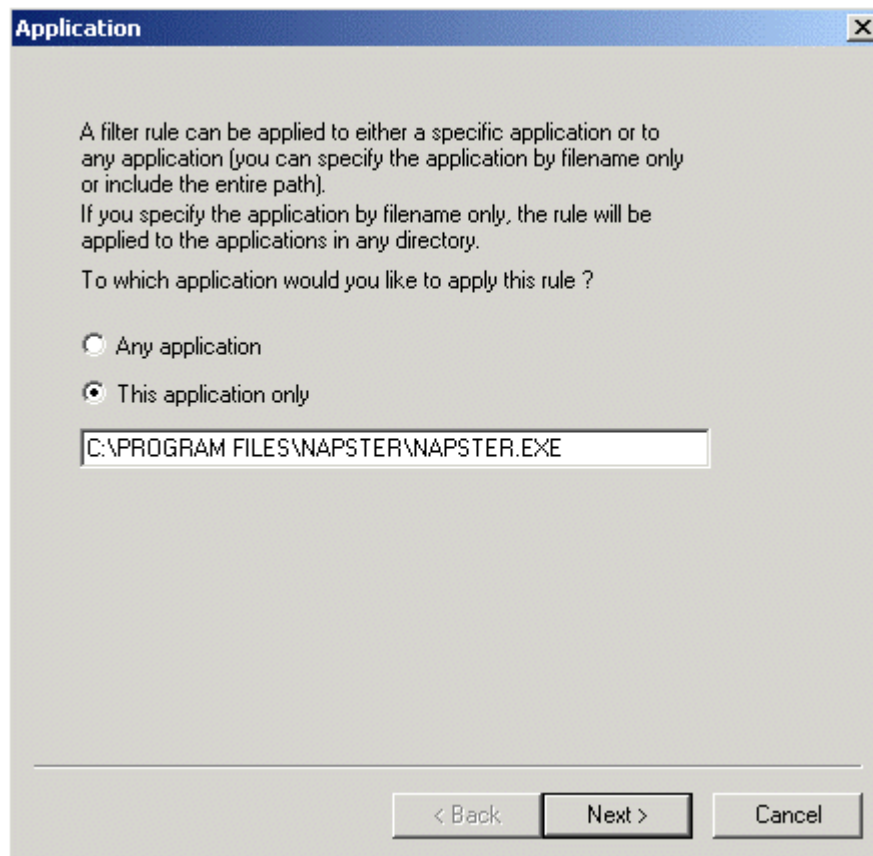
Minimum security permits any service or application to transmit and receive data. It is intended for those users who want to create their own filter rules. Note: the rule "ask for action if no rule is found" is not applicable to minimal security.

