SSH Sentinel 1.2
User Manual

October 2001

This document describes the SSH Sentinel software, an IPsec client product by SSH Communications Security Corp, providing secure communications over a TCP/IP connection.

SSH Communications Security Oyj

Fredrikinkatu 42
FIN-00100 Helsinki
FINLAND

SSH Communications Security Inc.
1076 East Meadow Circle
Palo Alto, CA 94303
USA

SSH Communications Security K.K.
House Hamamatsu-cho Bldg. 5F
2-7-1 Hamamatsu-cho, Minato-ku
Tokyo 105-0013, JAPAN

http://www.ssh.com/

ipsec-sales@ssh.com (sales), ipsec-support@ssh.com (technical support)

Phone: +358 20 500 7030 (Finland), +1 650 251 2700 (USA), +81 3 3459 6830 (Japan)

Fax: +358 20 500 7031 (Finland), +1 650 251 2701 (USA), +81 3 3459 6825 (Japan)

Chapter 1 **Introduction**

1.1 About This Document

This document describes SSH Sentinel, an IPSec client product by SSH Communications Security Corp, providing secure communication over a TCP/IP connection.

More information on the SSH Sentinel software is available on the SSH Communications Security web site (http://www.ssh.com/).

1.2 About SSH Sentinel

SSH Sentinel is a software product for securing Internet Protocol (IP) based traffic using the IPSec protocol< as specified by Internet Engineering Task Force (IETF) standards.

SSH Sentinel is an easy-to-use product designed for end users. It allows you to encrypt and authenticate important network connections, like remote access to corporate networks remote administration, file transfer, sending and receiving email (SMTP, POP) and IP telephony.

SSH Sentinel software currently supports the following Microsoft Windows operating systems: Windows 95, Windows 98, Windows NT4, Windows Me and Windows 2000. Next, the software will be available on the Linux platform.

SSH Sentinel is designed to be a *client* type IPSec application. The features are designed for a single user workstation using a single network adapter and the Internet Protocol (IP). SSH Sentinel supports all network connection types, including dial-up. The product is designed to be secure and robust, easy to use, and quick to adapt to the environment at hand. Key characteristics include intuitive installation and configuration, as well as an easy way to use certificates for authentication.

SSH Sentinel was implemented due to numerous customer and end-user requests to bring out a real IPSec solution for commercial platforms and to enable full-scale network encryption with strong authentication.

1.3 Internet Protocol

The open architecture of the Internet Protocol (IP) makes it a highly efficient, cost-effective and flexible communications protocol for local and global communications. It is widely adopted, not only on the global Internet, but also in the internal networks of large corporations.

The Internet Protocol was designed to be highly reliable against random network errors. However, it was not designed to be secure against a malicious attacker. In fact, it is vulnerable to a number of well-known attacks. This is preventing it from being used to its fullest for business and other purposes involving confidential or mission-critical data. The most common types of attacks include:

- Eavesdropping on a transmission, for example, looking for passwords, credit card numbers, or business secrets.
- Taking over communications, or hijacking communications, in such a way that the attacker can inspect and modify any data being transmitted between the communicating parties.
- Faking network addresses, also known as IP spoofing, in order to fool access control mechanisms based on network addresses, or to redirect connections to a fake server.

1.4 Internet Protocol Security (IPSec)

Internet Engineering Task Force (IETF) has developed the Internet Protocol Security (IPSec) protocol suite to prevent misuse and attacks on IP. IETF is an international standards body with representation from hundreds of leading companies, universities, and individuals developing Internet-related technologies. Its track record includes the Internet Protocol itself and most of the other protocols and technologies that form the backbone of the Internet.

The IPSec protocol suite adds security to the basic IP version 4 protocol and is supported by all leading vendors of Internet products. IPSec is a mandatory part of the next generation of IP protocol, IP version 6.

The IPSec protocol works on the network level. It adds authentication and encryption to each data packets transmitted. It protects each packet against eavesdropping and modification, and provides authentication of the origin of the packet.

IPSec works independently of any application protocol. Thus, all applications that use IP protocol for data transfer are equally and transparently protected. IPSec makes it safe to use the Internet for transmitting confidential data. By doing so, it solves the main obstacle that is slowing down the adoption of the Internet for business use.

However, IPSec alone does not solve the security problems in many operating systems and network applications. It often offers some protection against these problems, and often makes a break-in attempt much more traceable. Nonetheless, it must still be understood that operating system and application security cannot be overlooked. Furthermore, for smooth operation, IPSec requires a public key infrastructure. Such infrastructures are still in their infancy, and wide-scale

key infrastructures are just emerging on the Internet. All in all, the management of security policies and access policies is an extremely complicated field, and there are no magical solutions.

IPSec does, however, solve some of the most critical Internet security problems. It renders most of the commonly used attack methods completely ineffective. It does this by providing confidentiality, integrity and authentication of traffic.

1.5 SSH Communications Security

SSH Communications Security Corp (SSH) is an engineering organization whose core competence is in designing and implementing network security protocols. SSH is currently the world's largest company dedicated to the development of IPSec, and the established world leader in IPSec technology. SSH does not only implement the protocol, but is also one of the developers of the IPSec standard. SSH is actively involved in the Internet Engineering Task Force.

SSH develops and licenses IPSec technology to the world's leading information technology corporations, telecommunications companies, router manufacturers and operating system vendors, providing the most complete, standards-compliant and widely deployed IPSec solution available.

Chapter 2 Installing and Removing SSH Sentinel

The installation of the SSH Sentinel software is a straightforward process guided by an installation wizard, and you should be able to complete it without studying this manual. This chapter in the manual explains the process and the terms encountered in more depth.

This beginning of this chapter describes the first installation of the SSH Sentinel software. During the installation, you create an authentication key pair and a matching certificate to be used for authentication. However, if a previous version of the software is already installed on your computer, then launching the installation only updates the software. The security policy rules and the authentication keys that you have configured with the previous version of the software, are preserved. Naturally, you can always remove the software completely and then reinstall it.

2.1 Requirements

SSH Sentinel is available on the most popular Microsoft Windows platforms. The supported platforms and their versions are listed in the table below.

Platform	Version Build	Notes
Windows 95	OSR1, OSR2	Winsock2 required
Windows 98	SE	-
Windows NT 4.0	SP3 to SP6	-
Windows Me		-
Windows 2000	SP1	-

SSH Sentinel is a client-type implementation of IPSec, not an IPSec gateway software, even though some of the Windows platforms are capable of functioning as routers.

Before starting SSH Sentinel installation, make sure that there are no other IPSec implementations, network sniffers, NAT applications, firewalls, or third party intermediate network drivers installed. SSH Sentinel may affect the functionality of such software.

To run SSH Sentinel, you need a personal computer with the following minimum configuration:

Processor	Pentium 100 MHz
Memory (RAM)	32 MB for Windows 9x, or 64 MB for Windows NT4/2000
Hard disk space	10 megabytes of free disk space
Network connection	TCP/IP network protocol

2.2 Starting the Installation

The SSH Sentinel installation requires you to have full access rights for the system files on your computer. On a Windows NT system, you will have to log in with administrator rights.

To begin the installation, double click the SSH Sentinel installation package `Sentinel.exe` icon in Windows Explorer. The package can be found on the SSH Sentinel CD or in the directory where you have downloaded SSH Sentinel. The self-extracting package will automatically initiate InstallShield(R) software to install and set up SSH Sentinel.

Figure 2.1 The SSH Sentinel installation package icon.

The installer will run Installation Wizard, which creates the initial configuration and sets up the SSH Sentinel software. N.B. If a previous version of the SSH Sentinel software is installed on your computer and you try to install a new version, the wizard updates the software and the steps described here are skipped.

When started, the Installation Wizard will first go through a sequence of basic installation dialogs, displaying the licensing agreement and allowing you to select the installation directory and the program folder.

Figure 2.2 The installation begins.

The installation can only be performed on the local computer. Remote installation of SSH Sentinel is not possible, because the installation program updates kernel mode components related to networking and remote access.

Note that the installation will terminate immediately if you do not accept the licensing agreement.

2.3 Authentication Key Generation

The SSH Sentinel Installation Wizard generates a primary authentication key for IPSec peer (host) authentication purposes. The primary authentication key is a 1024-bit RSA key pair that is used for digital signatures and strong authentication.

Authentication key generation begins with random seed generation. A random pool of data is collected from the user moving the mouse or typing in random text. The data is then used as a seed to ensure that all authentication keys will be unique. With this method, the likelihood of generating two identical authentication keys is infinitesimal.

The general level of security that can be provided with 1024-bit RSA authentication keys is considered military strength. The Internet Key Exchange (IKE) protocol used in key negotiation is better by design and security than most of the other solutions that currently exist.

The SSH Sentinel key generation process will take some 30 seconds and may momentarily use most of the computer's CPU resources.

Once the authentication key generation is complete, you may proceed with the installation.

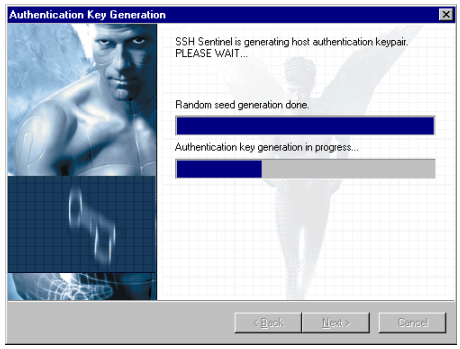

Figure 2.3 Generating the authentication key.

2.4 Identity Information

SSH Sentinel uses certificates and digital signatures as its primary authentication method. SSH Sentinel processes certificates according to the IETF Public-Key Infrastructure X.509v3 standards, allowing you to take advantage of the public-key infrastructure (PKI). SSH Sentinel supports certificate revocation lists (CRLs) and authority revocation lists (ARLs, that is, CRLs for CAs) and is very configurable. However, you can run the software as stand-alone, separately from any public-key infrastructure.

The setup requires host identity information that is to be associated with the authentication key pair and its certificate. A commonly preferred identity is the host DNS name, also referred to as the Fully Qualified Domain Name (FQDN). The DNS name should be used as the identity whenever the host has a static DNS name and whenever it is safe to assume that name service will be available. If the host does not have a static DNS name its static IP address may be used as the host identity.

If neither static DNS name nor IP address is available, you may use an email address as the identity. However, using an email address as the identity makes it difficult for remote hosts to bind IPSec rules for the host, since rules are normally bound to a host name or an IP address.

2.5 Choose the Enrollment Method

A certification request can be created as part of the installation process. You can either enroll online, in other words create and send the request immediately, or save the request in a file and deliver it later to the certification authority (CA). If there is no certification authority available or you for some reason want to postpone the creation of the request, create a self-signed certificate. It should be noted that once you've installed the software, you can create as many certification requests as you wish with the SSH Sentinel user interface but you cannot create a self-signed certificate with it.

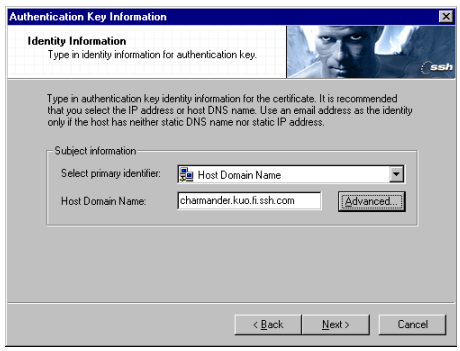

Figure 2.4 Inquiring certificate identity.

To create a self-signed certificate, select the topmost option *(CA is unavailable. SSH Sentinel uses..)*. The keys created in the previous step are used when the system creates the certificate.

To issue an online certification request, select the option *Request certificate online. Generate an online...* The installation wizard shows you next a dialog where further information on the certification authority and the enrollment protocol is asked for. See section2.5.1 'Online Enrollment Information' for reference.

To create a certification request and save it in a file, select the option *Request certificate off-line. Generate a ...* See section 2.5.2 'Off-line Certification Request' for the next step.

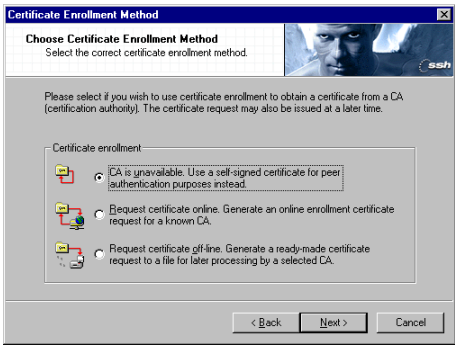

Figure 2.5 Choosing the enrollment method.

2.5.1 Online Enrollment Information

To enroll online, you need to be able to locate the certification authority server and you need to possess the certification authority certificate. Most often, you can fetch the certificate of the certification authority from its web site.

You must also specify the enrollment protocol. On top of all this, you may configure the socks and proxy settings to get through the firewall if the local server is protected by one.

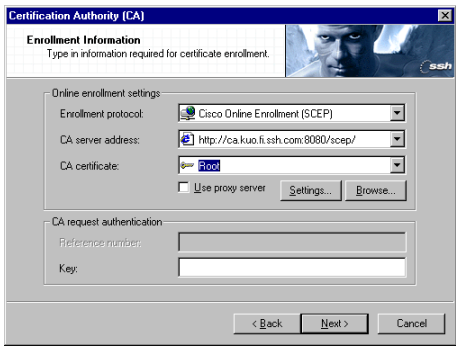

Figure 2.6 Inquiring information on the certification authority and the enrollment.

Enrollment Protocol (A)

Select the enrollment protocol from the drop-down list. Naturally, you should choose a protocol that is supported by the certification authority. The following protocols are available: Simple Certificate Enrollment Protocol (SCEP) and Certificate Management Protocol (CMP).

CA Server Address (B)

Specify the address (URL) of the certification authority web site.

CA Certificate (C)

The certificate of the certification authority is needed to encrypt the certification request before sending it to the certification authority. You can usually fetch it from the authority's Web site.

In the drop-down menu, you see the possibilities on how to import the certification authority certificate into the request: The most convenient way is to specify here the URL where the certificate is located. In this case, the certificate needs to be in PEM encoded format. SSH Sentinel then automatically fetches the certificate from the web site. You can also have fetched the certificate earlier using a web browser and either saved it in a file or copied the contents of it to the Windows clipboard. In a file, the certificate may be in binary (X.509), PEM (Privacy Enhanced Mail) or HEX format. Pasted from the clipboard, the certificate needs to be in PEM encoded format.

Advanced button (D)

Opens a dialog box for configuring the socks and proxy settings.

Reference Number (E)

Only in connection with the Certificate Management Protocol (CMP). The key identifier is used along with the key to identify the user requesting a certificate.

Key (F)

Only in connection with the CMP protocol. A shared secret granted by the certification authority to be used in the certification request. Used for verification of the user requesting a certificate.

2.5.2 Off-line Certification Request

An off-line certification request is simply a file, where the request is stored for later use. The request is of PKCS#10 format and saved in Privacy Enhanced Mail (PEM) encoded format.

To complete the enrollment, you need to deliver the request to the certification authority. You might save the request on a floppy disk and deliver the floppy to the authority or you might prefer sending the request via email or using an enrollment service on the Web.

Select PKCS#10 request file location

In the text field (A), specify the path and the name of the file where the certification request will be stored.

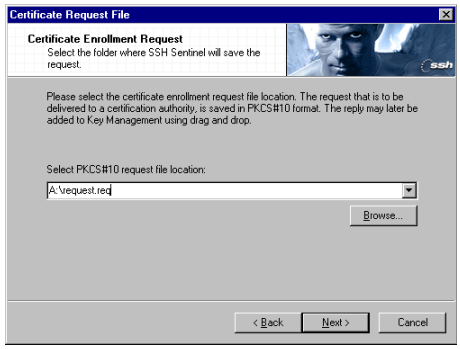

Figure 2.7 Off-line request: specifying where the request file is temporarily stored.

2.6 Encryption Speed Diagnostics

SSH Sentinel runs diagnostics on the encryption algorithms as the last step of the installation. You can bypass this step by clicking the `Skip` button on the dialog box.

The diagnostics reveal the speed of the encryption algorithm compared to each other. SSH Sentinel supports the following ciphers: Rijndael, Twofish, Blowfish, Cast, 3DES and DES. With the exception of DES, all these can be considered secure for commercial use. The DES encryption algorithm is supported as a fallback option for interoperability reasons. Rijndael, an encryption algorithm widely considered fast, secure and reliable, is used as the default cipher by SSH Sentinel.

The diagnostics also reveal the relative speed of your computer running the algorithms. There is a lot of contradictory information available on encryption speeds. The diagnostics give you the chance to use your own judgment.

The diagnostics measure the encryption speed of your computer within the memory. The data packets are not transmitted to the network. This is a common way to measure performance amongst the encryption hardware vendors. It has the advantage of giving simple figures on the speed: Due to a number of variables that affect the final result, it would be very complicated to define a standard environment where to reliably measure the overall network throughput. Moreover, the real-world network throughput simply cannot be measured during the installation, because the kernel mode IPSec engine is not available before the first reboot.

An Intel P3 personal computer with processor speed of 800 MHz should be able to provide a maximum IPSec throughput of over 40 Mbit/s on the preferred cipher. However, other variables, such as the operating system, network bandwidth and CPU load, naturally set limitations to the throughput.

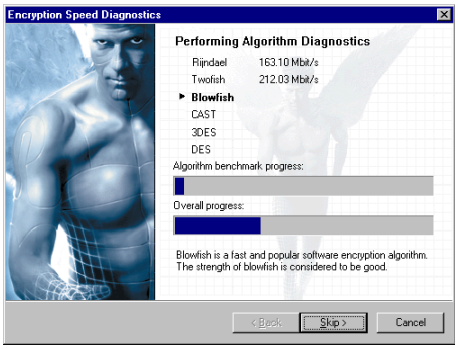

Figure 2.8 Measuring the encryption speed.

2.7 Completing the Installation

The installation of the SSH Sentinel software adds kernel-mode components to the operating system network management. For this reason, you have to restart the computer before using the software.

Figure 2.9 SSH Sentinel installation is now complete.

2.8 Updating SSH Sentinel

The system automatically updates the software, if you launch the installation package when there is a previous version of the SSH Sentinel software on your computer. The contents - the policies, the rules, the authentication keys etc. - are preserved. Only the software version in updated.

2.9 Removing SSH Sentinel

Before removing the software, you are adviced to do the following:

1. Export and save all such data in the SSH Sentinel that you might need in the future. For example, you might want to save the trusted root certificates for later use. Since removing the software will delete all files in connection to the software, save the data in a separate folder.
2. To be on the safe side, save all unsaved data in other applications and close all open applications.

To remove the software, use the Windows standard procedure: Open `Add/Remove Programs` under `Settings` in the *Start* menu. Select *SSH Sentinel* in the listing. Complete the removal by restarting the computer.

You can re-install the software after completely removing it. Import the saved data to your security policy after installation.

Chapter 3 **Policy Editor**

After you have installed the SSH Sentinel software and rebooted your computer, the software will run on the background and control the incoming and outgoing data traffic based on the rules that you configure. You can set up the rules - the policy - using the SSH Sentinel policy editor. You also use the policy editor to manage the authentication keys of both the remote hosts that you encounter and the local host.

In this chapter, the policy editor is described on a general level. The management of multiple policy layers is covered in Chapter 4 'Managing Multiple Policies'. Chapter 5 'Configuring Policy Rules' and Chapter 6 'Managing Authentication Keys' describe the usage of the policy editor in a more detailed level.

3.1 Policy Management

The security policy acts as a safeguard against harmful attacks from the network. The attacks may result in data losses, secret information being exposed to outsiders and distribution of unwanted, even damaging messages, for example. A successful security policy is able to prevent any malicious attack while still ensuring that the normal, everyday communications run smoothly. A very strict policy, although advantageous from the security point of view, might be considered cumbersome by the user. In worst case, the user might even decide not to apply the policy, in order to be able to communicate over the network. When designing a security policy, both aspect, the security itself and the usability, need to be taken into account.

SSH Sentinel is based on the concept of IPSec policy, which can be broken down to rules that define the policy. The IPSec policy dictates how network IP packages are encrypted, authenticated and compressed and are they passed through or dropped. The defined policy affects all IP communications of the local host. In SSH Sentinel, the user defines the policy rules using a graphical user interface. Administrative user rights are required for changing the IPSec policy - that is, setting up new rules and updating and removing old ones.

The IPSec policy rule set consists of virtual private network, secured connection and secured network rules. Further, packet filtering can be performed both before and after the IPSec transformation is performed on the data packet. In addition, there is the default traffic handling that dictates the actions if no specific rule matches the situation at hand.