Medium Security

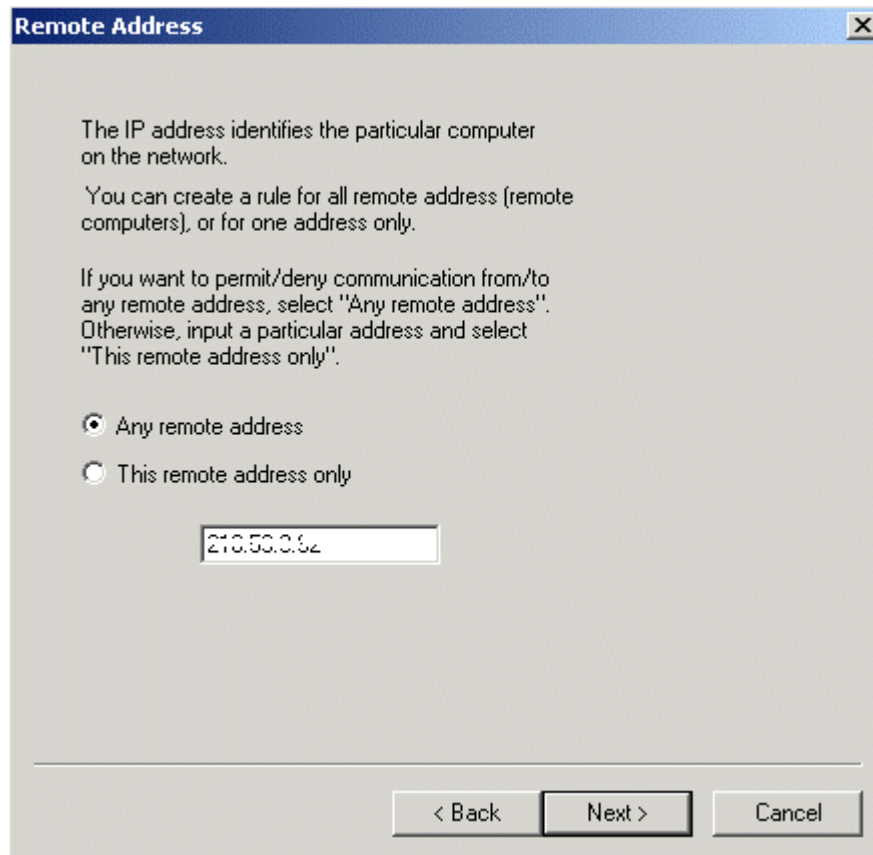
Medium security is the default security level and is recommended for most users. Contrary to minimal security, medium security blocks all IP traffic and expects the user to supply permission policies. To help the user out, we have provided a few predefined filter rules, which may be removed at the users discretion. The option "ask for action when no rule is found," should be enabled. This option is at the bottom of the filter rules list found in the advanced menu. Enabling it will prompt a wizard script to help walk you through the setup of new filter rules when an unknown packet is encountered. If you disable this option, any packet that does not fall within any filter rules or the known applications database will be dropped.



The wizard provides four options as shown in the previous screen shot. The two options on the left will only permit or deny the single packet that prompted the wizard. In most cases you will need to select one of the two options to the right and create a static rule for other packets of the same nature.



By selecting "any application," the port will be opened for any application. In most cases you will specify the default option "this application only."



The IP address identifies the particular computer on the network.

You can create a rule for all remote address (remote computers), or for one address only.

If you want to permit/deny communication from/to any remote address, select "Any remote address". Otherwise, input a particular address and select "This remote address only".

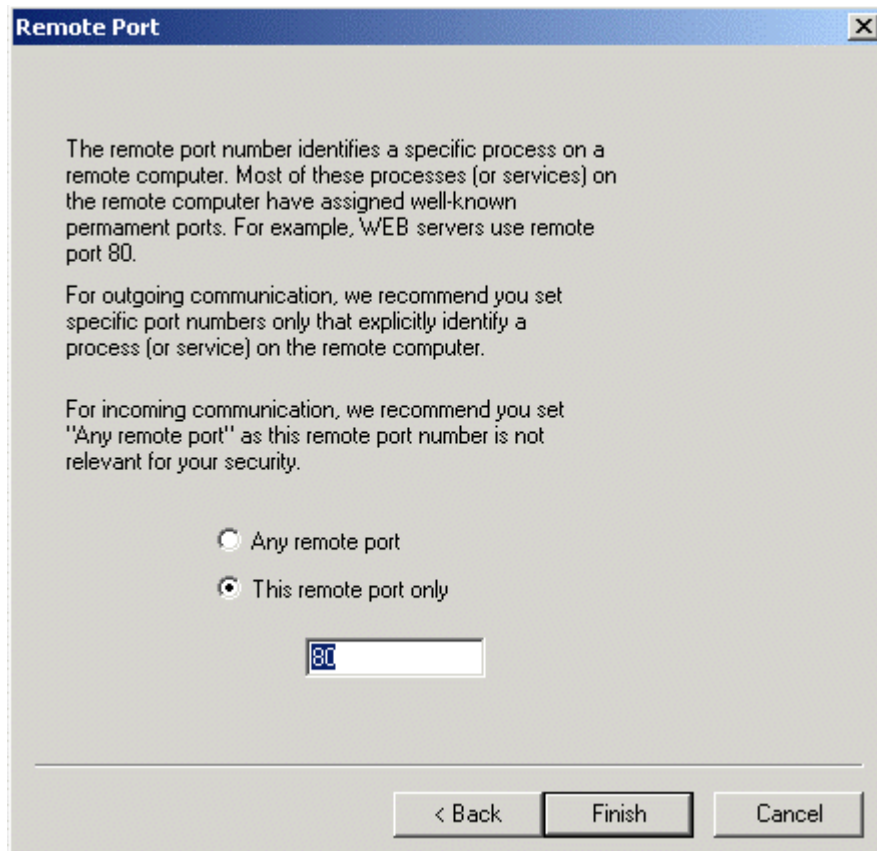
☒ Any remote address

☐ This remote address only

< Back Next > Cancel

The dialog box is titled "Remote Address" and has a close button in the top right corner. It contains instructional text about IP addresses and two radio button options. The first option, "Any remote address", is selected. The second option, "This remote address only", is unselected. Below the second option is a text input field containing the IP address "210.50.0.64". At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

You can limit outbound/inbound packets to/from an individual IP address. Most applications will communicate with several addresses so you will specify "any remote address." If you know that the application will communicate with a particular subnet or group of IP addresses you can specify the single IP address supplied by the wizard, and later edit the filter rule so that the remote endpoint will include additional IP addresses or a predefined list specified in the "trustful addresses."



The Remote Port dialog box is a standard Windows-style window with a title bar that says "Remote Port" and a close button (X) in the top right corner. The main area of the dialog contains three paragraphs of text explaining the purpose of the remote port number. Below the text are two radio buttons: "Any remote port" and "This remote port only". The "This remote port only" option is selected. Below the radio buttons is a text input field containing the number "80". At the bottom of the dialog, there are three buttons: "< Back", "Finish", and "Cancel".

Remote Port

The remote port number identifies a specific process on a remote computer. Most of these processes (or services) on the remote computer have assigned well-known permanent ports. For example, WEB servers use remote port 80.

For outgoing communication, we recommend you set specific port numbers only that explicitly identify a process (or service) on the remote computer.

For incoming communication, we recommend you set "Any remote port" as this remote port number is not relevant for your security.

☐ Any remote port

☒ This remote port only

80

< Back Finish Cancel

As indicated by the text, we recommend that the user specify a single remote port for each application. Many applications will send to several ports. For example, a web browser may also send to port 21 for FTP, or to 443 for HTTPS. By specifying "any remote port" you will eliminate the need to create other permitting rules; however, this may compromise security.

Maximum Security

Maximum security will incorporate all the filter rules and permissions from the known applications; however, the wizard is not available. If a packet does not apply to any rules it will be dropped. Maximum security works just the opposite of Minimum security. All traffic is blocked, therefore any acceptable communication must be specified through filter rules.

Creating Filter Rules

Intro to TCP/IP

In order to personalize your firewall it is important to understand the basics of data transmission. The Transmission Control Protocol/Internet Protocol is the only means of communication across the Internet. This is to ensure standardization so that routers can identify every packet and properly route them. The Internet Protocol is responsible for the addressing scheme of every packet that is to leave your computer. This addressing information is contained in a standard format header, that varies depending on the protocol.

ICMP (Internet Control Message Protocol) As the full name implies, ICMP is a messaging protocol. The 'ping' command is most commonly associated with ICMP packets. There are several flags that represent ICMP responses. The most common reply is the 'echo.' Personal Firewall can block both incoming and outgoing 'echo' replies.

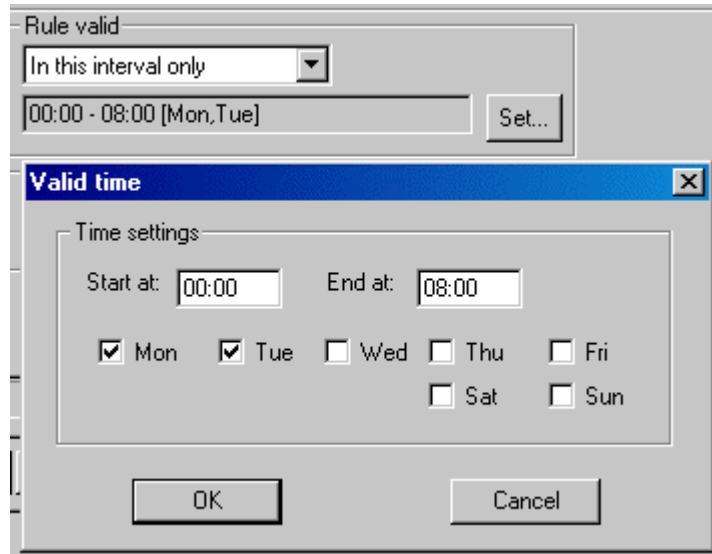
UDP (User Datagram Protocol) Uses a special type of packet called a datagram. Datagrams do not require a response, they are one way only. Datagrams are usually used for streaming media because an occasional packet loss will not affect the final product of the transmission.