Several policies can co-exists the system. However, only one of them can be active, applied, at any given time. The details of handling multiple policy layers are presented in Section Chapter 4 'Managing Multiple Policies'.

3.1.1 Symmetric Security Policy

The target of the conventional IPSec policy management is to build a secure tunnel over an insecure network from a known source to a known destination. It is assumed that both ends apply a similar policy when communicating with each other. The policies can be called *symmetric*. The thinking is justified when building virtual private networks and when a common policy needs to be enforced in an internal network, for example.

The SSH Sentinel software supports working with symmetric policies by supporting centralized policy management. A policy is created and managed in a single point and distributed throughout the network. The key advantages of a centrally managed policy are the ability to establish a common and consistent security policy throughout the network and the ease of configuring and managing multiple IPSec clients and servers from a single location.

3.1.2 Asymmetric Security Policy

While being suitable in some situations like when establishing virtual private networks, the requirement of symmetric policies is a rather inflexible approach when securing everyday end-to-end Internet connections. In the Internet, it is quite unlikely that you can agree on a common policy with each and every end point that you wish to communicate with. Since SSH Sentinel is designed to be a stand-alone software, it pursues a concept of *asymmetric policies*.

In the asymmetric security policy environment, each node in the network is assumed to define its own security policy based on its own needs regardless of the needs and policies of other nodes in the network. Consequently, if the local host requires a certificate based authentication when communicating with, say, host1, then the connection is established only if host-1 submits a certificate that is trusted by the local host. The local host submits its own certificate to host1. Whether host1 is interested in it, is not a concern of the local host. The same goes the other way round, too: The remote host requires certificate based authentication, but the local host would do without authentication. The local host submits its certificate but ignores the remote host's certificate it receives.

Although the opposite is sometimes claimed, two-way authentication is not mandatory in IPSec. Actually one-way authentication is most often sufficient in everyday Internet communications. The SSH Sentinel software supports the above thinking. The existence of symmetric policies is not required.

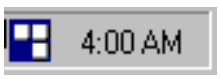

Figure 3.1 The SSH Sentinel system tray icon.

3.2 SSH Sentinel Agent

Once you have installed the software the SSH Sentinel icon appears on the system tray, on the right hand side of the Windows taskbar when the software is running. The icon is shown in Figure Figure 3.1.

If the software is disabled for some reason, the policy rules are not applied on the network data traffic and the icon is dimmed.

Double-clicking the tray icon opens the Sentinel Statistics window. Clicking the right mouse button opens a floating menu with the following items:

View Statistics

Opens the SSH Sentinel statistics display. Double-clicking the tray icon has the same effect. See Section 7.4 'Statistics'.

Run Policy Editor

Open the policy editor.

Auditing

View Audit Log

Open the audit logs to be viewed. For more on audit logs, see Section 7.1 'Auditing'.

View IKE Log Window

Open the IKE Log window that can be used for monitoring traffic and troubleshooting. See Section 7.2 'IKE Log Window' for further reference.

Audit Options

Open the *audit options* property sheet to view and modify the settings. For details, see Section 7.1 'Auditing'.

Select Active Policy

Select the applied policy from the list of policies available on your host.

Enable IPSec

Enable SSH Sentinel: The active policy is applied.

Disable IPSec

Disable SSH Sentinel: No policy is applied.

Help

Open SSH Sentinel online help.

About

Show general information on the SSH Sentinel software.

Online Support

Readme

Open the SSH Sentinel online support pages.

Support Request

Fill in a support request.

Hide Tray

Hide the tray icon. You can make the icon reappear by starting the SSH Sentinel Agent (from the SSH Sentinel folder).

Even if hidden, the tray icon will appear on the screen after each re-boot of your computer. To prevent it from appearing after reboot, remove *SSH Sentinel Agent* from the *Startup* folder under the Windows *Start menu*. Be careful NOT to remove the item from the SSH Sentinel main folder, though! The exact location of the *Startup* folder varies according to the Windows version.

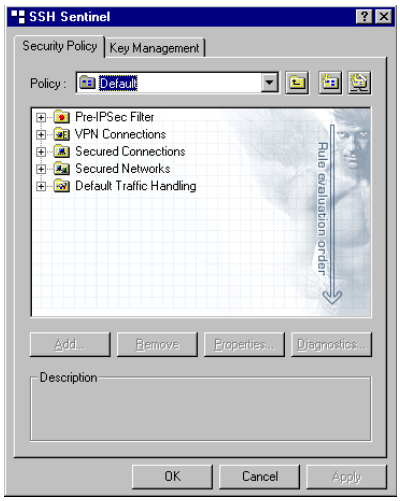

Figure 3.2 The policy editor window.

3.2.1 Starting Policy Editor

To start the policy editor, do one of the following:

1. Click the SSH Sentinel tray icon with the right mouse button. Select `Sentinel policy editor` in the menu that opens.
2. Open the Windows *Control Panel*. Double-click the *SSH Sentinel* icon, or select it with the right mouse button and choose `Open` in the menu that opens.
3. Open the Windows Start menu. Amongst *Programs* you find the *SSH Sentinel* folder. Select *SSH Sentinel Policy Editor* found there.

3.2.2 Policy Editor Window

The policy editor seen in Figure 3.2 is the main window where you can view and manage the policy rules and the authentication keys. There are two sheets, the *Security Policy* and the *Key Management* for management of the rules and the keys, respectively.

Chapter 4 **Managing Multiple Policies**

The SSH Sentinel software supports a multi-layer policy structure which means that a single host can possess an unlimited number of security policy layers. A policy layer is a complete security policy, with packet filtering, IPSec rules, trust policy and possibly a position in a public key infrastructure. Even though several of such layers can exists on a single host simultaneously, naturally only one is active at any given time.

A *local policy layer* is a security policy that is locally created, maintained and stored. The user can modify the policy freely.

A *centrally managed policy layer* is a security policy that is created and stored in a remote location. The centrally managed policy layer is seen as a cached copy on the local host. It cannot be modified on the local host. All changes made to the policy by the centralized management are automatically downloaded to the local host. Since the centrally managed policy is distributed throughout the network, all policy object and the whole policy data are protected by digital signatures. Only a trusted policy certification authority (PCA) is allowed to define and maintain a centrally managed policy layer shared by the local host. See details on policy certification authorities in Chapter 6.1 'Trusted Certificates'.

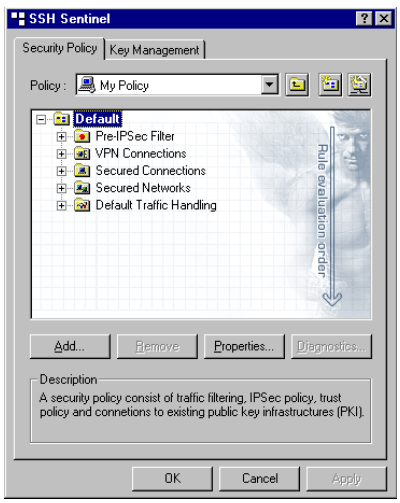

Figure 4.1 Adding a new policy.

4.1 Adding Policies

You can create a new policy from scratch or you can import a configured policy from a file. SSH Sentinel also supports centralized policy management by allowing you to share a policy from a server.

To add a policy, go through the following steps.

1. Set the focus on the header of an existing policy like shown in Figure 4.2 Click the `Add` button.
2. The dialog box called *Add Policy Layer* opens. See Figure 4.2 for reference. In the lower part of the dialog box, specify if you are creating a new local policy from scratch, importing it from a file or sharing a centrally managed policy. Also, give the policy a name in (A). For details, see the sections below.
3. Once ready, click `OK`. To cancel, click `Cancel`.
4. Back on the policy editor, click `Apply` or `OK` to put the new policy into effect. To discard the change, click `Cancel`. **Note!** To pursue the new policy, you need to set it *active*. See Section 4.4.1 'Active Policy' for details.

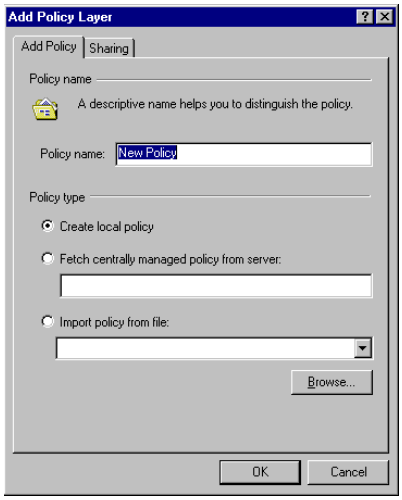

Figure 4.2 The dialog box for adding a new policy layer.

4.1.1 Creating Local Policies

Creating a new local policy means configuring a new policy from scratch. You can freely update the rule set and share the policy with other users if you wish.

To create a new policy from scratch, select the option *Create local policy* on the *Add Policy Layer* dialog. You can make the new policy *sharable* on the *Sharing* sheet of the same dialog.

4.1.2 Importing Policies

A policy that you import from a file, becomes a normal local policy once you have added it. You can freely update the rules and you can share the policy to other users.

To import a policy from a file, select the option *Import Policy from File* on the *Add Policy Layer* dialog. Write the name and path of the file in (B). You can also locate the file by clicking the `Browse` button and navigating in the file system.

You can make the new policy *sharable* on the *Sharing* sheet of the *Add Policy Layer* dialog.

4.1.3 Sharing Policies

To support centralized policy management, SSH Sentinel offers an easy way to share a policy from a server. A shared, centrally managed policy is not locally updateable.

To share a policy, select the option *Add Centralized Policy Layer From Server* on the *Add Policy Layer* dialog. Specify the server in the appropriate field.

4.2 Removing Policies

To delete a policy, select it on the policy editor and click the `Remove` button. To make the removal permanent, click `Apply` or `OK`. To cancel the removal, click `Cancel`.

4.3 Viewing Policy Properties

The properties of a policy can be viewed on the *Properties* dialog box. There are two sheets, the *General* and *Sharing*.

To open the dialog box, select the policy you want to view, and do one of the following:

1. Click the `Properties` button.
2. Click the right mouse button, and select `Properties` in the menu that opens.
3. Click the right mouse button, and select `Sharing` in the menu that opens. The `Properties` dialog opens with the *Sharing* sheet visible.

4.4 Policy Properties

Type

The type of a policy is either local or shared.

Created

The date and time of the creation of the policy.

Modified

The date and time of the last update on the policy.

Active Policy

Checked, if this is the executed policy. See Section 4.4.1 'Active Policy'.

Not Shared/Shared

(On the *Sharing* sheet.) If you allow others to share your local policy, the tag *Shared* is checked. If not, *Not Shared* is checked.

4.4.1 Active Policy

Out of the different policy layers in your system, the *active* policy is executed. The active policy has the property tag *Active Policy* set on. Also, the active policy is shown in bold letters in the policy editor. Only one policy can be active at a time.

To change the active policy, do one of the following:

1. From the policy editor: Select the policy you want to activate with the right mouse button. Select `Set as Active Policy` in the menu that opens. The menu item is inactive, if the policy you are handling is the active policy or if you have not commited your changes after adding a new policy.
2. From SSH Sentinel tray icon: Open the tray icon menu with the right mouse button. Select the menu item `Select Active Policy` with the left mouse button. A list of policies is shown. Select the correct policy.

Chapter 5 Configuring Policy Rules

The security policy can be seen and modified on the *Security Policy* sheet of the policy editor. You cannot modify a centrally managed policy that you share from a server.

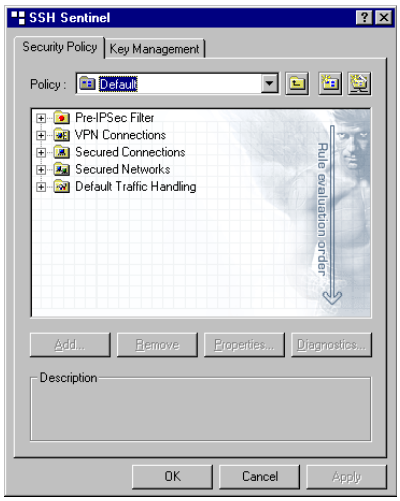

Figure 5.1 The security policy sheet on the policy editor.

5.1 Rule Evaluation

The security policy consist of IPSec connection rules (virtual private network connections, secured connections and secured networks), traffic filtering and default response rules. The rules are evaluated in the order they appear in the user interface.

Unprotected IP traffic is used here as the term for "non-IPSec" data packets. Since SSH Sentinel only supports ESP out of the IPSec protocols - and not AH - the unprotected IP traffic would mean all other IP traffic but ESP.

5.1.1 Traffic Filters

The traffic filter consist of simple rules that stop or pass a data packet. The decision is based on pre-defined selectors that include the traffic protocol as well as the IP address and the port of both the local and remote end. Also, the traffic direction can be specified; a rule can only match the incoming or outgoing packets - from the local host point of view - or both directions.

The filter rule specifies, what to do with the data packet that matches the criteria defined by the selectors. If passed by, the data packet is not affected by any other rules, including remaining filter rules, the IPSec rules or default response rules. If the data packet is dropped or rejected, however, it is not let through. Dropping just discards the packet, rejecting means sending a standard-form message of denial to the initiator while still discarding the data packet.

The traffic filters cannot trap data packets with IPSec encapsulation. To be able to filter IPSec transmissions in all situations, two separate filters are available: the pre-IPSec filter to be applied before the IPSec transformation is performed and the post-IPSec filter to be applied after the transformation. The existence two separate filters can undeniably cause confusion and consequently the filter rules need to be configured with care to avoid ambiguity. Often, the filtering of the outbound traffic is performed with pre-IPSec filter and the filtering of the inbound traffic with post-IPSec filter. However, both filters are applied on all traffic (unless, of course, a particular data packet is, for example, bypassed by a filter rule).

5.1.2 Default Response Rules

A default response rule is evaluated if none of the rules - traffic filter rules or IPSec rules - match the data packet. There is naturally separate rules for IPSec traffic and unprotected IP traffic.

Suppose a situation where a remote host tries to establish an IPSec connection with the local host. None of the rules matches the connection proposed. Thus, the default response rule is applied. It contains the default identity - a certificate or a pre-shared key - which to use to authenticate the local host. Also, the local host can be configured to by default trust all certificates received from remote hosts when applying the default response rule.

If the IPSec connection can be negotiated, then the security associations are set up. The security associations have a lifetime. Within this lifetime, IPSec traffic from the remote host are accepted.

The default response for unprotected IP traffic basically denies or bypasses all such traffic that was not explicitly denied or accepted by the filter rules.

5.1.3 Controlling Incoming Unprotected IP Traffic

When a remote host sends an unprotected IP packet to the local host, the packet is first filtered by the rules in the pre-IPSec filter. If there is a matching bypass rule, the data packet skips all other rules and is forwarded to further processing in the IP stack. If there is a matching drop or reject rule the data packet is discarded - and if it was rejected a message of denial is delivered to the

remote end. However, it may well be that none of the rules match. Then, the data packet is forwarded to the post-IPSec filter.

In the post-IPSec filter, processing similar to the pre-IPSec filter, occurs. The packet can again be bypassed, meaning that it is accepted and forwarded to the IP stack, or rejected or dropped. Again, it might well be that none of the rules matches. To handle these packets that the system actually does not know how to handle, there is the default response rule. Basically, you can either pass all such traffic by or drop it all. If you choose to drop, you must be careful enough to explicitly pass all acceptable traffic by. However, if you choose to bypass, you must create explicit drop rules in your filters to discard all potentially harmful traffic.

5.1.4 Controlling Incoming IPSec Traffic

The incoming IPSec traffic passes the pre-IPSec filter unaffected since the packet is still IPSec-encapsulated. Next, the IPSec rules are evaluated, first the virtual private network rules, then the secured connection rules and last the secured network rules. If a rule matches, then the data packet is passed to the IPSec de-capsulation and to the IP stack.

If none of the rules matches, then the default response rule is evaluated, unless a matching security association, as a result of a previous application of the default response rule, exists. The data packet is then decapsulated and passed on to the post-IPSec filter. It might actually happen, that even if the IPSec negotiations with the default response rule were successful, the data packets later delivered are denied by a rule in the post IPSec filter.

5.1.5 Controlling Outgoing Traffic

An outgoing data packet enters first the pre-IPSec filter. The packet might be bypassed and thus let through unprotected. It is forwarded to the IP stack and no other rules will affect the packet. On the other hand, the packet might be dropped or rejected. In the latter case, a message of denying the packet is delivered. If none of the rules in the pre-IPSec filter matches, next the IPSec rules are evaluated.

If one of the IPSec rules - a virtual private network rule, a secured connection rule or a secured network rule - matches, then an IPSec connection is negotiated (if the security associations do not already exist), IPSec transformation is performed on the data packet and the resulting IPSec packet is forwarded to the IP stack.

If, however, none of the IPSec rules matches, the data packet is next filtered by the post-IPSec filter rules. Here again, the data packet can be bypassed, dropped or rejected, or no matching rule can be found. If so, the data packet is forwarded to the IP stack.

A word of caution is here appropriate: Suppose you have created a secured connection rule to a particular remote host. In other words, you want all communications between your local host and the remote host to be IPSec protected. However, you have accidentally created a general bypass rule in the pre-IPSec filter to bypass all traffic from the remote host in question. Now, the outgoing IP packets are bypassed and the secured connection rules are never even evaluated. As a result,

you assume that the communications to the remote host is IPSec protected even if it is not. Naturally, this same mechanism is exploited to pass, for example, IKE negotiations which are needed to establish the IPSec sessions. But, all in all, the importance of careful configuration of the filter rules is, once again, emphasized.

5.2 Traffic Filters

The SSH Sentinel policy management includes tools for comprehensive IP packet filtering. By filtering IP packets, you can discard connections from and to potentially hostile end points and to filter out unwanted traffic. The traffic filtering rules allow the system administrator to limit access to open services and ports and to define arbitrary allow and drop rules in general. The filtering rules complement the protection provided by IPSec.

The traffic filtering in the SSH Sentinel software is broken down to filtering taking place before and after the IPSec transformation on the data packet transmitted. As far as the incoming traffic is concerned, the pre-filter is applied before the IPSec encapsulation is removed and the post-filter after removing it. Following the same logic, the pre-filter is applied on the outgoing traffic before the IPSec encapsulation is inserted and the post-filter after inserting it.

This kind of a setup is required to be able to filter IPSec encrypted traffic. Since IPSec encrypted traffic is of ESP protocol, it obviously is not affected by filter rules that handle TCP and UDP traffic. Consequently, the generic filtering of the outbound traffic is often performed with the pre-IPSec filter and of the inbound traffic with post-IPSec filtering. However, the post-IPSec filtering is executed on all traffic, no matter if it is IPSec encrypted or not.

Traffic filtering may be performed according to the selectors listed below:

Direction

Inbound, outbound or both. The IP packets may be filtered separately depending on the traffic direction: coming in from the network or going out to the network.

Local Address

The local host. In outbound traffic, this is the IP address of the source, in inbound traffic, the IP address of the destination of the traffic.

Remote Address

The remote host. In inbound traffic, this is the IP address of the source, in outbound traffic, the IP address of the destination.

Local Port

The port used in the local host. In outbound traffic, this is the source port, in inbound traffic, the destination port.

Remote Port

The port used in the remote host. In inbound traffic, this is the source port, in outbound traffic, the destination port.

Protocol

The traffic protocol: TCP, UDP, ICMP or *all* to include every protocol.

Since the user interface and the management of both filters is almost identical, they are explained here together. The examples and pictures are taken from the pre-IPSec filter. The differences are highlighted.

5.2.1 Listing and Viewing Filter Rules

Expand the `Pre-IPSec Filter` branch in the policy tree by clicking the plus sign to the left from it, and you see a listing like in the Figure 5.2, on the rules configured. A few fundamental properties are shown. Select a rule and you can see the description of it on the bottom of the policy editor window. The description is user updateable. You get a more complete listing by double-clicking `Pre-IPSec Filter`.

To view the details of a particular rule, open the Filter Rule Properties window: Select the rule you want to view in detail and click the `Properties` button **OR** double-click the rule you want to investigate.

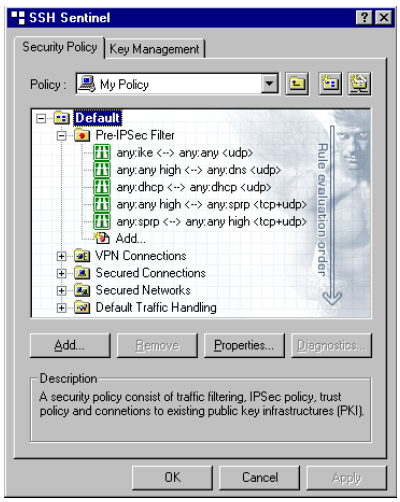

Figure 5.2 The basic listing of the pre-IPSec filter rules.