PPTP (Point-to-Point Tunneling Protocol) Is used mostly by companies, who have set up a Virtual Private Network (VPN). A VPN creates a tunnel between a client and server that uses various forms of encryption. Individual rules cannot be made for PPTP, however Personal Firewall can support pass through of any PPTP packet.

TCP (Transmission Control Protocol) The Transmission Control Protocol is host to a large suite of protocols such as FTP, Telnet, HTTP, SMTP, POP3 and so on. TCP is based on two-way communication, meaning that the sending party requires an acknowledgement before sending more packets.

Time Intervals

You can set time frames on individual filter rules. In the following example I have validated a filter rule only on Mondays and Tuesdays from midnight until 8:00 a.m. Personal Firewall always uses military time.



Endpoints

The **local endpoint** always refers to the computer where Tiny Personal Firewall resides. For **outgoing** TCP or UDP packets, the local endpoint can be defined by a single or range of ports. In most cases you would leave this option as 'any' because the source port may be dynamically picked from a range. A better option is to select filter rules for a specific **application**. This is necessary if you are in minimal security level, or you have disabled the option to ask for an action if no rule is found, or you want to allow communication from a particular application during specified intervals. For **incoming** packets, the local endpoint port number would be the destination port of the arriving packet. This value is, under most circumstances, only relevant for server applications that are listening for packets destined to a particular port. Some services and server applications listen on multiple or non-standard ports, therefore it is recommended that you specify the application and leave the port number as 'any.'

The **remote endpoint** always refers to the other party. For example, if you request a URL from CNN.com, the CNN.com web server is the other party. Or, if you were running an FTP server from your local machine, any person connecting to your FTP server is the remote endpoint. Since your computer is secured by personal firewall, you would need to allow access to your FTP site. To do this you would need to specify that for **incoming** packets that will be received by the FTP server application (which you will specify in the application field of the local endpoint) must be permitted always, or only during certain hours and days (which you will specify in 'rule valid'). You can also specify the remote endpoint by IP address and/or port. For **outgoing** packets (those initiated locally such as a URL request) you can permit or deny access to certain destinations. For example, you can deny any web communication over HTTP by specifying the remote endpoint by port number 80. Port 80 is the standard for HTTP communication. This and other common protocols are described in the Port Addressing chapter.

Ordering and Precedence

On an incoming or outgoing packet the personal firewall driver captures the packet and compares it against several rules. If the packet falls within a particular rule's conditions, then it follows the action provided by the rule and is dismissed from any further interrogation. For this reason it is important to order rules appropriately.

An **incoming** packet is first compared against the filter rules. **Filter rules** are compared against in **top down** order. If a packet does not meet the criteria on any user defined filter rules, then it is compared against its record table. If a match is found in the record table then the packet is permitted as a returning packet to a recent request that was initiated locally.

An **outgoing** packet is also compared first against the filter rules. If you are in minimal security then no further checks are performed. If you are in medium security, and the outbound packet did not apply to any filter rules, then the packet is compared against the database of permitted applications.

Port Addressing

This section does not contain a direct application to Personal Firewall. However, it is important to understand the following concepts.

The Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) both require port addressing. Applications and services bind to particular ports, so when a TCP or UDP packet arrives at your computer, the operating system knows to forward the packet to the application listening on the port that the packet was directed to.