5.2.2 Modifying Evaluation Order

The filter rules are listed in the evaluation order, the order being from top to bottom. The first rule that matches the connection will be applied and no more rules evaluated. Thus, if you want to reject ftp connections in general, but to accept it from one particular host, called host-1, you create two rules: one to bypass ftp from host-1 and another to reject it from any host. The more specific bypass rule must be evaluated before the general rejection rule.

To modify the evaluation order, do the following:

1. Select the rule you want to reposition. Two arrow buttons appear now below the policy tree. See Figure Figure 5.3.
2. Move the rule upwards or downwards with the arrows.
3. Once ready, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

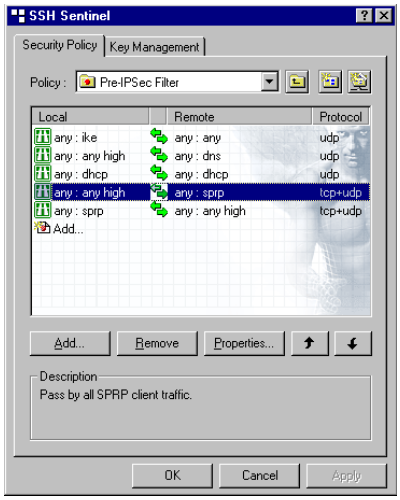

Figure 5.3 The detailed listing of filter rules. When a rule is selected, the arrows that move the rule up and down in the listing appear.

5.2.3 Adding Filter Rules

To add a new filter rule, go through the following steps:

1. Open the Filter Rule Properties window: Select `Pre-IPSec Filter` in the policy tree. Click the `Add` button. **OR** Select `Pre-IPSec Filter` in the policy tree with the right mouse button. In the menu that opens, select `Add New Rule`. **OR** Double-click the `Add` branch hanging under *Traffic Filter* in the policy tree.

2. Fill in the appropriate values. See Section 5.2.8 'Filter Rule Properties' for assistance.
3. Accept the changes by clicking `OK`. To not to add the rule after all, click `Cancel` and your changes will disappear.
4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

5.2.4 Removing Filter Rules

To remove a filter rule, select it on the policy editor and click the `Remove` button. To make the removal permanent, click either `OK` or `Apply`. `OK` also closes the policy editor. You can restore the rule set after even several removals by clicking `Cancel`, provided that you haven't yet committed the changes with `Apply` or `OK`.

5.2.5 Modifying Filter Rules

To modify a filter rule, go through the following steps:

1. Open the Filter Rule Properties window: Either select the rule you want to modify and click the `Properties` button **OR** double-click the rule you want to modify.
2. On the Filter Rule Properties window, edit the values. See Section 5.2.8 'Filter Rule Properties' for assistance.
3. Accept the changes by clicking `OK`. To not to add the rule after all, click `Cancel` and your changes will disappear. You return to the policy editor.
4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

5.2.6 Enabling and Disabling Filter Rules

You can disable a single filter rule - and naturally later enable it again. To disable (enable), select the rule with the right mouse button and click `Rule Enabled` in the menu that opens. A check mark appears next to the menu item text when the rule is enabled. When disabled, a little x-mark appears on the rule icon. Remember to commit you changes by clicking the `Apply` button.

5.2.7 Auditing Filter Rules

To audit a rule, do one of the following:

1. Select the rule with the right mouse button. Click `Audit Rule` in the menu that opens. A check mark appears to indicate that the rule is now audited.
2. Select the rule and open the *Connection Properties* window. Click the *Advanced* tab to view the audit options. Select the check box with label *Audit this rule*. Click the `OK` button to return.

Remember to commit your changes by clicking the `Apply` button.

To stop auditing a rule, do the opposite: Select the rule with the right mouse button and click `Audit Rule` again. The check mark now disappears. **OR** On the *Advanced* sheet of the *Properties* window, clear the check box *Audit this rule*. Remember to commit your changes.

To learn the details of auditing, see Section 7.1 'Auditing'.

5.2.8 Filter Rule Properties

The rule properties are shown on the *Filter Rule Properties* dialog box shown in Figure 5.4 and Figure 5.5.

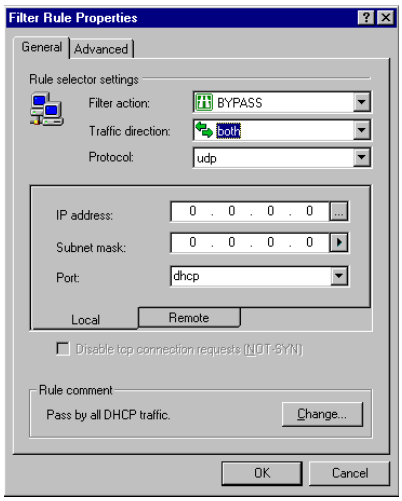

Figure 5.4 The filter rule properties.

Rule Selector Settings

Filter Action

There are three options on what to do with the incoming or outgoing data packet: pass it by, drop it or reject it. *Bypass* means, naturally, that the data packet is allowed to leave or enter as it is. (When configuring post-IPSec filter rules, the term *allow* is used.) *Dropping* traffic means that the data packet is not let through but the sender of the packet is not

informed on dropping. *Rejecting* the traffic also stops the data packet, but in this case, the sender is informed that the transmission was rejected.

Traffic Direction

A rule can only affect the *inbound* traffic, that is the data packet coming from an external origin into the local host or the *outbound* traffic, data packets from the local host to an external target. The third option is, that the rule affects both directions, *inbound and outbound*.

Protocol

The upper level (transport layer) protocol. The available options are listed in the drop down list: *udp*, *tcp*, *icmp* and *any*. Selecting *any* means that the rule matches traffic of any protocol.

Local and Remote Ends

You specify the local and remote ends by describing the host or network in question and stating the port through which the traffic flows in each end. You can shift between specifying the local end and the remote end by clicking the `Local` and `Remote` tabs (A) in Figure 5.2.8 'Filter Rule Properties'. By clicking the arrow button (B), you can select how to describe the host or network: by IP address, by IP network or by IP address range. The descriptions below are on these menu items.

Before proceeding, consider for a moment the fact that you are able to specify the local end exactly with IP address. This feature is significant, if your local host has several network cards attached to it.

IP Address

The local / remote end is a single host identified by its IP address. You can look up the IP address up by clicking the button (C) to the right of the text box and typing the host DNS name in the text box that appears.

IP Network

The network IP address and the subnet mask. You can look the gateway IP address up by clicking the button (C) to the right of the text box and typing the host DNS name in the text box that appears.

IP Address Range

The IP address range occupied by the network.

Port

The port used by the traffic. Since each service (application) has a distinct port associated with it, you can reject one service while allowing others to enter by specifying the port. There are a few alternative ways to specify the port. You can select the service from the

drop down list that opens when you click the arrow button (D). The system automatically fills in the corresponding port. Also, you can simply type in the port number. To specify multiple ports, number ranges are also acceptable.

Rule Comment

Free-form text that appears on the bottom of the policy editor window when you have the rule selected. You can edit the text by clicking the `Change` button and typing the new description in the text field that appears.

Audit Options

The audit options are seen on the *Advanced* sheet of the *Filter Rule Properties* window in Figure 5.5.

Audit this rule

To audit the rule, select this check box. For more information on auditing, see Section 7.1 'Auditing'.

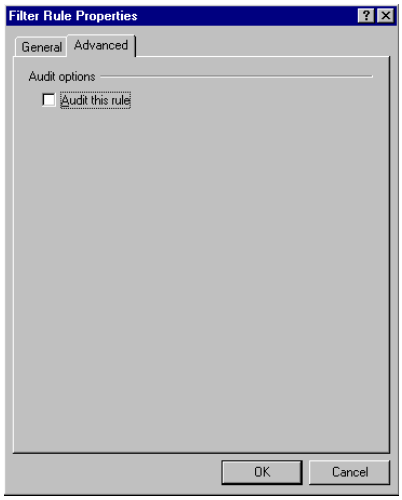

Figure 5.5 The advanced properties.

5.3 Virtual Private Network Connections

A virtual private network is a setup where a private network is made accessible to authorized users connecting to it from a remote location. The private network is protected by a security gateway. All traffic in to and out of the network traverses the gateway. Between the security gateway and the

user host, the connection is established over the public Internet. The virtual private network connections are typically used for enabling remote access to a network and for remote working.

A virtual private network connection is a peer-to-security-gateway tunnel mode IPSec connection. The data transmitted is encrypted and the communicating parties are authenticated. An example of a virtual private network setup is shown in Figure 5.6.

Figure 5.6 A typical virtual private network layout.

5.3.1 Adding VPN Connection Rules

In order to insert a new virtual private network connection rule, go through the following steps:

1. Initiate the adding of a new rule: Select *VPN Connections* in the policy tree and click the `Add` button. **OR** Select *VPN Connections* with the right mouse button and click `Add New Rule` in the menu that opens.
2. The dialog box shown in Figure 5.7 opens. Type the name or IP address of the remote security gateway in the field provided. You can switch from host name to IP address and back with the button (E) to the right of the edit control. Next, type in the network IP address and the subnet mask the remote network. Alternatively, you may click to specify the IP address range in question. To do this, open the menu with the arrow button and click `IP Address range`. To specify the IP address and subnet mask, click `IP Network` in the same menu. Select the authentication key that you wish to use to identify the local host, using the list box. If there is a certificate on smart card available, the list box will show the certificate of the certification authority found on the smart card. If you select it, the actual authentication key used will be the end entity certificate issued by the certification authority. For details on managing the authentication keys and smart cards, see Chapter 6 'Managing Authentication Keys'. If you wish to use the short form of the proposal, select the *Use legacy proposal* check box. For more on proposal types, see the discussion under Section 5.3.7 'VPN Connection Rule Properties'. When ready, click the `OK` button.

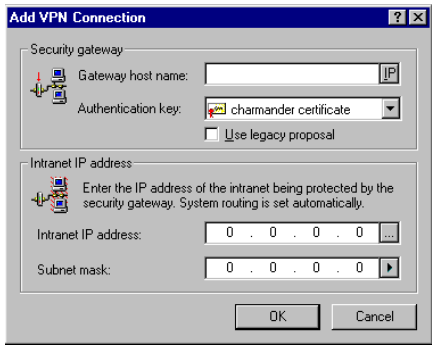

Figure 5.7 Adding a rule on a virtual private network connection.

3. The system then goes on to probe the connection parameters in order to establish the connection parameters and to test the functionality of the connection. The dialog box is shown in Figure 5.8. Once ready, the system informs you, if the probing was successful. The dialog box shown after having successfully probed the parameters, can be seen in Figure 5.9. To view the connection parameters established, click the `Details` button. If the probing fails, you can still add the rule and later return to edit the parameters manually, see Section 5.3.3 'Viewing and Editing VPN Connection Rules'.

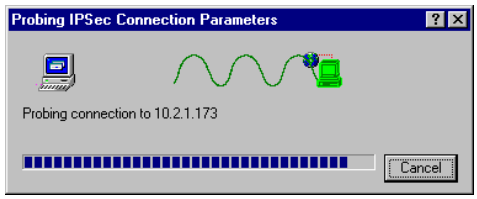

Figure 5.8 Probing the connection parameters.

4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

5.3.2 Removing VPN Connection Rules

To remove a rule, select it on the policy editor and click the `Remove` button. To make the removal permanent, click either `OK` or `Apply`. `OK` also closes the policy editor. You can restore the rule set after even several removals by clicking `Cancel`, provided that you haven't yet committed the changes with `Apply` or `OK`.

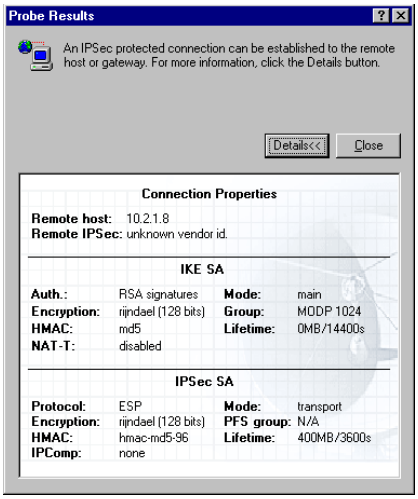

Figure 5.9 The probe results.

5.3.3 Viewing and Editing VPN Connection Rules

Once you have created an IPSec rule on a virtual private network connection, you can view the rule parameters and modify them at any time. You might want to change the rule parameters after having unsuccessfully probed the parameters or because, simply, you know, that some setting has changed.

To see the basic listing of the configured virtual private network connection rules, like seen in Figure 5.10, expand the `VPN Connections` branch. To get a more complete listing, double-click `VPN Connections`. A new view, like in Figure 5.11, containing the listing with a few more details on each rule, opens. To view the details of a rule, select the rule you want to view, and click the `Properties` button **OR** double-click the rule you want to view in detail. The *Connection Properties* dialog box opens.

To edit a rule, do the following:

1. Open the *Connection Properties* dialog box: Select the rule you want to edit, and click the `Properties` button. **OR** Double-click the rule you want to edit.
2. The *Connection Properties* dialog box appears with the *General* sheet visible. You can shift to the *Advanced* and *SA Lifetimes* sheets by clicking the respective tabs. Edit the values. See Section 5.3.7 'VPN Connection Rule Properties' for assistance.
3. When ready, accept the changes by clicking `OK`. To discard your changes, click `Cancel`. Both buttons will take you back on the policy editor.
4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

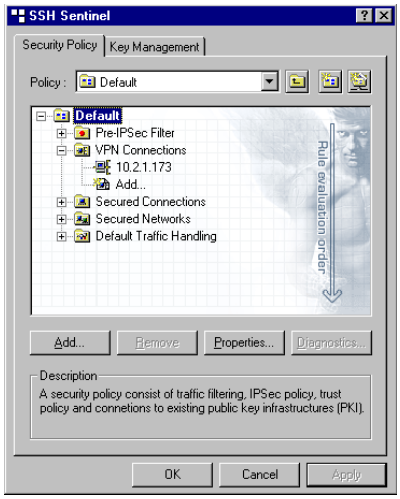

Figure 5.10 The basic listing of the virtual private network connection rules.

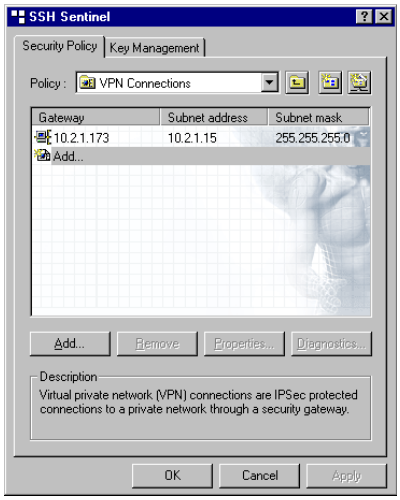

Figure 5.11 The detailed listing of the virtual private network connection rules.

5.3.4 VPN Connection Diagnostics

To test a virtual private network connection, select the rule you have created and do one of the following:

1. Click the `Diagnostics` button.
2. Click `Diagnostics` in the menu that opens with the right mouse button.

While the system is probing the connection the dialog box seen in Figure 5.8 is shown. The connection is negotiated and the connection parameters agreed on with the other end are established. The probing may change the parameters that you have initially set on the rule. The parameters you set are used as initial proposals, however, if the other does not support, say, you're first choice of encryption algorithm, then some other algorithm is chosen.

You are clearly informed if the probing was successful. If it is, then you are able to establish a connection with the other end. The dialog box shown after having successfully probed the parameters is shown in Figure 5.9.

5.3.5 Enabling and Disabling VPN Connection Rules

You can disable a single rule - and naturally later enable it again. To disable (enable), select the rule with the right mouse button and click `Rule Enabled` in the menu that opens. A check mark appears next to the menu item text when the rule is enabled. When disabled, a little x-mark appears on the rule icon. Remember to commit you changes by clicking the `Apply` button.

5.3.6 Auditing VPN Connection Rules

To audit a rule, do one of the following:

1. Select the rule with the right mouse button. Click `Audit Rule` in the menu that opens. A check mark appears to indicate that the rule is now audited.
2. Select the rule and open the *Connection Properties* window. Click the *Advanced* tab to view the audit options. Select the check box with label *Audit this rule*. Click the `OK` button to return.

Remember to commit your changes by clicking the `Apply` button.

To stop auditing a rule, do the opposite: Select the rule with the right mouse button and click `Audit Rule` again. The check mark now disappears. **OR** On the *Advanced* sheet of the *Properties* window, clear the check box *Audit this rule*. Remember to commit your changes.

To learn the details of auditing, see Section 7.1 'Auditing'.

5.3.7 VPN Connection Rule Properties

The rule properties can be seen on the *Properties* dialog box, seen in Figure 5.12. There are three sheets, the *General*, *SA Lifetimes* and *Advanced*. The latter two sheets can be seen in Figure 5.13 and Figure 5.14.

IP Address settings

You describe the private network you wish to access by specifying the gateway host and the private network. You can specify the network segment in two ways. You either specify the network IP address and the subnet mask or you specify the IP address range occupied by the network segment in question. To specify one or the other, click the arrow button (A) and select `IP Network` or `IP Address Range` in the menu that opens. The latter two of the descriptions below are on these menu items:

Gateway host

The DNS name or the IP address of the security gateway protecting the network segment you want to access. You can shift from host name to IP address by clicking the button (B) to the right of the text box.

IP Network

The network IP address and the subnet mask. You can look the gateway IP address up by clicking the button (C) to the right of the text box and typing the host DNS name in the text box that appears.

IP Address Range

The IP address range occupied by the network.

Proposal Parameters

Authentication key

You can select the authentication key used to authenticate your local host in this connection from the list of available keys. The available keys include both the certificates and the pre-shared keys. If there is a smart card reader attached to your computer and a smart card in the reader, also the certificates on the smart card are available. However, the listing shows the certificate of the certification authority not the actual end entity certificate. When the connection is negotiated, the smart card is read again and the end entity certificate issued by the certification authority in the listing is actually included in the proposal. Consequently, in a large corporation, an administrator needs not to configure each instance of the software individually since the rule is not bound to a certain smart card and a certificate found on it. For more on management of the smart cards and authentication keys in general, see Chapter 6 'Managing Authentication Keys'.

Proposal type

You can select the type of the proposal that your host is going to make when negotiating a connection with another host. There are two options: *normal* and *legacy*. A normal proposal contains all the possible parameter values supported by SSH Sentinel thus making the proposal very long. Some security software have difficulties with such long proposals. The legacy proposal is a fallback option to be used in such situations. The main differences in the two types of proposals are highlighted in the table below.

Encryption

The encryption algorithm. Select the one you want to use from the list of supported algorithms. Naturally, the algorithm can be applied only if both ends support it. SSH Sentinel uses your choice as its initial proposal. If the other end does not support that particular algorithm, the hosts will negotiate and find a mutually supported one. The following are supported by the SSH Sentinel Software: Rijndael, Twofish, Blowfish, Cast, 3DES and DES.

IPSec mode

A VPN connection is inherently a *tunnel* mode connection. Consequently, you can not change this setting.

IKE mode

The Internet Key Exchange (IKE) mode setting has to do with negotiating the connection. Each time you wish to create a connection, the system contacts the other end in order to establish the connection parameters or to check that they are correct and the connection works. There are two modes in the negotiations: *main* and *aggressive*. The latter is faster, whereas the former is more flexible and thus more likely to succeed. You can change the setting of the IKE mode by selecting the correct mode from the list.

IKE group

When negotiating a connection, the communicating parties also settle the actual keys used to encrypt the data. The keys are based on the private keys of each party and some random data. The random data generation is based on so-called pool-bits. The IKE group basically tells the number of pool bits. The more pool bits, the larger numbers in question. The larger the number, the harder to break. Consequently, more pool bits means more secure.

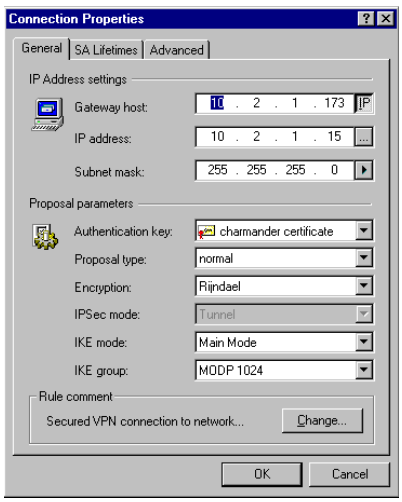

Figure 5.12 The general properties.

Rule comment

Free-form text that appears on the bottom of the policy editor window when you have the rule selected. You can edit the text by clicking the `Change` button and typing the new description in the text field that appears.

Security Association Lifetimes

You can use the four sliders on the *SA Lifetimes* sheet seen inFigure 5.13 to control the lifetimes of the established security associations. There are distinct but similar controls for IKE and IPSec security associations. You can control the lifetime of either type of security association by time (minutes) or transferred data (megabytes). Whenever either limit is fulfilled, the existing security association is destroyed and a new, eventually, established.

Use the `Defaults` button to return the original values.

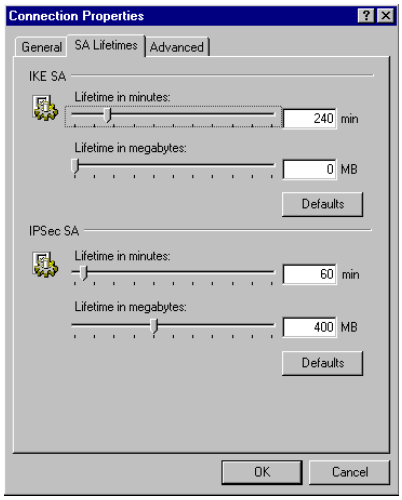

Figure 5.13 The security association lifetime settings.

Proposal Parameter	Normal Proposal	Legacy Proposal
Encryption Algorithm	Rijndael, Twofish, Blowfish, Cast, 3DES, DES	3DES, DES
IPSec Mode	Tunnel	Tunnel
IKE Mode	Main, aggressive	Main, aggressive
IKE Group	MODP 768, MODP 1024 MODP 1536	MODP 1024

Audit Options

Audit this rule

To audit the rule, select this check box. For more information on auditing, see Section 7.1 'Auditing'.

Advanced Options

Apply IP compression

Each data packet transmitted is compressed. Consequently, transmitting it is faster. The compression is performed only if the other end also supports IP compression.

Auto-discover path maximum transfer unit (PMTU)

The system will find out the maximum transfer unit (MTU) size in the connection. This is done to avoid data fragmentation. In other words, the system finds out how big a data packets it can transmit along the connection. Then, it makes sure to send maximum size data packets. Consequently, the data is distributed to a minimal number of packets and transmission is faster and less prone to errors.

Use perfect forward secrecy (PFS) in IKE rekey

When negotiating the actual keys again, the new key will in no way be dependent on the old key. This is, naturally, more secure, because even if somebody found out the old key, it does not reveal anything on the new key. On the other hand, the perfect forward secrecy slows down the key negotiations.

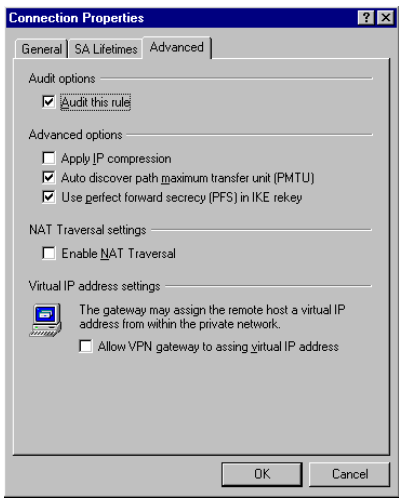

Figure 5.14 The advanced properties.

NAT Traversal Settings

Enable NAT Traversal

Network Address Translation (NAT) is a method by which IP addresses are mapped from one network realm to another. A situation like this occurs, when a private network, with private addresses, is connected to the Internet. Traditional virtual private network solutions cannot work, however, with network address translation due to security problems. NAT Traversal is a solution to this.

Virtual IP Address Settings

Allow VPN gateway to assign virtual IP address

When selected, the local host accepts a virtual IP address assigned by the security gateway to the local host. The assigned virtual IP address belongs to the internal private network. It is used exclusively when communicating with the private network to create an illusion of a seamless connection to the network.

5.4 Secured Connections

Secured connections are simple peer-to-peer IPSec connections. Such important IP connections as remote administration of computers, file transfer or download, sending and receiving email, web browsing and IP telephony are commonly protected by IPSec.

If you establish a secured connection with a remote host, the data transmitted is encrypted. Also, both ends are authenticated.

Figure 5.15 The peer-to-peer connection layout.

5.4.1 Adding Secured Connection Rules

In order to insert a new secured connection rule, go through the following steps:

1. Initiate the adding of a new rule: Select *Secured Connections* in the policy tree and click the `Add` button. **OR** Select *Secured Connections* with the right mouse button and click `Add New Rule` in the menu that opens.
2. The dialog box shown in Figure 5.16 opens. Type in the name or IP address of the remote host (A). You can switch between the host name and the IP address with the button (C). Next, select the authentication key that you wish to use to identify the local host, using the list box (B). If there is a certificate on smart card available, the list box will show the certificate of the certification authority found on the smart card. If you select it, the actual authentication key used will be the end entity certificate issued by the certification authority. For details on managing the authentication keys and smart cards, see Chapter 6 'Managing Authentication Keys'. When ready, click the `OK` button. *Tip: To use the legacy proposal, e.g. the short form of the proposal, add the rule and set the connection properties properly. See Section 5.4.7 'Secured Connection Rule Properties' for reference.*

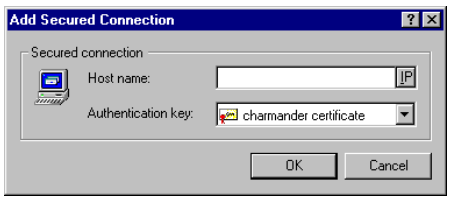

Figure 5.16 Adding a new peer-to-peer secured connection.

3. The system then goes on to probe the connection parameters and to test the connection. The dialog box that appears is shown in Figure 5.17. Once ready, the system informs you if the probing was successful. The dialog box shown after successfully probing the parameters, can be seen in Figure 5.19. To view the connection parameters established, click the `Details` button. If the probing fails, you can still add the rule and later edit the parameters manually. See Section 5.4.3 'Viewing and Editing Secured Connection Rules'.

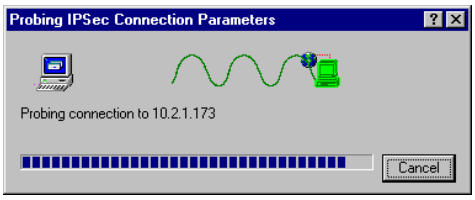

Figure 5.17 Probing the connection parameters.

4. If certificates are used to authenticate the communicating parties, then one of the results of the probing is the remote end certificate. In some cases, the certificate can automatically be trusted by the system. It can be a certificate that you have received in some earlier occasion and classified as a trusted certificate or it might be a certificate issued by a certification authority that you trust. See Chapter 6 'Managing Authentication Keys' for details on

managing certificates. However, if the remote end certificate cannot be automatically trusted, the system will ask you to make the decision, see Figure 5.18.

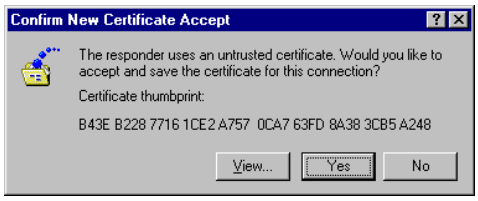

Figure 5.18 Accepting a remote end certificate.

5. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

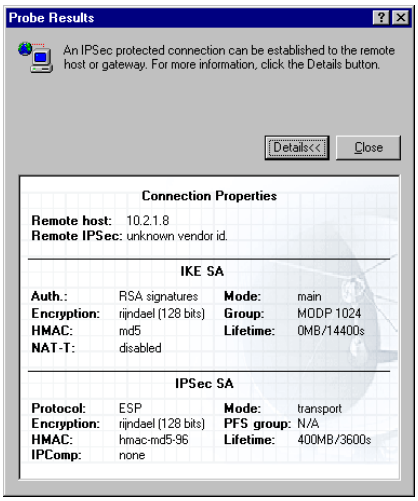

Figure 5.19 The probe results.

5.4.2 Removing Secured Connection Rules

To remove a rule, select it on the policy editor and click the `Remove` button. To make the removal permanent, click either `OK` or `Apply`. `OK` also closes the policy editor. You can restore the rule set after even several removals by clicking `Cancel`, provided that you have not yet committed the changes with `Apply` or `OK`.

5.4.3 Viewing and Editing Secured Connection Rules

Once you have created an IPSec rule on a peer-to-peer connection, you can view the connection parameters and modify them at any time. You might want to change the connection parameters after having unsuccessfully probed the parameters or because, simply, you know that some setting has changed.

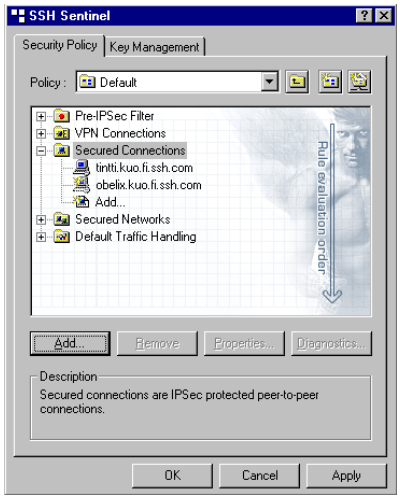

Figure 5.20 The basic listing of the secured connection rules.

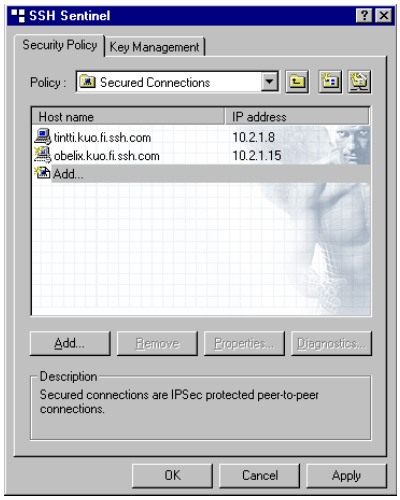

Figure 5.21 The detailed listing of the secured connection rules.

To see the basic listing of the configured secured connection rules, seen in Figure 5.20, expand the `Secured Connections` branch. To get a more complete listing, seen in Figure 5.21, double-click `Secured Connections`. A new view containing the listing with a few more details on each rule opens. To view the details of a rule, select the rule you want to view, and click the `Proper-`

ties button or double-click the rule you want to view in detail. The *Connection Properties* dialog box opens.

To edit a rule, do the following:

1. Open the Connection Properties dialog box: Select the rule you want to edit, and click the Properties button **OR** double-click the rule you want to edit.
2. The *Connection Properties* dialog box appears with the *General* sheet visible. You can shift to the *Advanced* and *SA Lifetimes* sheets by clicking the respective tab. Edit the values. See Section 5.4.7 'Secured Connection Rule Properties' for assistance.
3. When ready, accept the changes by clicking OK. To discard your changes, click Cancel. Both take you back on the policy editor window.
4. Back on the policy editor, click Accept or OK to accept the changes and put them into effect. OK will also close the policy editor window. To discard the changes, click Cancel. **Note!** These actions affect all modifications in the rule set and the key management. Thus Apply and OK will commit and Cancel discard all changes made so far.

5.4.4 Secured Connection Diagnostics

To test a secured connection, select the rule you have created and do one of the following:

1. Click the Diagnostics button.
2. Select Diagnostics in the menu that opens with the right mouse button.

While the system is probing the connection the dialog box seen in Figure 5.17. The connection is negotiated and the connection parameters agreed on with the other end are established. The probing may change the parameters that you have initially set on the rule. The parameters you set are used as initial proposals, however, if the other does not support, say, you're first choice of encryption algorithm, then some other algorithm is chosen.

You are clearly informed if the probing was successful. If it is, then you are able to establish a connection with the other end. The dialog box shown after having successfully probed the parameters is shown in Figure 5.4.1 'Adding Secured Connection Rules'.

5.4.5 Enabling and Disabling Secured Connection Rules

You can disable a single rule - and naturally later enable it again. To disable (enable), select the rule with the right mouse button and click Rule Enabled in the menu that opens. A check mark appears next to the menu item text when the rule is enabled. When disabled, a little x-mark appears on the rule icon. Remember to commit you changes by clicking the Apply button.

5.4.6 Auditing Secured Connection Rules

To audit a rule, do one of the following:

1. Select the rule with the right mouse button. Click `Audit Rule` in the menu that opens. A check mark appears to indicate that the rule is now audited.
2. Select the rule and open the *Connection Properties* window. Click the *Advanced* tab to view the audit options. Select the check box with label *Audit this rule*. Click the `OK` button to return.

Remember to commit your changes by clicking the `Apply` button.

To stop auditing a rule, do the opposite: Select the rule with the right mouse button and click `Audit Rule` again. The check mark now disappears. **OR** On the *Advanced* sheet of the *Properties* window, clear the check box *Audit this rule*. Remember to commit your changes.

To learn the details of auditing, see Section 7.1 'Auditing'.

5.4.7 Secured Connection Rule Properties

The properties of a rule can be seen on the *Connection Properties* dialog box, seen in Figure 5.22. There are three sheets, the *General*, *SA Lifetimes* and *Advanced*. The latter two sheets can be seen in Figure 5.23 and Figure 5.24.

IP Address Settings

Remote host

The DNS name or the IP address of the remote host. You can shift between DNS name and IP address with the button to the right of the text box.

Proposal Parameters

Authentication key

You can select the authentication key used to authenticate your local host in this connection from the list of available keys. The available keys include both the certificates and the pre-shared keys. If there is a smart card reader attached to your computer and a smart card in the reader, also the certificates on the smart card are available. However, the listing shows the certificate of the certification authority not the actual end entity certificate. When the connection is negotiated, the smart card is read again and the end entity certificate issued by the certification authority in the listing is actually included in the proposal. Consequently, in a large corporation, an administrator needs not to configure each instance of the software individually since the rule is not bound to a certain smart card and a certif-

icate found on it. For more on management of the smart cards and authentication keys in general, see Chapter 6 'Managing Authentication Keys'.

Proposal type

You can select the type of the proposal that your host is going to make when negotiating a connection with another host. There are two options: *normal* and *legacy*. A normal proposal contains all the possible parameter values supported by SSH Sentinel thus making the proposal very long. Some security software have difficulties with such long proposals. The legacy proposal is a fallback option to be used in such situations. The main differences in the two types of proposals are highlighted in the table below.

Encryption

The encryption algorithm. Select the one you want to use from the list of supported algorithms. Naturally, the algorithm can be applied only if both ends support it. SSH Sentinel uses your choice as its initial proposal. If the other end does not support that particular algorithm, the hosts will negotiate and find a mutually supported one. The following are supported by the SSH Sentinel software: Rijndael, Twofish, Blowfish, Cast, 3DES and DES.

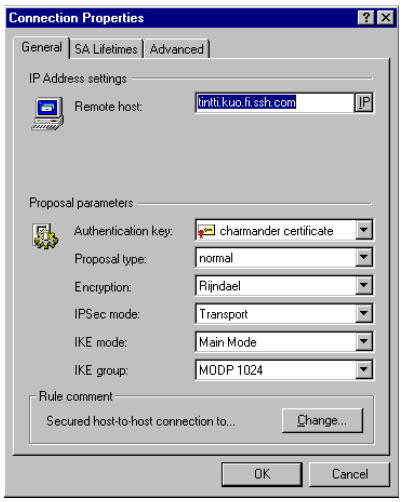

Figure 5.22 The general properties.

Proposal Parameter	Normal Proposal	Legacy Proposal
Encryption Algorithm	Rijndael, Twofish, Blowfish, 3DES, DES	Cast, 3DES, DES
IPSec Mode	Tunnel, transport	Tunnel
IKE Mode	Main, aggressive	Main, aggressive
IKE Group	MODP 768, MODP 1024	MODP 1024 MODP 1536

IPSec mode

There are two alternatives on the IPSec mode: *transport* and *tunnel*.

IKE mode

The Internet Key Exchange (IKE) mode setting has to do with negotiating the connection. Each time you wish to establish a connection, the system contacts the other end in order to establish the connection parameters or to check that they are correct and the connection works. There are two modes in the negotiations: *main* and *aggressive*. The latter is faster, whereas the former is more flexible and thus more likely to succeed. You can change the setting of the IKE mode by selecting from the list.

IKE group

When negotiating a connection, the communicating parties also settle the actual keys used to encrypt the data. The keys are based on the private keys of each party and some random data. The random data generation is based on so-called pool bits. The IKE group basically tells the number of pool bits. The more pool bits, the larger numbers in question. The larger the number, the harder to break. Consequently, more pool bits means more secure.

Rule Comment

Free-form text that appears on the bottom of the policy editor window when you have the rule selected. You can edit the text by clicking the `Change` button and typing the new description in the text field that appears.

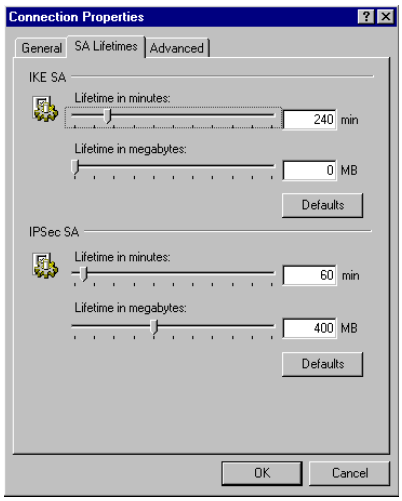

Figure 5.23 The settings for security association lifetimes.

Security Association Lifetimes

You can use the four sliders on the *SA Lifetimes* sheet seen in Figure 5.23 to control the lifetimes of the established security associations. There are distinct but similar controls for IKE and IPSec security associations. You can control the lifetime of either type of security association by time (minutes) or transferred data (megabytes). Whenever either limit is fulfilled, the existing security association is destroyed and a new, eventually, established.

Use the `Defaults` button to return the original values.

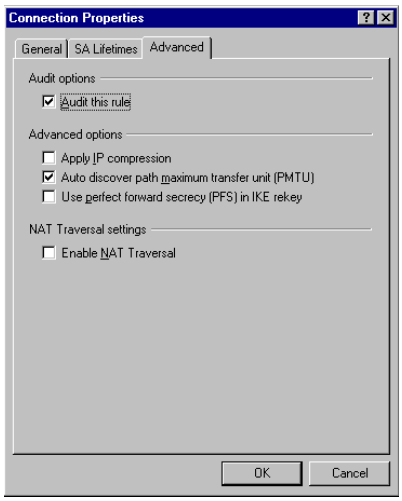

Figure 5.24 The advanced properties.

Audit Options

Audit this rule

To audit the rule, select this check box. For more information on auditing, see Section 7.1 'Auditing'.

Advanced Options

Apply IP compression

Each data packet transmitted is compressed. Consequently, transmitting it is faster. The compression is performed only if the other end also supports IP compression.

Auto-discover path maximum transfer unit (PMTU)

The system will find out the maximum transfer unit (MTU) size in the connection. This is done to avoid data fragmentation. In other words, the system finds out how big data packets it can transmit along the connection. Then it makes sure to send maximum size data

packets. Consequently, the data is distributed to a minimal number of packets and transmission is faster and less prone to errors.

Use perfect forward secrecy (PFS) in IKE rekey

When negotiating the actual keys again, the new key will in no way be dependent on the old key. This is, naturally, more secure, because even if somebody found out the old key, it does not reveal anything on the new key. On the other hand, the perfect forward secrecy slows down the key negotiations.

NAT Traversal Settings

Enable NAT Traversal

Network Address Translation is a method by which IP addresses are mapped from one network realm to another. A situation like this occurs, when a private network, with private addresses, is connected to the Internet. Traditional virtual private network solutions cannot work, however, with network address translation due to security problems. NAT Traversal is a solution to this.

5.5 Secured Networks

A secured network can be seen as an expansion of the concept of secured connection. A single policy rule affects the data traffic within an entire network segment. All traffic is encrypted and the communication endpoints authenticated. Secured network provides protection against active malicious attacks, unauthenticated network access, network sniffing and data tampering. Even wireless local area networks (WLAN), which are extremely vulnerable due to their open nature can be effectively protected.

Figure 5.25 The secured network layout.

It is recommended that when secure network rules are applied, a WINS server is used to avoid extra broadcast traffic. The Microsoft NetBIOS protocol, which is commonly encapsulated in TCP and UDP headers (and therefore subject to IPSec rules), is a rather noisy protocol and may generate IPSec/IKE negotiation "storms" between the hosts.