Ports are a driving force behind firewalls and routers. Packets are typically routed or denied based on source and/or destination port information. The Internet Assigned Numbers Authority (IANA) is responsible for maintaining the association of each port. Ports 0 to 1023 are the well known ports. The most common TCP protocols such as HTTP, FTP, Telnet, POP3, SMTP and so on are within the well known ports and, therefore, use standard port allocations across wide area networks. For example, HTTP uses port 80, so a web server requires port 80 to be open in order to receive web requests. For a list of all port assignments refer to the IANA website at <http://www.isi.edu/in-notes/iana/assignments/port-numbers>.

Every TCP and UDP packet contains a source and destination port. Referring to the previous example, let's say you request information from <http://www.cnn.com>. Your web browser sends the information to the TCP/IP stack. The Internet Protocol (IP) recognizes the request as HTTP and assigns the destination port as 80. An available source port is also assigned to the packet and remains used for a period of time, or until a packet returns on that port. This returning packet will have a reversed port assignment, meaning that the source port would be 80 and the destination port as the port your web browser is expecting a return from. Note that [cnn.com](http://www.cnn.com) is a domain name, which means that before you can send any request to [cnn's](http://www.cnn.com) web server, your web browser must first send information to a domain name server, requesting the IP address associated with the provided domain name.

LOGS AND PACKET ANALYSIS

In This Chapter

Creating Useful Logs	32
Interpreting the Logs	33
Syslog	34

Creating Useful Logs

Each individual filter rule offers the option to log information pertaining to each packet that matches the rule. This way you can isolate log information. In the following chapter a filter rule was created to deny and log any incoming packet arriving on port 139 (NetBios TCP). NetBios is the Microsoft protocol that enables sharing over TCP/IP.

Interpreting the Logs

The following diagram describes three lines taken from a sample filter.log. All logs follow this format.

[07/Nov/2000 15:45:27]	Rule 'NetBT Session': Blocked:	
[07/Nov/2000 15:45:30]	Rule 'NetBT Session': Blocked:	
[07/Nov/2000 15:45:36]	Rule 'NetBT Session': Blocked:	
<u>Time Stamp</u>	<u>Name of rule</u>	<u>Action</u>
TCP 192.168.10.144:1368->localhost:139, Owner: SYSTEM		
TCP 192.168.10.144:1368->localhost:139, Owner: SYSTEM		
TCP 192.168.10.144:1368->localhost:139, Owner: SYSTEM		
<u>Remote Address</u>	<u>Local Address</u>	<u>Listening Service\Application</u>
<u>Protocol</u>	<u>Remote Port</u>	<u>Local Port</u>

Syslog

You can choose to send all of your log information to a syslog server. The Beta version allows the user to enable this feature from the advanced miscellaneous window; however, there is no support within the administration window to specify the IP address of the syslog server. In the config file located in the Tiny Personal Firewall folder you will find the following information. Change the value of "LogIntoSyslog" to "1" and "SyslogIpAddress" to the IP of the syslog server.

```
[HKEY_LOCAL_MACHINE\Software\TinySoftware\PersFw\Config\FwCfg]
```

```
"FirewallEnabled"="1"
```

```
"LogIntoFile"="1"
```

```
"UserResultRequired"="1"
```

```
"LogIntoSyslog"="1"
```

```
"SyslogIpAddress"="192.168.10.2"
```

```
"SecurityLevel"="3"
```

```
"FilterFragmentedPackets"="0"
```

INDEX

A

Administration • 7

C

Creating Filter Rules • 25

Creating Useful Logs • 32

E

Endpoints • 28

F

Firewall Description • 15

G

Getting Started • 3

I

Installation • 5

Interpreting the Logs • 33

Intro to TCP/IP • 26

L

LAN users • 17

Logs and Packet Analysis • 31

M

Maximum Security • 24

MD5 Signature • 18

Medium Security • 20

Minimal Security • 19

O

Ordering and Precedence • 29

Overview • 8

P

Personal Firewall Architecture • 6

Personal Firewall Status • 10

Port Addressing • 30

R

Remote Administration • 12

Run as Service • 9

S

Security Levels • 19

Setting up Security • 14

Syslog • 34

System Requirements • 4

T

Time Intervals • 27

Trustful Address Group • 16