5.5.1 Adding Secured Network Rules

In order to insert a new secured connection rule, go through the following steps:

1. Initiate the adding of a new rule: Select *Secured Networks* in the policy tree and click the `Add` button. **OR** Select *Secured Networks* with the right mouse button and click `Add New Rule` in the menu that opens.
2. The dialog box shown in Figure 5.26 opens. Type in a descriptive name for the rule (A). It will be used as the rule name shown in the policy tree. Next, type in the network IP address and the subnet mask (B and C, respectively). Alternatively, you may choose to specify the IP address range occupied by the network section question. To do this, click button (F) and click `IP Address range` in the menu that opens. Finally, select the authentication key that you wish to use to identify the local host, using the list box (D). If there is a certificate on smart card available, the list box will show the certificate of the certification authority found on the smart card. If you select it, the actual authentication key used will be the end entity certificate issued by the certification authority. For details on managing the authentication keys and smart cards, see Chapter 6 'Managing Authentication Keys'. When ready, click the `OK` button. *Tip: To use the legacy proposal, e.g. the short form of the proposal, add the rule and set the connection properties properly. See Section 5.5.6 'Secured Network Rule Properties' for reference.*

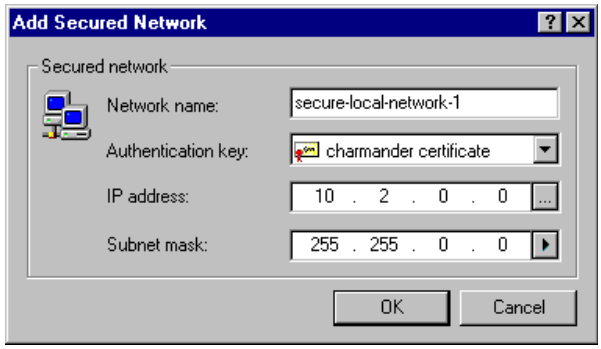

Figure 5.26 Adding a secured network rule.

3. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

5.5.2 Removing Secured Network Rules

To remove a rule, select it on the policy editor and click the `Remove` button. To make the removal permanent, click either `OK` or `Apply`. `OK` also closes the policy editor. You can restore the rule set after even several removals by clicking `Cancel`, provided that you haven't yet committed the changes with `Apply` or `OK`.

5.5.3 Viewing and Editing Secured Network Rules

Once you have created an IPSec rule on a network section, you can view the connection parameters and modify them at any time. You might want to change the connection parameters after unsuccessfully probing the parameters or because, simply, you know, that some setting has changed.

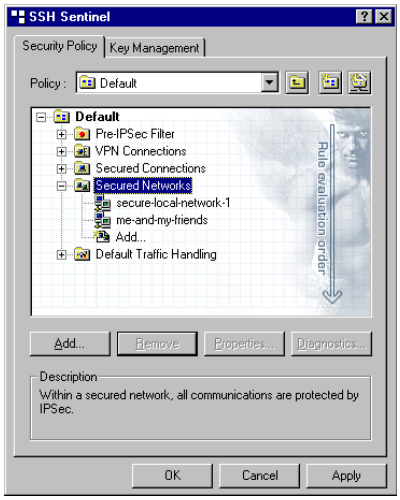

Figure 5.27 The basic listing of the secured network rules.

To see the basic listing of the configured secured network rules, seen in Figure 5.27 expand the `Secured Networks` branch in the policy tree. To get a more complete listing, seen in Figure Figure 5.28, double-click `Secured Networks`. A new view containing the listing with a few more details on each rule, opens. To view all properties of a rule, select the rule you want to view, and click the `Properties` button or double-click the rule you want to view in detail. The *Connection Properties* dialog box opens.

To edit a rule, do the following:

1. Open the Connection Properties dialog box: Select the rule you want to edit, and click the `Properties` button **OR** double-click the rule you want to edit.
2. The *Connection Properties* dialog box appears with the *General* sheet visible. You can shift to the *SA Lifetimes* and *Advanced* sheets by clicking the respective tab. Edit the values. See Section 5.5.6 'Secured Network Rule Properties' for assistance.

3. When ready, accept the changes by clicking `OK`. To discard your changes, click `Cancel`. Both buttons will take you back on the policy editor window.

4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

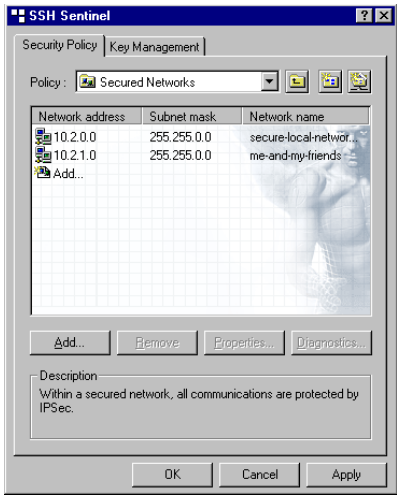

Figure 5.28 The detailed listing of the secured network rules.

5.5.4 Enabling and Disabling Secured Network Rules

You can disable a single rule - and naturally later enable it again. To disable (enable), select the rule with the right mouse button and click `Rule Enabled` in the menu that opens. A check mark appears next to the menu item text when the rule is enabled. When disabled, a little x-mark appears on the rule icon. Remember to commit you changes by clicking the `Apply` button.

5.5.5 Auditing Secured Network Rules

To audit a rule, do one of the following:

1. Select the rule with the right mouse button. Click `Audit Rule` in the menu that opens. A check mark appears to indicate that the rule is now audited.

2. Select the rule and open the *Connection Properties* window. Click the *Advanced* tab to view the audit options. Select the check box with label *Audit this rule*. Click the `OK` button to return.

Remember to commit your changes by clicking the `Apply` button.

To stop auditing a rule, do the opposite: Select the rule with the right mouse button and click `Audit Rule` again. The check mark now disappears. **OR** On the *Advanced* sheet of the *Properties* window, clear the check box *Audit this rule*. Remember to commit your changes.

To learn the details of auditing, see Section 7.1 'Auditing'.

5.5.6 Secured Network Rule Properties

The properties of a rule can be seen on the *Connection Properties* dialog box, seen in Figure 5.29. There are three sheets, the *General*, *SA Lifetimes* and *Advanced*. The latter two sheets can be seen in Figure 5.23 and Figure 5.24.

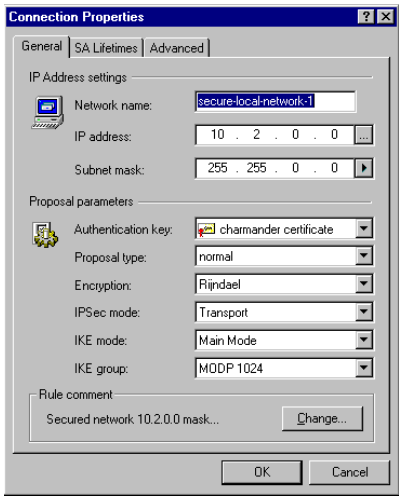

Figure 5.29 The general properties.

IP Address Settings

You describe network segment by specifying either the network IP address and the subnet mask or the IP address range occupied. To specify one or the other, click the arrow button (A) and select `IP Network` or `IP Address Range` in the menu that opens. The latter two of the descriptions below are on these menu items:

Network name

> An arbitrary name by which you wish to call the network. Something descriptive is recommended.

IP Network

The network IP address and the subnet mask. You can look the gateway IP address up by clicking the button to the right of the text box (B) and typing the host DNS name in the text box that appears.

IP Address Range

The IP address range occupied by the network.

Proposal Parameters

Authentication key

You can select the authentication key used to authenticate your local host in this connection from the list of available keys. The available keys include both the certificates and the pre-shared keys. If there is a smart card reader attached to your computer and a smart card in the reader, also the certificates on the smart card are available. However, the listing shows the certificate of the certification authority not the actual end entity certificate. When the connection is negotiated, the smart card is read again and the end entity certificate issued by the certification authority in the listing is actually included in the proposal. Consequently, in a large corporation, an administrator needs not to configure each instance of the software individually since the rule is not bound to a certain smart card and a certificate found on it. For more on management of the smart cards and authentication keys in general, see Chapter Chapter 6 'Managing Authentication Keys'.

Proposal type

You can select the type of the proposal that your host is going to make when negotiating a connection with another host. There are two options: *normal* and *legacy*. A normal proposal contains all the possible parameter values supported by SSH Sentinel thus making the proposal very long. Some security software have difficulties with such long proposals. The legacy proposal is a fallback option to be used in such situations. The main differences in the two types of proposals are highlighted in the table below.

Proposal Parameter	Normal Proposal	Legacy Proposal
Encryption Algorithm	Rijndael, Twofish, Blowfish, Cast, 3DES, DES	3DES, DES
IPSec Mode	Tunnel, transport	Tunnel
IKE Mode	Main, aggressive	Main, aggressive
IKE Group	MODP 768, MODP 1024 MODP 1536	MODP 1024

Encryption

The encryption algorithm. Select the one you want to use from the list of supported algorithms. Naturally, the algorithm can be applied only if both ends support it. SSH Sentinel uses your choice as its initial proposal. If the other end does not support that particular

algorithm, the hosts will negotiate and find a mutually supported one. The following are supported by the SSH Sentinel Software: Rijndael, Twofish, Blowfish, Cast, 3DES and DES.

IPSec mode

There are two alternatives on the IPSec mode: *transport* and *tunnel*.

IKE mode

The Internet Key Exchange (IKE) mode setting has to do with negotiating the connection. Each time you wish to create a connection, the system contacts the other end in order to establish the connection parameters or to check that they are correct and the connection works. There are two modes in the negotiations: *main* and *aggressive*. The latter is faster, whereas the former is more flexible and thus more likely to succeed. You can change the setting of the IKE mode by selecting the correct from the list.

IKE group

When negotiating a connection, the communicating parties also settle the actual keys used to encrypt the data. The keys are based on the private keys of each party and some random data. The random data generation is based on the so-called pool bits. The IKE group basically tells the number of pool bits. The more pool bits, the larger numbers in question. The larger the number, the harder to break. Consequently, more pool bits means more secure.

Rule Comment

Free-form text that appears on the bottom of the policy editor window when you have the rule selected. You can edit the text by clicking the `Change` button and typing the new description in the text field that appears.

Security Association Lifetimes

You can use the four sliders on the *SA Lifetimes* sheet seen in Figure 5.23 to control the lifetimes of the established security associations. There are distinct but similar controls for IKE and IPSec security associations. You can control the lifetime of either type of security association by time (minutes) or transferred data (megabytes). Whenever either limit is fulfilled, the existing security association is destroyed and a new, eventually, established.

Use the `Defaults` button to return the original values.

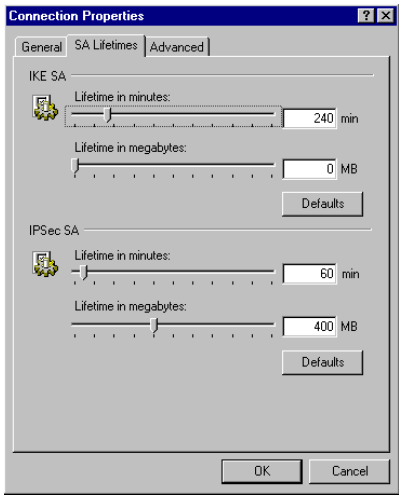

Figure 5.30 The security association lifetime settings.

Audit Options

Audit this rule

To audit the rule, select this check box. For more information on auditing, see Section 7.1 'Auditing'.

Advanced Options

Apply IP compression

Each data packet transmitted is compressed. Consequently, transmitting it is faster. The compression is performed only if the other end also supports IP compression.

Auto-discover path maximum transfer unit (PMTU)

The system will find out the maximum transfer unit (MTU) size in the connection. This is done to avoid data fragmentation. In other words, the system finds out how big a data packets it can transmit along the connection. Then, it makes sure to send maximum size data packets. Consequently, the data is distributed to a minimal number of packets and transmission is faster and less prone to errors.

Use perfect forward secrecy (PFS) in IKE rekey

When negotiating the actual keys again, the new key will in no way be dependent on the old key. This is, naturally, more secure, because even if somebody found out the old key, it does not reveal anything on the new key. On the other hand, the perfect forward secrecy slows down the key negotiations.

NAT Traversal Settings

Enable NAT Traversal

Network Address Translation is a method by which IP addresses are mapped from one network realm to another. A situation like this occurs, when a private network, with private addresses, is connected to the Internet. Traditional virtual private network solutions cannot work, however, with network address translation due to security problems. NAT Traversal is a solution to this.

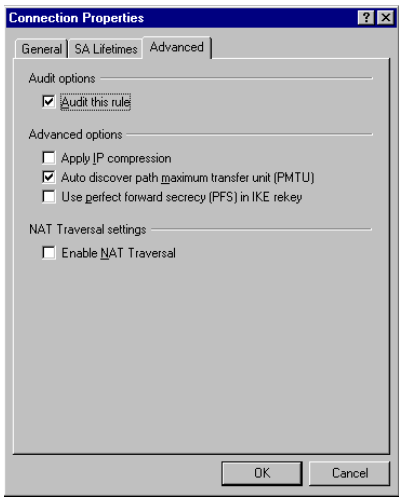

Figure 5.31 The advanced properties.

5.6 Default Response Rules

The default response rules can be found in the policy tree under *Default Traffic Handling* by expanding the *Default IPSec Response* branch. The responses to unprotected IP traffic and IPSec transmissions are configured separately. As the name implies, the default response rules only affect the incoming traffic. The rules determine, what to do, if none of the specific rules (filter rules or connection rules) matches the connection proposed by a remote host.

5.6.1 Default IPSec Response

To modify the default IPSec response, do the following:

1. Select any of the three default parameters hanging under `Default IPSec Response` in the policy tree. Either double-click it **OR** click the `Properties` button.

2. The *Default Settings* dialog box opens with the IPSec Response sheet visible. Reset the values. For reference, see the descriptions below.
3. Once ready, click the `OK` button to accept the changes. If you want to discard the changes, click `Cancel`. Both buttons will close the dialog box and take you back to the policy editor.
4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

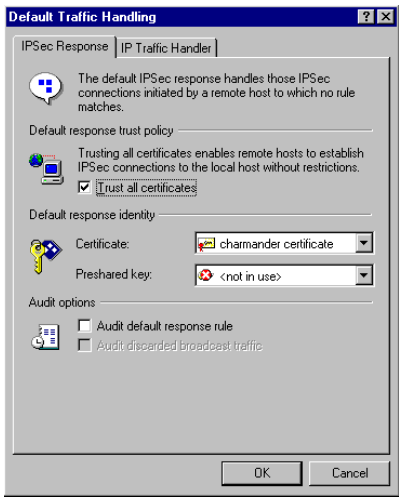

Figure 5.32 The default traffic handling: the default IPSec response.

IPSec Default Response Settings

Default response trust policy

Selecting the option `Trust All Certificates in IPSec Default Response` means that all certificates received from any remote host will be trusted by your host without further verification. If you don't check this option and your host receives a certificate that cannot be trusted, the traffic is dropped. Of course, this only affects the default IPSec traffic handling.

Default response identity

The default response identity describes how your host will authenticate itself. There are two option, a certificate and a pre-shared key. The default certificate is created during the installation of the SSH Sentinel software. You can manually change it to any available local certificate or choose the setting *not in use*. The latter would mean, that no certificate will be submitted to the remote end. You may reset the setting on the pre-shared key, too, to be any of the keys available or *not in use*. The system will submit the authentication key that is asked for by the remote end, the certificate or the pre-shared key.

Audit options

Select the check box *Audit default response rule* to audit the default response rule. For more information on auditing, see Section 7.1 'Auditing'.

5.6.2 Default IP Traffic Handling

To modify the default IP response open the *IP Traffic Handler* sheet on the *Default Traffic Handling* window, see Figure 5.33. You can choose between two options: to deny or allow all such unprotected incoming traffic to which no specific rule matches.

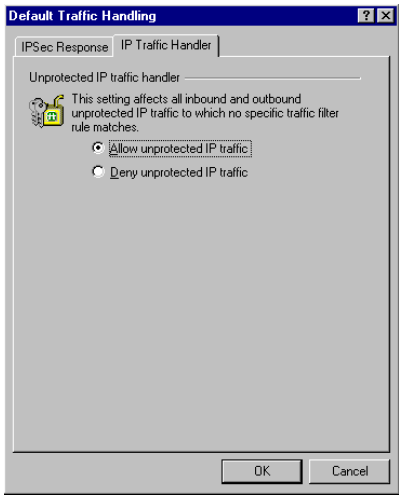

Figure 5.33 The default traffic handling: the default IP traffic handler.

Chapter 6 Managing Authentication Keys

The authentication keys are managed using the *Key Management* sheet of the policy editor, seen in Figure 6.1. The management is divided into three categories: the *trusted certificates*, the *authentication keys* and the *directory services*. The local authentication keys are common to all policy layers while the trusted certificates and directory services are distinct to each policy. You can manage - view and edit - the authentication keys of the active policy.

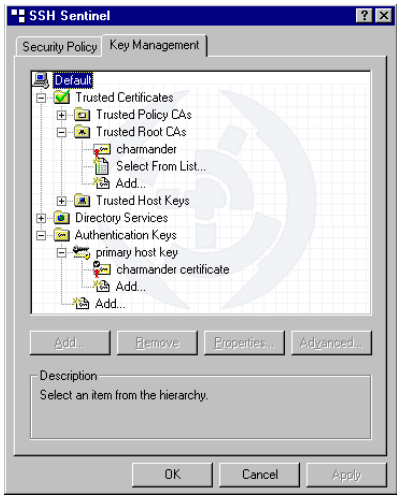

Figure 6.1 The authentication key management.

6.1 Trusted Certificates

The trusted certificates are certificates of the remote hosts that you have encountered. Despite the name, the listing might include certificates that you can not trust. Most often, such a certificate has previously been trusted by you but for some reason it has lost its credibility. For example, it validity period might have ran out. In such a situation, it is crucial that the certificate is still in your system and that the system clearly indicates the lost of credibility.

The trusted remote certificates are further classifies into three groups:

Trusted Policy CAs

You can share a policy certified by a certification authority listed here. In practice, if you try to share a policy signed by an authority that does not appear on this list, the system will prompt you to verify that it should trust the policy and the authority. If your answer is positive, then, on top of sharing the policy, the authority will be added on the list of trusted policy CAs.

Trusted Root CAs

You automatically trust all remote end certificates issued by a certification authority listed here. So, if you are establishing a connection to a remote host that provides you a certificate certified by an authority that you trust, the system does the decision on trusting the certificate automatically without prompting you. The remote end certificate itself will not be listed in your system. On the other hand, the self-signed certificate of your own local host is by default classified as a trusted root certificate. This way, you can certify remote end certificates and then trust them. In practice, you certify a remote certificate each time the system prompts you about a not trusted certificate and asks you the make the decision on trusting it. These remote end certificates appear in the listing of trusted host keys.

Trusted Host Keys

These are certificates of remote hosts. Most often these are self-signed certificates from a remote host. As explained in the previous section, when a remote end provides you such a certificate, the system will prompt you to make the decision on trusting it. If you decide to trust the certificate, you sign it yourself and the remote end certificate will be listed here.

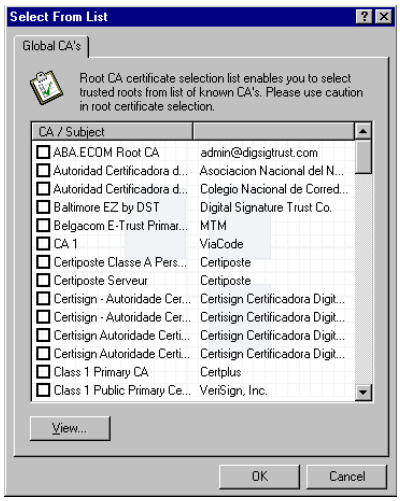

Figure 6.2 The selection list of the root CA certificates.

6.1.1 Viewing Certificates

You can view remote certificates in greater detail on the *Certificate Properties* dialog box. To open the dialog box, do one of the following:

1. On the policy editor, select the certificate you want to investigate and click the `Properties` button.
2. Double-click the certificate you want to view.
3. Select the certificate on the policy editor with the right mouse button, and select `Properties` in the menu that opens.

On the *Certificate Properties* dialog box, there are two sheets, *General* and *Details*. For information on the properties, see Section 6.1.4 'Certificate Properties'. For documentation on the `Export` button, see Section 6.1.3 'Exporting Certificates'.

6.1.2 Importing Certificates

To manually import a certificate, you must know where to find the certificate in your host or network. To import only valid certificates, the system does some elementary checking on them. For example, if you try to import a certificate of a trusted certification authority, the system checks that the certificate really is certificate of a certification authority.

There are several ways to import a remote end certificate. You can *add* a certificate, *drag-and-drop* a certificate into the key management tree, *paste* a certificate from the Windows clipboard or, when importing a root CA certificate, *choose it from a list* of available certificates.

To *add* a certificate, it must be in PEM encoded or binary format. Follow these step:

1. Initiate the adding of a new certificate: Select the category, where you want to add a new certificate and click the `Add` button. **OR** Select the category, where you want to add a new certificate, with the right mouse button and select `Add New Certificate` in the menu that opens. **OR** Double-click the selected category branch.
2. In the dialog box that opens, browse for the certificate in your file system. Once you have located it, click `Open`.
3. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

Alternatively, you can *drag and drop* a PEM encoded certificate file to the key management tree. Simply, drag the file over the branch where you want it to be placed and drop it there.

You can also *paste* the certificate from the Windows clipboard. To do this, select the correct branch from the key management tree with the right mouse button. In the menu that opens, select `Paste`. The normal key bindings, `Shift+Insert` and `Ctrl-v` also work. The certificate in the clipboard

needs to be in PEM encoded form. Naturally, to perform this operation, you must have copied the certificate to the clipboard in advance.

Selecting Root CA from List

When adding a new root certification authority, you can select it from a list provided by the system. The listing contains the most popular certification authorities available. To do this, go through the following steps:

1. Double-click the *Select from List* branch hanging under *Trusted Root CAs* in the key management tree.
2. A dialog box, seen in Figure 6.2 containing a listing on a number of certification authorities opens. Mark those you want to add by checking the box to the left of the CA name. You can investigate the certificate by selecting the authority and clicking the `View` button. For information on the certificate properties, see Section 6.1.4 'Certificate Properties'.
3. When you have selected all the CAs you want to add, click the `Add` button. To not to add any, click `Cancel`. Both buttons return you to the policy editor.
4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

6.1.3 Exporting Certificates

To export a certificate, in other words, to save it to a file, do the following:

1. Select the certificate you want to export and open the *Certificate Properties* dialog box. You can do this by clicking the `Properties` button.
2. On the *Certificate Properties* dialog box, click the `Export` button.
3. A standard dialog box for saving a file opens. Select the location of the certificate file and the file format. The PEM and binary formats are recommended as they are the most commonly supported.
4. Once you have saved the certificate, the dialog box disappears and you find yourself back on the *Certificate Properties* dialog. Close it by clicking `Close` and you return to the policy editor.

6.1.4 Certificate Properties

The certificate properties are shown on the *Certificate Properties* dialog box seen in Figure 6.3. There are two sheets, the *General* and *Details*. The latter can be seen in Figure 'Details'. The certificate information is shown in X.509 format.

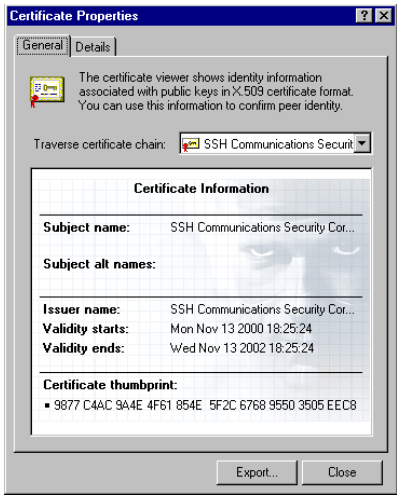

Figure 6.3 The general properties of a certificate.

General Properties

Traverse Certificate Chain

If the certificate you are viewing is not trusted, the text *"Not Trusted Certificate"* is shown in this drop-down list box.

However, the usual purpose of this box and the listing it contains, is to illustrate the connection between the particular certificate and the certification authority that you trust. Suppose, that the certificate is issued by a certification authority called A1. Then A1 is shown in the list. And, suppose that A1 has received its certificate from an upper-level certification authority called B2. Then, both B2 and A1 would be listed. Since you trust B2, you can trust the certificate it has issued to A1. Thus, you trust A1 and also the certificates issued by it.

Subject Name

The so-called common name of the certificate holder. The host DNS name, for example.

Subject Alt. Name

An alternative name of the host. Typically the IP address, the DNS name or email address.

Issuer Name

The issuer of the certificate. If the certificate is issued by a certification authority, the authority name is shown here. If this is a self-signed certificate, then, naturally, the name of the certificate holder is shown here.

Validity Starts

The beginning of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Validity Ends

The end of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Certificate Thumbprint

A checksum on the certificate calculated by the system. With the thumbprint, it is easy and fast to verify if two people are talking about the same certificate. They only need to check that the algorithm used and the thumbprint itself are the same.

Details

On the details sheet, shown in Figure 6.4, more information on the certificate is shown. If you select a field, then the lower part of the dialog box shows the value in its entirety.

Serial Number

The serial number of the certificate, given by the issuing authority. The certification authority binds a consecutive number to each certificate it issues.

Issuer

The issuer of the certificate. If the certificate was received from a certification authority, the authority name is shown here. If this is a self-signed certificate, then, naturally, the name of the certificate holder is shown here.

Subject

The so-called common name of the certificate holder. The host DNS name, for example.

Valid From

The beginning of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Valid To

The end of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Subject Alt. Names

An alternative name of the host. Typically the IP address, the DNS name or email address.

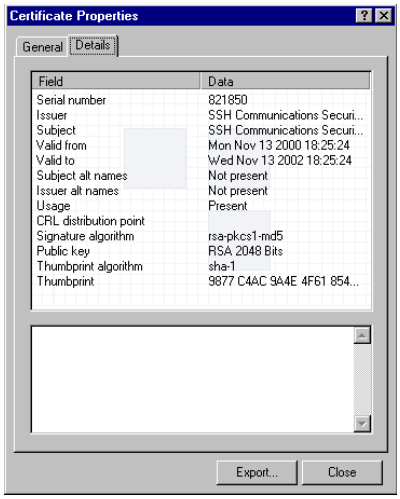

Figure 6.4 The details of a certificate.

Issuer Alt. Names

An alternative name of the certification authority. Typically the IP address, the DNS name or email address.

Usage

An description on how the certificate is intended to be used. Common examples are authentication and certifying other certificates.

CRL Distribution Point

Information on how the certificate revocation list (CRL) can be found. Revocation lists, published by certification authorities, list all certificates, issued by the authority in question, that for some reason have lost their credibility. The value in this field is merely a search key. The directory services, see Section 6.3 'Directory Services' are taken advantage of, when actually locating the revocation list.

Signature Algorithm

The algorithm applied when the keys associated with the certificate were created.

Public Key

The public key associated with the certificate. The data field shows the algorithm and key length whereas the key itself can be seen on the lower half of the window.

Thumbprint Algorithm

The algorithm used to calculate the certificate thumbprint, a checksum of the certificate. With the thumbprint, it is easy and fast to verify if two people are talking about the same certificate. They only need to check that the algorithm used and the thumbprint itself are the same.

Thumbprint

The checksum of the certificate.

For documentation on the `Export` button, see Section 6.1.3 'Exporting Certificates'.

6.2 Authentication Keys

The branch called *Authentication Keys* refers to the local host authentication keys. The local authentication keys characterize the local host. Consequently, they are common to all policies you have configured with SSH Sentinel on your host. For example, if you receive a new certificate from a certification authority, it is visible in all policies.

There are two types of local authentication keys: certificates and pre-shared keys. The local certificates are similar to the trusted remote certificates, presented in Section 6.1 'Trusted Certificates' as far as viewing and saving them and the certificate properties and formats are concerned. However, since local certificates are in question, you also can create the key pair associated with a certificate and request a certificate from a certification authority.

The pre-shared keys are, as the name implies, secrets shared by the communicating parties. To authenticate, the parties prove that they know the secret.

6.2.1 Certificate Enrollment Process

In order to enroll for a new certificate, you need an authentication key pair to be submitted to the certification authority. You can send your request online or save the request in a file and deliver it to the authority by some other means. How to create the key pair and the request as well as how to contact the authority using SSH Sentinel software is explained in Section 6.2.6 'Creating Certificates'.

After you have sent you request but before you have received the actual certificate from the authority, your request is in *pending* status. If you enrolled online, you can poll the certification authority, to find out the current status on your request. See Section 6.2.11 'Polling Certification Requests' for details.

The authority can return the certificate to you either online or in a file by some other means, in email, for example. If you receive it online, then the pending certificate is automatically updated to a real certificate. If you receive it in a file, you can import it from the file or you can copy it to the

clipboard and paste it into your key management tree. For details on importing and pasting, see Section 6.2.8 'Importing Certificates'

6.2.2 Viewing Local Authentication Keys

The local authentication keys can be found hanging under the *Authentication Keys* branch in the key management tree. A sample tree is shown in Figure 6.5.

The certificates are always situated under a key pair. The key pairs are given consecutive names: *primary host key*, *secondary host key* etc. Several certificates can share the same key pair. In Figure 6.5, there is a key pair with two certificates (A).

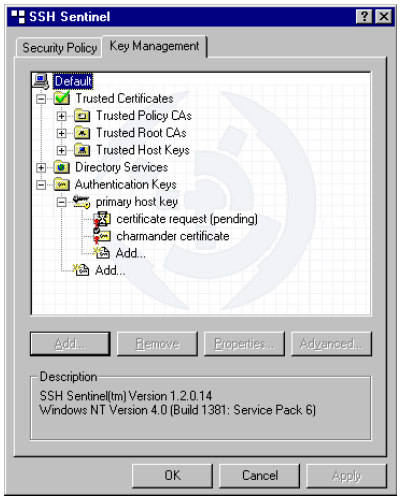

Figure 6.5 The local authentication keys.

When installing SSH Sentinel, you can create a self-signed certificate. When using the product, however, you only can create a key pair and request certificate from a certification authority.

After requesting, it takes some time before the certification authority returns the certificate to you. A certificate in pending status, is shown in (B).

The pre-shared keys are listed below the key pairs. When you add new key pairs or pre-shared keys, the order might be temporarily mixed. However, when you commit the changes, that is, click `Apply` or `OK` in the policy editor, the order is restored: first the key pairs and certificates, then the pre-shared keys.

If there is a smart card reader attached to your host and a smart card in the reader, there is an extra branch called *Smart Card Keys* in the authentication keys tree. Under that branch, you can see the key pairs and related certificates found on the smart card. For details on working with certificates on smart cards, see Section 6.2.3 'Certificates on Smart Cards'.

You can view the details of a certificate or a pre-shared key on the *Properties* dialog box. To open it, do one of the following:

1. Select the certificate or pre-shared key you want to investigate and click the `Properties` button.
2. Double-click the certificate or pre-shared key you want to view.
3. Select the certificate or pre-shared key with the right mouse button, and select `Properties` in the menu that opens.

For assistance on the certificate properties, see Section 6.2.12 'Certificate Properties', and on pre-shared keys properties, Section 6.2.13 'Pre-shared Key Properties'.

6.2.3 Certificates on Smart Cards

The authentication keys and the related certificates found on a smart card are mostly handled in the same way as other certificates. However, there are some differences which are explained in this section.

The requirements for smart cards and readers are as follows: The smart card readers need to be compliant with the PC/SC standard. Readers from Setec, Bull and Towitoko have been tested. The smart card itself has to be of format PKCS#15. There has to be the CA certificate available on the smart card along with - naturally - the end entity certificate and key pair.

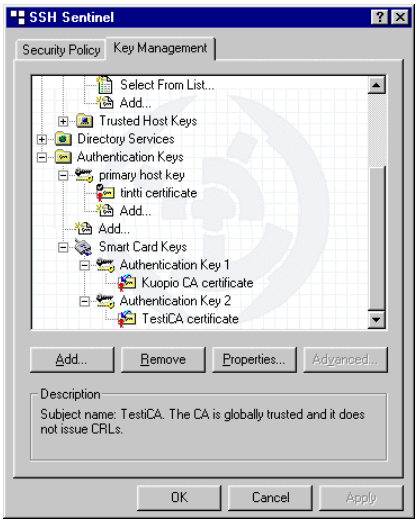

Figure 6.6 A smart card is available and thus shown in the key management tree.

If a smart card is available, the key pairs and certificates found on it can be used like other authentication keys in connection rules etc. The smart card icon appears in the key management tree, see Figure 6.6. If a smart card reader and a smart card are unavailable, the smart card icons and features related to smart cards are invisible. The availability of the reader and the card is checked

every five seconds. Thus, you can put the card in the reader during your session and the smart card icon soon appears.

Since certificates in general cannot be modified by SSH Sentinel, you cannot modify a certificate found on a smart card either. In addition, you cannot delete a smart card certificate or key pair. Further, you cannot add or import a new key pair and a related certificate on the smart card. Neither can you create a new certificate based on a key pair found on the smart card. However, you can export a certificate from the smart card into a file. Naturally, you can view the properties of the certificate. The smart card key pairs and certificates are not stored in the policy database; they appear and disappear with the smart card.

SSH Sentinel demands that the smart card certificates are issued by a certification authority (CA). It is recommended that you add that CA certificate to your trusted root CAs. Then, the certification authority is unambiguously trusted in all situations. If you don't, however, the certification authority in question is only trusted in connection to the rules where the particular certificate is used as the authentication method. This would mean, that if the remote host responds with a certificate issued by the same certification authority, the certificate would be trusted, even though the CA is not a trusted root CA. In all other situations, the certification authority is not regarded trustworthy. As this may result in awkward situations, it is recommended to add the CA certificate to the trusted root CAs. The CA certificates found on the smart card are listed on the *Smart Card* sheet of the *Select from List* dialog seen in Figure 6.7.

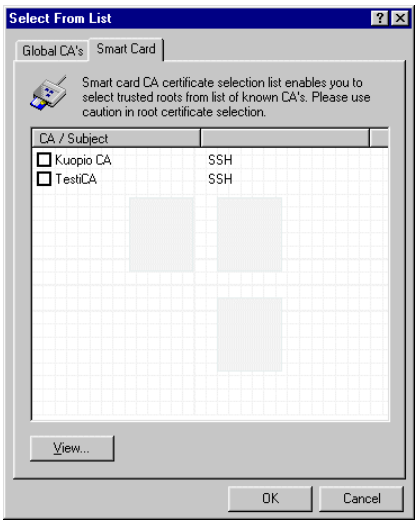

Figure 6.7 The certification authorities available on the smart card.

When you configure a policy rule and select the certificate on the smart card as the authentication method, the certificate shown is actually the one of the certification authority. The certificate submitted to the remote host, is, however, the local end entity certificate. Configuring the rule with the CA certificate provides for centrally managed policies and policy management by a system administrator. Suppose the following situation: In a company, all employees have smart cards with certificates issued by a common corporate CA. The system administrator can configure the rules with his own smart card. In plain English, the rule states that the certificate is found on the smart card and is issued by the corporate CA. When the rule is applied on an end user host, the end user smart

card in the reader is read and the certificate that matches the criteria is actually used as the authentication method. If the rules were configured with the actual end entity certificates, the administrator would have to configure each host individually.

If there are several certificates on the same smart card, then the system does not necessarily know which certificate it should use when a connection is proposed by a remote host. SSH Sentinel then prompts the user to make the decision. This may result in seemingly random prompts on the screen.

6.2.4 Editing Pre-shared Keys

You cannot edit a certificate or the authentication key pair, but you can edit a pre-shared key. Naturally, only change the secret, if you actually have agreed on the modification with the other end.

Editing the secret is performed on the *Pre-shared Secret Properties* dialog box.

1. Open the *Pre-shared Key Properties* dialog as guided in Section 6.2.2 'Viewing Local Authentication Keys'.
2. Update the secret, see Section 6.2.13 'Pre-shared Key Properties' for assistance.
3. Once ready, click the `OK` button. To discard the updates, click `Cancel` instead. You return to the policy editor.
4. Back on the policy editor, click `Accept` or `OK` to accept the changes and put them into effect. `OK` will also close the policy editor window. To discard the changes, click `Cancel`. **Note!** These actions affect all modifications in the rule set and the key management. Thus `Apply` and `OK` will commit and `Cancel` discard all changes made so far.

6.2.5 Creating Authentication Keys

To create a new authentication key, launch the wizard that guides you through the process:

1. Select the *Authentication Keys* branch from the key management tree and click the `Add` button.
2. Select the *Authentication Keys* branch from the key management tree with the right mouse button, and select *Add New Auth Key* in the menu that opens.
3. Select the *Add* branch, hanging lowest under the *Authentication Keys*, below all key pairs and pre-shared keys.
4. Select any key pair, certificate or pre-shared key and click the `Add` button.
5. Select any key pair, certificate or pre-shared key with the right mouse button, and select `Add New Auth Key` in the menu that opens.

The wizard displays the window shown in Figure 6.8. You can do three things with the wizard: Create a new key pair and enroll for a certificate, use an existing key pair and enroll for a certifi-

cate and create a new pre-shared key. It is worthwhile to notice that you are not able to create a new self-signed certificate.

The first window in the wizard, shown in Figure 6.8, appears a little different depending on where your focus is when you launch it. If the focus is on a key pair - on the key-pair itself, on a certificate hanging under it or the *Add* branch under it - all three options are available. If you choose to request a new certificate, the system uses the key pair where your focus is, in enrollment. The net effect is that you will receive a new certificate based on that key pair, and it will be situated under the key pair. But if you choose to create a new key pair and certificate or to create a new pre-shared key, then the certificate or pre-shared key will be independent of the key pair in focus and situated hanging last in the *Authentication Keys* branch.

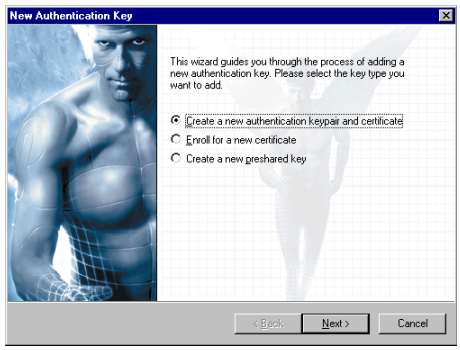

Figure 6.8 Adding a new authentication key.

If your focus is on something else than a existing key pair, you only have two options available: Creating a new key pair and certificate and creating a new pre-shared key. The system cannot just enroll for a certificate since it has no key pair associated with the request. The outcome of the process is a new key pair and certificate or a new pre-shared key, depending on your choice.

The creation of a certificate and a pre-shared key are explained separately in the following sections. Quite naturally, you can not create a certificate from an authentication key pair found on a smart card.

6.2.6 Creating Certificates

Select Key Type

On the *New Authentication Key* window shown in Figure 6.8, select if you want to create a new key pair and certificate or just enroll for a new certificate using an existing key pair. If the enrollment option is not active, your focus in the key management tree is off. In that case, click `Cancel` and once back on the policy editor, select the key pair you want to use in enrollment. Then launch the wizard again. See Section 6.2.5 'Creating Authentication Keys' for a more detailed discussion.

For documentation on creating a new pre-shared key by choosing the option *Create New Pre-shared Key* (C), see Section 6.2.7 'Creating Pre-shared Keys'.

Click `Next` to proceed.

Generate New Key Pair

If you chose only to enroll for a new certificate - the middle option in Figure 6.8 - skip this step. If you chose to create a whole new key pair - the top option - the window shown in Figure 6.9 opens.

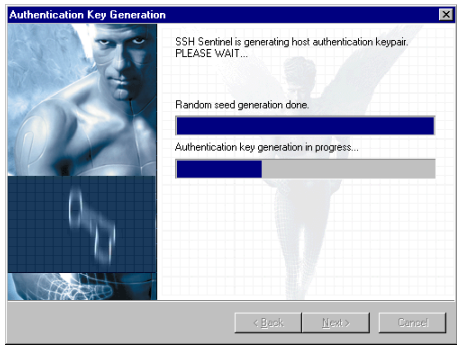

Figure 6.9 Generating the authentication key pair.

In this step, the key pair is generated. First, a random seed for the key generation is collected. The seed is given by you moving the mouse or typing in text. User random input is taken advantage of to ensure that the keys generated are unique. The chance that two inputs are alike and the keys generated similar, is infinitesimal.

Having collected enough data for the seed, the software goes on to calculate the actual key pair. This may take some twenty seconds to complete.

After the generation is complete, click `Next` to proceed.

Identity Information

The *Authentication Key Information* dialog box shown in Figure 6.10 opens. The data gathered on this dialog characterizes your local host. Fill in the fields with appropriate values (see below). To proceed, click `Next`.

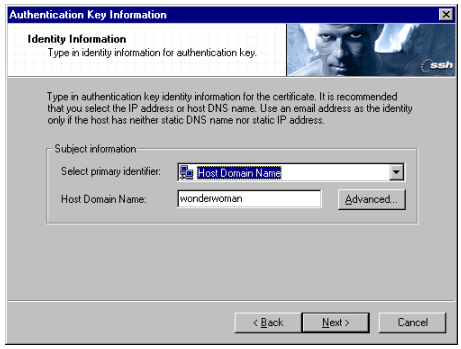

Figure 6.10 The identity information.

Primary Identifier

The primary piece of information that identifies your host. You can choose between the IP address, the DNS name and the email address. It is recommended that you choose either of the first two, preferably the DNS name. Only static DNS names and IP addresses may be used.

If your host lacks both static IP address and static DNS name, use email address instead. However, since IPSec rules are normally bound to DNS names or IP addresses, this might cause problems in interoperability with the software used in the remote end.

The field below changes its name according to your choice on primary identifier. Type in it the actual identifier, that is the IP address of your host, the DNS name of your host or the email address. Whatever you specify here will appear under the header *Subject Name* on the *Certificate Properties* dialog box. See Section 6.2.12 'Certificate Properties' for further reference.

Advanced

By clicking the `Advanced` button, the dialog box shown in Figure 'Identity Information' opens. You can specify the `organization` and the `country`.

To proceed, click

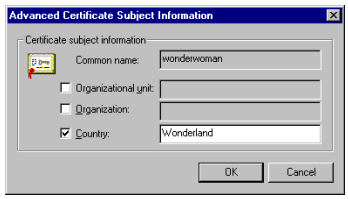

Figure 6.11 Further information on the certificate subject.

Enrollment Information

The *Certification Authority* dialog box is shown in Figure 6.12. On this dialog, you specify data associated with the certification authority to whom you send your certificate request.

Enrollment Protocol

Two online enrollment protocols are supported: Simple Certificate Enrollment Protocol (SCEP) and Certificate Management Protocol (CMP). Alternatively, you may choose to deploy an off-line protocol, PKCS#10.

CA Server Address

Available only if either online enrollment protocol is chosen. The URL of the online enrollment service.

CA Certificate

Available only if either online enrollment protocol chosen. The certificate of the certification authority. The certificate is needed to encrypt your certificate request before sending it to the certification authority.

You can type the name of the certificate in this field in which case the system fetches a certificate by that name from the URL in the previous field. Alternatively, you can type in the URL where the certificate can be found. The system then fetches the certificate from that URL. Or, you can use the menu item *Paste from Clipboard* with obvious functionality. Naturally, in this case you must have copied the certificate to the clipboard in advance.

Advanced

The button opens the *Advanced Settings* dialog box where you can set the socks and proxy settings. These are needed to get through the firewall, if the certification authority server is protected by one.

Reference Number

Needed if CMP protocol is applied in requesting the certificate. The reference number along with the key (see below) is used to identify the user requesting the certificate.

Key

Needed if CMP protocol is applied in requesting the certificate. A shared secret, granted by the certification authority. The key identifies the user requesting the certificate.

Click `Finish` to complete the process.

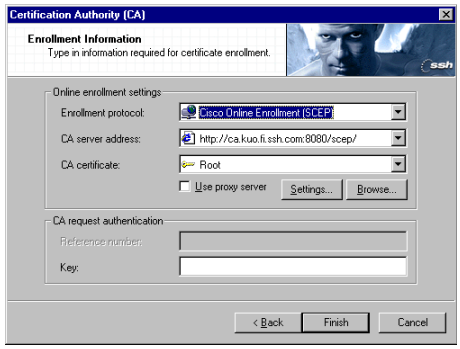

Figure 6.12 Certificate enrollment information.

The Finishing Steps

If you are doing file-based enrollment, the request is saved into a file. You must take care of delivering your request to the certification authority yourself.

If an online request is in question the system tries to contact the certification authority server. If it is not possible, an error message is shown.

If you don't possess the certificate of the certification authority from whom you a requesting a certificate online the system fetches the certificate from the remote server. Since the authority is not listed as a trusted root CA, the system can't accept the certificate automatically. It prompts you to make the decision on trusting the certificate. By common sense, you should trust the CA and thus answer *Yes* here. The CA certificate is then imported. Finally, if all went right, the system tells you that the certificate you requested is in *pending* status.

The certification authority will, if it decides to issue you a certificate, return the certificate to you. If you have requested online, you can *poll* the certification authority from time to time to check the status on your request. See Section 6.2.11 'Polling Certification Requests' for reference. Once the authority has returned the certificate, the status is changed automatically by the system and you can start using the certificate. In case of an off-line request, the certificate is most likely returned to you in a file. You must manually import the certificate into the system. See Section 6.2.8 'Importing Certificates' for reference.

6.2.7 Creating Pre-shared Keys

To create a new pre-shared key, click `Add` on the key management sheet. The dialog box shown in Figure 6.8 opens.

Select Key Type

Select the option *Create New Pre-shared Key*. For documentation on the other two option see Section 6.2.6 'Creating Certificates'. Click `Next` to proceed.

Pre-shared Key Information

The dialog box shown in Figure 6.13, you specify information concerning the pre-shared key itself.

Figure 6.13 Creating the shared secret.

Pre-shared Key Name

Give the pre-shared key an arbitrary name. Something descriptive is recommended.

Enter Shared Secret

Fill in the secret shared by you and the other end. The text will not be shown in plain text. The general rules on choosing passwords are valid here, too.

Confirm Shared Secret

Re-type the string to detect potential typos.

SHA-1 Thumbprint

The system automatically calculates a checksum using SHA-1 MAC algorithm on the shared secret while you type it. The thumbprint is unique and the calculation is extremely hard to invert, meaning that it is virtually impossible to figure out the shared secret, even if you know the algorithm and the thumbprint. You can confirm with the other communicating party, that the secret that you are using is the same, by exchanging the thumbprint on the phone, for example. That way, there is no need to reveal the actual secret. Naturally, both ends need to apply the same algorithm in order to get similar thumbprints.

Finally, click `Finish` to complete the creation of a pre-shared key.

The identity of the created pre-shared key is set to *no identity* by default. You can update the identity information, as well as other properties of a pre-shared key, on the *Properties* dialog box. See Section 6.2.13 'Pre-shared Key Properties' for details.

6.2.8 Importing Certificates

If the certification authority returned the certificate you've requested in a file, you can *import* the file or *copy* the certificate to the clipboard and *paste* it to the key management tree. You can also *drag and drop* the certificate file to the key management tree.

To import a certificate from a file, it needs to be saved in the PEM encoded format. Select the key pair corresponding to the certificate in question in the key management tree with the right mouse button. Select `Import New Auth Key` in the menu that opens. The standard Windows dialog box for opening a file opens. Locate the certificate file, and click `OK`.

Alternatively, you can drag and drop a PEM encoded certificate file to the key management tree. Simply, drag the file over the branch where you want it to be placed and drop it there.

To paste the file from Windows clipboard, first copy the contents of a PEM encoded certificate file to the clipboard. Then select the branch where you want the certificate to appear on the key management tree with the right mouse button. Select `Paste` from the menu that opens. Normal key bindings also work, thus `Shift+Insert` and `Ctrl+v` are equivalent to the menu item `Paste`.

6.2.9 Exporting Certificates

To export a certificate, in other words, to save it in a file, do the following:

1. Select the certificate you want to export and open the *Certificate Properties* dialog box. You can do this by clicking the `Properties` button, for example.
2. On the *Certificate Properties* dialog box, click the `Export` button.
3. A standard dialog box for saving a file opens. Select the location for the certificate file and the file format. The PEM and binary formats are the most commonly supported.
4. Once you have saved the certificate, the *Save* dialog box disappears and you find yourself back on the *Certificate Properties* dialog. Close it by clicking `Close` and you return to the policy editor.

6.2.10 Removing Certificates and Pre-shared Keys

To remove a certificate or a pre-shared key, select it on the policy editor and click the `Remove` button. To make the removal permanent, click either `OK` or `Apply`. `OK` also closes the policy editor. You can restore the key management tree after even several removals by clicking `Cancel`, provided that you haven't yet committed the changes with `Apply` or `OK`.

6.2.11 Polling Certification Requests

Once you have sent a certification request to an authority, the certificate appears in *pending* status. You can *poll* the status of your request. In Figure 6.5 there is a certificate in pending status. If you select a pending certificate, the `Poll` button appears. Click it, and the system starts to poll.f the certification authority has already issued you the certificate, the poll will return the certificate to you. The status of the certificate is then changed. If the request is still pending, the status is not updated. The system automatically polls the requests from time to time. As soon as the authority issues the certificate, the status of it is updated.

6.2.12 Certificate Properties

The certificate properties are shown on the *Certificate Properties* dialog box seen in Figure 6.3. There are two sheets, the *General* and *Details*. The latter can be seen in Figure 6.4. The certificate information is shown in X.509 format.

General Properties

Traverse Certificate Chain

If the certificate you are viewing is not trusted, the text *"Not Trusted Certificate"* is shown in this drop-down list box.

However, the usual purpose of this box and the listing it contains, is to illustrate the connection between the particular certificate and the certification authority that you trust. Suppose, that the certificate is issued by a certification authority called A1. Then A1 is shown in the list. And, suppose that A1 has received its certificate from an upper-level certification authority called B2. Then, both B2 and A1 would be listed. Since you trust B2, you can trust the certificate it has issued to A1. Thus, you trust A1 and also the certificates issued by it.

Subject Name

The so-called common name of the certificate holder, the local host. The host DNS name, for example.

Subject Alt. Name

An alternative name of the local host. Typically the IP address, the DNS name or email address.

Issuer Name

The issuer of the certificate. If the certificate was received from a certification authority, the authority name is shown here. If this is a self-signed certificate, then, naturally, the name of the local host is shown here.

Validity Starts

The beginning of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Validity Ends

The end of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Certificate Thumbprint

A checksum on the certificate calculated by the system. With the thumbprint, it is easy and fast to verify, if two people are talking about the same certificate. They only need to check that the algorithm used and the thumbprint itself are the same.

Details

On the details sheet, shown in Figure 6.4, more information on the certificate is shown. If you select a field, then the lower part of the dialog box shows the value in its entirety.

Serial Number

The serial number of the certificate, given by the issuing authority. The certification authority binds a consecutive number to each certificate it issues.

Issuer

The issuer of the certificate. If the certificate was received from a certification authority, the authority name is shown here. If this is a self-signed certificate, then, naturally, the name of the local host is shown here.

Subject

The so-called common name of the local host. The host DNS name, for example.

Valid From

The beginning of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Valid To

The end of the validity period, date and time. Outside of the validity period, the certificate is classified as not trusted.

Subject Alt. Names

An alternative name of the host. Typically the IP address, the DNS name or email address.

Issuer Alt. Names

An alternative name of the certification authority. Typically the IP address, the DNS name or email address.

Usage

An description on how the certificate is intended to be used. Common examples are authentication and certifying other certificates.

CRL distribution point

Information on how the certificate revocation list (CRL) can be found. Revocation lists are published by certification authorities. They list all certificates, issued by the authority in question, that for some reason have lost their credibility. The value in this field is merely a search key. The directory services (see Section 6.3 'Directory Services') are taken advantage of, when actually locating the revocation list.

Signature Algorithm

The algorithm used to create the keys associated with the certificate.

Public Key

The public key associated with the certificate. The data field shows the algorithm and key length whereas the key itself can be seen on the lower half of the window.

Thumbprint Algorithm

The algorithm used to calculate the certificate thumbprint, a checksum of the certificate. With the thumbprint, it is easy and fast to verify, if two people are talking about the same certificate. They only need to check that the algorithm used and the thumbprint itself are the same.

Thumbprint

The checksum of the certificate.

For documentation on the `Export` button, see Section 6.2.9 'Exporting Certificates'.

6.2.13 Pre-shared Key Properties

General Properties

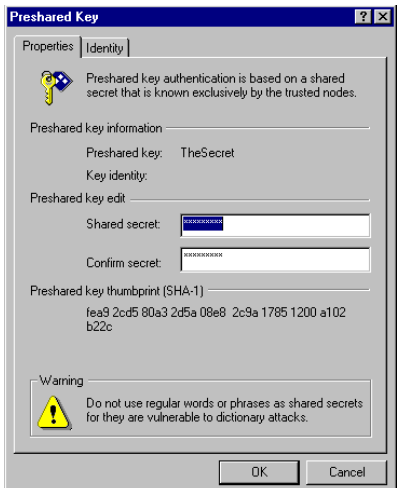

Figure 6.14 The general properties of a pre-shared key.

Pre-shared Key

The name of the pre-shared key. It is given by the user when creating the pre-shared key. Not updateable.

Key Identity

The local host identifier. You can change the identity on the *Identity* sheet.

Shared Secret

The secret in non-readable form.

Confirm Secret

The check field to avoid typos.

Thumbprint

The checksum of the shared secret calculated by the system using SHA-1 MAC algorithm. The thumbprint is unique and the calculation is extremely hard to invert, thusr it is virtually impossible to figure out the shared secret, even if knowing the algorithm and the thumbprint. You can confirm with the other communicating party the same secret is used by exchanging the thumbprint.

Identity

By default, the identity associated with the pre-shared key is set to *no identity*. You can change the identity to the host *IP address*, the host *domain name* or *email address*. Select the type of the identifier from the drop-down list in Figure 6.15. The field below, where the actual identifier is typed, changes its title and layout according to your selection.

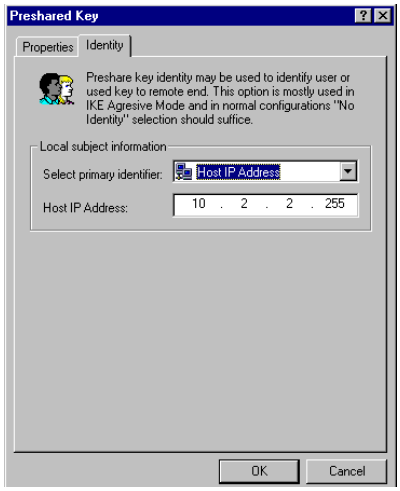

Figure 6.15 The identity properties of a pre-shared key.

6.3 Directory Services

The certification authorities publish revocation lists on certificates that for some reason have lost their credibility. Before establishing a connection, your system should check the revocation list, to verify that the remote end's certificate is still trusted by the authority that has issued it.

To be able to locate the revocation list that a certification authority has placed on a remote server, you need to define a directory service.

6.3.1 Adding Directory Services

To add a new directory service, do the following:

1. Initiate the adding of a new directory service: Double-click the *Add* branch hanging under *Directory Services* **OR** select the *Directory Services* branch with the right mouse button and select `Add New Directory Service` in the menu that opens **OR** select the *Directory Services* branch, and click the `Add` button.
2. The *Directory Service* dialog box opens. Fill in the values. For assistance, see Section 6.3.4 'Directory Service Properties'.

3. Accept the changes by clicking OK. To not to add the directory service after all, click Cancel. Both buttons will take you back to the policy editor.
4. Back on the policy editor, click Accept or OK to accept the changes and put them into effect. OK will also close the policy editor window. To discard the changes, click Cancel. **Note!** These actions affect all modifications in the rule set and the key management. Thus Apply and OK will commit and Cancel discard all changes made so far.

6.3.2 Viewing and Editing Directory Services

Expand the *Directory Services* branch and you see a listing on the existing directory services. The details of a service are shown on the *Directory Service* dialog box. To open and use it, do the following:

1. On the *Key Management* sheet, select the directory service you want to investigate and click the Add button **OR** double-click the directory service you want to view or edit.
2. The *Directory Service* dialog box opens. There are two sheets, the general and advanced, shown in Figure 6.16 and Figure 6.17. Investigate and update the values.
3. Once ready, click OK to accept your changes. Cancel discards the changes you made on this dialog box. If you made no updates, it doesn't matter which button you click. Both close the *Directory Service* dialog box and take you back to the policy editor.
4. Back on the policy editor, click Accept or OK to accept the changes and put them into effect. OK will also close the policy editor window. To discard the changes, click Cancel. **Note!** These actions affect all modifications in the rule set and the key management. Thus Apply and OK will commit and Cancel discard all changes made so far.

6.3.3 Removing Directory Services

To remove a directory service, select it on the policy editor and click the Remove button. To make the removal permanent, click either OK or Apply. OK also closes the policy editor. You can restore the key management tree after even several removals by clicking Cancel, provided that you haven't yet committed the changes with Apply or OK.

6.3.4 Directory Service Properties

General Properties

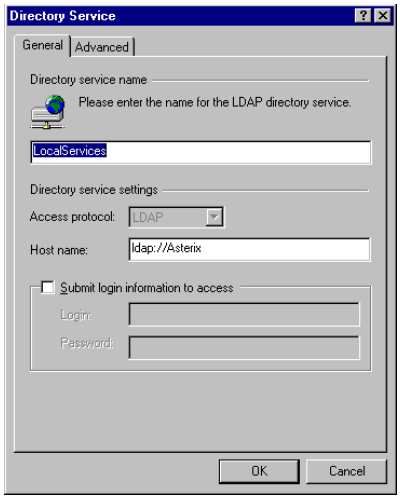

Figure 6.16 The general properties.

Directory Service Name

A descriptive name given by you to the directory service. It will be used as the name shown under the *Directory Services* branch in the key management.

Access Protocol

The protocol used in accessing the service. At the moment, only LDAP (Light-weight Directory Access Protocol) is supported.

Host Name

The server that provides the directory service.

Login Information

Check the box, if the server requires you to log in. Specify the login name and password in the respective fields.

Advanced Properties

Server Port Number

The number of the server port used.

Proxy Settings

The proxy and socks settings are taken advantage of when getting through the firewall if the server is protected by one. Click the `Settings` button to view and edit the settings.

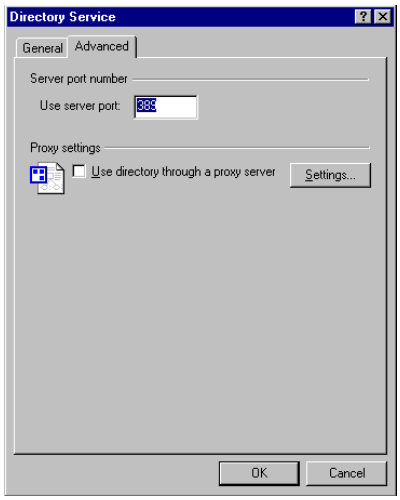

Figure 6.17 The advanced properties.

Proxy and Socks Settings

Click the `Auto-detect` button to automatically find out the settings. To accept the changes, click `OK`. To discard them, click `Cancel`. Both buttons return you to the *Directory Service* window. Naturally, you can also modify the settings by hand.

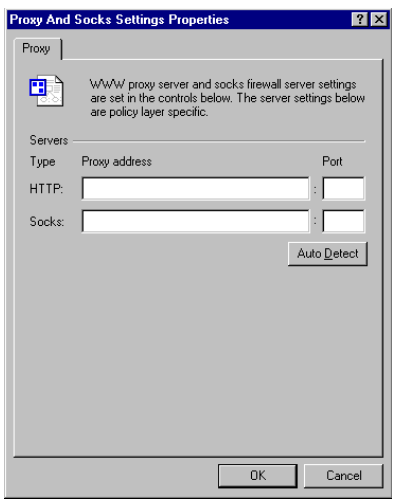

Figure 6.18 Settings for the proxy server and the socks firewall server.

Chapter 7 **Maintenance**

There is a number of tools available in SSH Sentinel to monitor the network traffic and the actions taken by the system. These tools range from connection diagnostics - simple IKE probing to check a single connection - to general network traffic statistics and to auditing rules.

7.1 Auditing

You can set any of the rules in the security policy to be audited. Being audited, each time the rule is applied, the event appears in an *audit log*. You can view the audit log through a web interface.

7.1.1 Auditing Rules

You can set any of the rules - filter rules, IPSec rules and default response rules - to be audited. Naturally, you can audit multiple rules simultaneously.

Figure 7.1 The audit options.

To audit a rule, do one of the following:

1. Select the rule with the right mouse button. Click `Audit Rule` in the menu that opens. A check mark appears to indicate that the rule is now audited.
2. Select the rule and open the *Connection Properties* window. Click the *Advanced* tab to view the audit options. Select the check box with label *Audit this rule*. Click the `OK` button to return.

Remember to commit your changes by clicking the `Apply` button.

To stop auditing a rule, do the opposite: Select the rule with the right mouse button and click `Audit Rule` again. The check mark now disappears. **OR** On the *Advanced* sheet of the *Properties* window, clear the check box *Audit this rule*. Remember to commit your changes.

7.1.2 Audit Options

To manage the audit logs, use the *Audit Options* property sheet, seen in Figure 7.1. Open it from the SSH Sentinel main menu: Select `Auditing`, and then select `Audit Options` in the submenu that opens.

Allowed Users

Allowed users

In the *allowed users* text box all the users that are granted access to your audit logs are shown. When opening the web interface to view the audit logs, user name and password are inquired.

You can add users by clicking the `Add` button and typing the user name and password (twice for checkup) in the text fields provided. You can remove allowed users with the `Remove` button and change the password by clicking the `Properties` button.

Audit server port

The port to access the audit server, the interface to the audit logs.

Allow remote access to audit server

If selected, remote access to the audit server is allowed. If not selected, the audit server can only be accessed from your local host.

Log file disk space usage

The log files tend to increase in size, especially if you audit many rules and there is a lot of network action on your host. It is crucial that you limit the size of the log files and that you remove old files regularly. Use these controls to achieve automatic clean-up and limitation of disk usage.

Min. free space

This setting limits the size of the audit log by dictating how much free space there should be left on the hard disk after the audit log is saved. If there isn't any space left, the log is not written and you miss the events.

To change the figure, either move the slider or type the figure in the text filed. The scale is from 0 to 1024 megabytes.

History length

This figure shows the clean-up interval. All logs older that this figure are automatically removed from the hard disk.

To change the figure, either move the slider or type the figure in the text filed. The scale is from 0 to 365 days.

Audit folder location

The path tells where the audit log is stored. You can specify another location by clicking the `Browse` button and navigating to the new folder in your file system.

7.1.3 Audit Logs

You view the audit logs with a web browser interface. To be able to access the resources, exclude the local host (127.0.0.1) from those hosts for which a proxy server is used. To view the logs, click `Audit Logs` sub menu item under `Auditing` in the SSH Sentinel main menu.

The interface first asks for your user name and password which are set on the *Audit Options* dialog. Having typed those correctly the main page with general information opens. Since you can grant the rights to view the logs to another users on remote hosts, the local host is identified on the first page by the IP address and the DNS name. Click `View audit logs` link to see a listing on available logs and select the one you want to investigate.

The system creates an audit log every day, even if there were no events to log, in which case the log is empty. A new event is written in the log each time a rule that is being audited is applied.

You can reduce the number of events shown by filtering the events by time and by the remote host. The listing shows the following information on each event:

Time

The logging time.

Event

The type of the event. If the event was triggered by an IPSec rule, the type of the event is *trigger*. The IPSec negotiations cause events *Phase1 / Phase 2 succeeded / failed* . If the event was triggered by a filter rule, the event is, depending of the action, *bypass*, *drop* or *reject*. If it was the default response rule that bypassed the traffic, the event is *allow*. In addition, there are various event caused by the IKE engine, like *delete payload received* or *invalid payload length* most often reporting a problem.

Rule/Source

If the event was caused by a rule that was applied, the group of the rule (e.g. pre-IPSec filter, secured connection) is shown in this field. Click to see the details of the rule, provided the rule has not been deleted or altered. But the event can also be triggered by the IKE engine, in which case the filed contains *IKE*.

Local

The IP address of the local host. Click to see details.

Direction

The arrow denotes the direction of the network traffic.

Remote

The IP address of the remote host. Click to see the details.

Protocol

The traffic protocol.

Count

The log is written every five seconds. If during that time, the same event is detected multiple times, the count is increased rather that the event written several times.

7.2 IKE Log Window

To detect and study problems in establishing connections to remote hosts, use the SSH Sentinel *IKE Log Window*. When the utility is set on, the window displays information on internet key exchange negotiations. The information can also be written into a file.

To open the *IKE Log* window, select `Open Log Window` in the SSH Sentinel system tray icon menu.

The amount of information you receive depends on your choice of logging level. The available levels are:

- If set *off*, no information will be logged.
- On the *low* level, you will get information about the success or failure of the negotiation. In case of a successful negotiation, the parameters established are shown. If the negotiation fails, you will get a rough idea of the reason.
- On the *moderate* level, more detailed information about the negotiation will be displayed. The moderate level is usually suitable for finding the reason for a failed negotiation.
- The *detailed* level gives you all the available information. In most cases, the detailed level should not be used because of the excessive amount of messages. Using the detailed level will also slow down the negotiations. However, this setting is useful if you need to know everything that is going on during a negotiation.

To write the messages to a file, mark the *Log to file* check box and browse for the file where to write the messages. Logging starts when you click the *Set* button. If you want to change the file while logging, just find the new file and click the *Set* button again.

If you close the log window while logging to a file, the system asks if you wish to keep on writing the messages to the file. If you choose to continue logging, it will not stop until you specifically turn logging off.

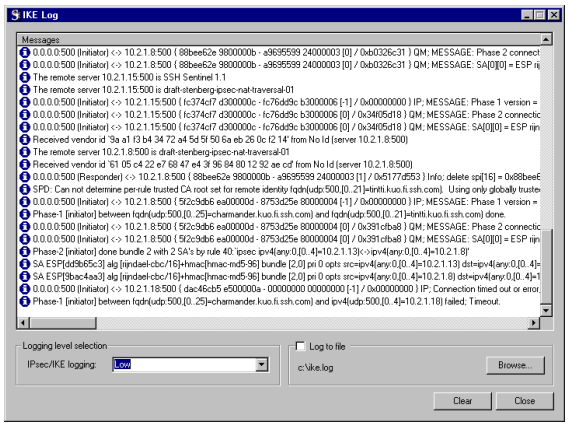

Figure 7.2 The IKE Log window.

7.3 Connection Diagnostics

You can test a connection using the `Diagnostics` button on the *Security Policy* sheet. The button is available in connection to secured connections and virtual private network connections. To run the diagnostics, select the rule and click the button. Alternatively, you can select `Diagnostics` in the menu that opens with the right mouse button.

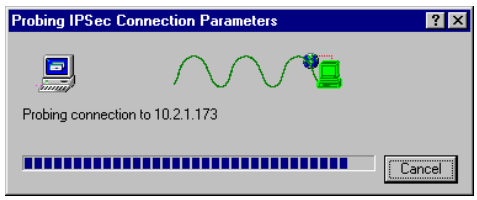

Figure 7.3 Probing the connection parameters.

When you run the connection diagnostics, the system actually probes the connection parameters. While probing, the connection the dialog box seen in Figure 7.3. The connection is negotiated and the connection parameters agreed on with the other end are established. The probing may change the parameters that you have initially set on the rule. The parameters you set are used as initial proposals, however, if the other does not support, say, you're first choice of encryption algorithm, then some other algorithm is chosen.

You are informed if the probing was successful. The dialog box shown after having successfully probed the parameters is shown in Figure 7.4.

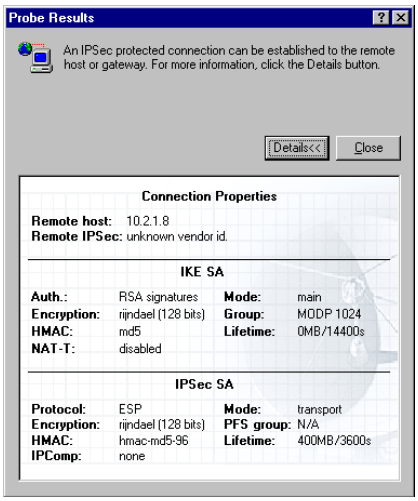

Figure 7.4 The probe results.

There are plenty of reasons why probing might fail, the most obvious being that the computer host in the other end is not up and running. By probing the connection you can easily verify if the other end is up.

7.4 Statistics

The SSH Sentinel Statistics window shows you information on data traffic to and from your host. You can view the established security associations, the data packets transmitted and the errors detected in the data traffic, for example.

You can open the *Statistics* window from the *SSH Sentinel tray icon menu*: Click the *tray icon* with the right mouse button and, keeping the button down, select the menu item `View Statistics`.

There are two sheets on the *SSH Sentinel Statistics* window that opens: *Security Associations* and *IPSec Statistics*.

7.4.1 Security Associations

The *Security Associations* sheet, seen in Figure 7.5, shows the established and currently valid security associations.

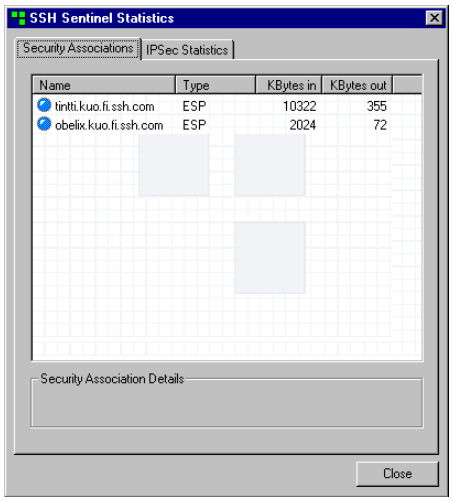

Figure 7.5 The established security associations are displayed on the Statistics window.

Name

The IP address or DNS name of the other end.

Type

The type of the security association, *ESP* or *ESP+IPComp*

KBytes in

The amount of data received.

KBytes out

The amount of data transmitted.

7.4.2 IPSec Statistics

On the *IPSec Statistics* sheet, shown in Figure 7.6 you can monitor the data traffic on your computer host.

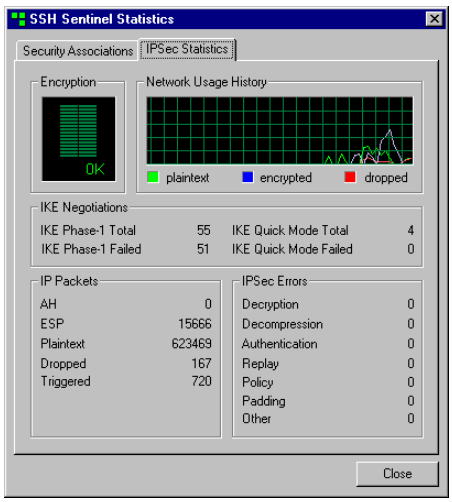

Figure 7.6 The IPSec statistics.

Encryption

The graphics indicate the data throughput in kilobytes.

Network Usage History

The diagram shows the history of data traffic. It distinguishes between plaintext packets, encrypted packets and dropped packets.

IKE Negotiations

IKE Phase-1 Total

The total number of first phases of the Internet Key Exchange negotiations started.

IKE Phase-1 Failed

The number of those Internet Key Exchange negotiations first phases that have ended in failure.

IKE Quick Mode Total

The total number of Internet Key Exchange negotiations started in quick mode.

IKE Quick Mode Failed

The number of the Internet Key Exchange negotiations in quick mode that have ended in failure.

IP Packets

In this box, you can see the number of IP packets in transmission. There are separate figures for *AH*, *ESP* and *plaintext* packets as well as for packets *dropped* by the filter rules and *triggered* by the IPSec engine.

IPSec Errors

With the information presented in this box, you can track errors in the data packets received. There are separate figures for errors in decrypting, decompressing, authenticating and padding the data packet as well as for situations where data replaying is suspected and policy errors. Errors that fall beyond these categories are classified as *other* errors.

Decryption and *decompression* errors are simply errors that occur when a data packet is decrypted or decompressed, respectively. An *authentication* error occurs when the origin of the data packet can not be authenticated, that is, you don't trust the other end's certificate, for example. *Replay* errors are associated with a well-known attack type in the Internet: some hostile party transmits the same data packets over and over again. When there is something wrong with the padding inserted to a data packet that would otherwise have been short, a *padding* error.

Chapter 8 **Glossary**

This glossary contains definitions of special terms and abbreviations used in this document.

access control A security measure that prevents unauthorized use of resources. In the IPSec context, the resources to which access is being controlled are usually the computing cycles of a host, the data stored in it, the network behind a security gateway, or the bandwidth on that network.

Authentication Header (AH) An upper level header located between the IP header and the payload within an IP packet. Typically, an AH includes an Integrity Check Value (ICV) of the transfer-independent contents of the IP packet. The exact nature of the checksum depends on the transformation used. An AH is used to ensure the integrity of the whole IP packet, including both the payload and the IP header. It does not provide data confidentiality. The AH transformation is defined in RFC 2402.

Address Resolution Protocol (ARP) The ARP protocol is used in Ethernet networks to find out the physical link-layer address (Ethernet address) of some other host on the network. The protocol is defined in RFC 826.

ASN.1 (Abstract Syntax Notation) Standard notation for explaining complicated data structures. The ASN.1 is used in X.500 directory to describe the directory objects. Further X.509 uses ASN.1 to describe certificates and CRLs. There are many ways to encode the ASN.1 described structures to binary, most common methods are BER and DER. The ASN.1 is described in ITU-T Recommendation X.680.

authentication The verification of the identity of a person or process. In a communications system, authentication verifies that messages come from their stated source. In this documentation, this term is used to refer to the combination of two nominally distinct security services, data authentication and connectionless integrity.

availability A security service that addresses the security concerns engendered by attacks on networks that deny or degrade service. For example, in the IPSec context, the use of replay prevention mechanisms in AH and ESP support availability.

Border Gateway Protocol (BGP) A routing protocol that is normally used between independent Internet Service Providers. The protocol is defined in RFC 1771.

block cipher A representative of symmetric (secret key) encryption algorithms that encrypts a fixed length block of plaintext (for example, 64 bits) at a time. With a block cipher, the same plaintext block will always encrypt to the same ciphertext block, under the same key.

Blowfish A block cipher with a block size of 64 bits and a key length of up to 448 bits.

Certificate Authority (CA) An entity that attests to the identity of a person or an organization. A CA can be an external company that offers certificate services or it can be an internal organization such as a corporate Management Information System (MIS) department. The chief function of the CA is to verify the identity of entities and issue digital certificates attesting to that identity.

CAST-128 A block cipher with a block size of 64 bits and a key length of up to 128 bits. CAST-128 is believed to be very strong. See RFC 2144 for more information.

Cipher Block Chaining (CBC) A way of using a block cipher. The basic idea is that the previous ciphertext block is XORed to the next block before encryption, thus making two identical plaintext blocks encrypt to different cipher texts. For more information, please see e.g. Bruce Schneier: "Applied Cryptography".

certificate A digital document which is used for verifying the identity of the other end of the transmission. Any type of address including domain name, IP, and email addresses can be authenticated using certificates. Current SSH products use X.509 certificates.

certificate enrollment Certificate enrollment is a process where an end entity requests certification for a key pair. In certificate enrollment the end entity communicates with the Certificate Authority (CA) or Registration Authority (RA).

certificate extension An optional field in X.509v3< certificate providing some further information of the usage or applicability of the certificate in a certain PKI.

certificate request A certificate request contains at least the public key and some identity information, and it is signed with the private key. Certificate requests are generated by end entities

or RAs and sent to the CA. If allowed by the certification policy of the CA, a certificate can then be issued.

certification policy A certification policy is a named set of rules that indicates the applicability of a certificate to a particular community.

ciphertext Text which has been encrypted by an encryption system. The opposite is plaintext.

Certificate Management Protocol (CMP) is a protocol defining the online interactions between the end entities and the certification authority in PKI. It is written by PKIX working group of IETF and is specified in document RFC 2510.

confidentiality A security service that protects data from unauthorized disclosure. Usually, unauthorized disclosure of application level data is the primary concern, but the disclosure of the external characteristics of communication can also be a concern in some circumstances. The traffic flow confidentiality service addresses this latter concern by concealing source and destination addresses, message length, or frequency of communication. In the IPSec context, using ESP in tunnel mode, especially at a security gateway, can provide some level of traffic flow confidentiality. See also traffic analysis.

connectionless integrity A service that detects the modification of an individualIP packet, regardless of the ordering of the IP packet in a stream of traffic.

Certificate Revocation List (CRL) Usual hierarchical certificate system is based on concept of a CA (Certificate Authority). A CA is a trusted party, and has a trusted certificate. The certificates the CA has issued, the end-user certificates, have finite validity period. Nevertheless, it happens that some certificates may need to be revoked before the end of the validity period, thus a frequently supplied list called CRL is issued by the CA. The CRL is a basic tool in X.509 to revoke certificates before their validity period has ended. See for more information.

Certificate Request Message Format (CRMF) is used as a request format in Certificate Management Protocol (CMP). CRMF is a PKIX specification (RFC 2511<).

cross certification Cross certification is a trust model where two certification authorities certify each other. It allows end entities in different certificate hierarchies to verify each other's certificates.

cryptology The branch of mathematics that studies the mathematical foundations of cryptographic methods.

DES and 3DES Data Encryption Standard, defined by the U.S. government. It was created in the 1970s by IBM assisted by the agency that is nowadays called NSA. Based on Horst Feistel's ideas, the team of scientists at IBM devised a cipher that has influenced the science of cryptology. The controversy around DES key length and design issues has developed many variants of the original algorithm. The 3DES (or triple-DES) is the most accepted. Most of what is known about block ciphers is due to analysis of DES.

Diffie-Hellman key exchange A method for key exchange between two parties. This method can be used to generate an unbiased secret key over an insecure medium. The method has many variants. A well known attack called the man-in-the-middle attack forces the use of digital signatures, or other means of authentication, with Diffie-Hellman protocol.

digital signature By encrypting a digest of a message with the private key, authentication can later be performed by applying the public key to an encrypted digest (digital signature) and comparing the result to the digest of the message.

Distinguished Name (DN) A distinguished name belongs to the X.500 directory terminology. It declares a name that can be distinguished from other names in the directory. In that sense that name needs not be unique. Often these names are seen encoded using the LDAP format e.g."CN=John Doe, O=Some Organization, C=US", however, the actual names are ASN.1 objects.

domain name A domain name is a textual name for an Internet host, e.g. `www.ssh.com`. The Domain Name System (DNS) infrastructure is used to map domain names to IP addresses. See RFC 1035 for more information.

Denial of Service (DoS) Denotes attacks that do not cause a security violation per se, but harm the availability of a service. For example, if an attacker sends lots of forged packets to an SSH IPSEC VPNhost, they may degrade the performance of the host. One of the design goals in the SSH IPSEC architecture has been to minimize the consequences of Denial of Service attacks.

Digital Signature Algorithm (DSA) DSA is a public key algorithm for digital signatures. For more information, see e.g. Bruce Schneier: "Applied Cryptography". See DSS.

Digital Signature Standard (DSS) The U.S. government digital signature standard. It is a standard for digital signatures using the DSA public key algorithm and the SHA hash algorithm.

encryption A security mechanism used for the transformation of data from an intelligible form (plaintext) into an unintelligible form (ciphertext), to provide confidentiality. The inverse transformation process is properly known as designated decryption, but encryption is often used to in a generic way refer to both processes.

end entity A human user or an application to whom a certificate is issued. The end entity has also the private key counterpart of the public key in the certificate.

Encapsulating Security Payload (ESP) An upper level IP header that denotes that the contents of the payload are encrypted and possibly also otherwise protected. An ESP may appear after the IP header, after an ESP header or theoretically also elsewhere within an IP packet. An ESP only protects the contents of the payload, not any associated header. Therefore it is possible, for example, to change any field in the header of the IP packet carrying an ESP without causing a security violation. The contents of the ESP header are unknown to anyone not possessing information about the transformation and SA needed to recover the protected data. An ESP may also contain integrity protection. The ESP protocol is defined in RFC 2406.

Ethernet Ethernet is the most widely used a Local Area Network< (LAN</) type in office networks. In Ethernet, each workstation has a unique 48-bit address assigned by the network adapter manufacturer. The ARP protocol is used to convert between Ethernet addresses and IP addresses. RFC 894 defines how to transmit IP packets over Ethernet.

firewall A node located on the perimeter of an administrative domain that implements the security policy of the domain. A firewall usually performs address and port-based packet filtering and usually has proxy servers for e-mail and other services.

Hash Message Authentication Code (HMAC) A secret key authentication algorithm. Data integrity and data origin authentication as provided by HMAC depend on the scope of the distribution of the secret key. If only the source and destination know the HMAC key, this provides both data origin authentication and data integrity for packets sent between the two parties. If the HMAC is correct, this proves that it must have been added by the source.

host Any node that does not forward packets that are not addressed to the node itself. Generally this term is used to refer to any computer or other computing device connected to an IP-based network.

Hypertext Transfer Protocol (HTTP) HTTP is the protocol used to transfer web pages from a WWW server to the browser. The HTTP client sends requests to the server, and gets some data as a response. HTTP identifies objects on the server using URIs or URLs. For more information, see RFC 2068.

Internet Control Message Protocol (ICMP) ICMP is a message control and error-reporting protocol between a host server and a gateway to the Internet.

Integrity Check Value (ICV) Usually, an HMAC algorithm using either Message Digest 5 (MD5) or SHA-1 hash functions, but possibly also a DES-MAC or HMAC-RIPEMD algorithm. See also integrity.

Internet Engineering Task Force (IETF) An international standards body that has standardized the IP protocol and most of the other successful protocols used on the Internet. The IETF web pages are available at http://www.ietf.org. .

Internet Key Exchange (IKE) The key exchange algorithm used with IPSec. This is a new name for the ISAKMP/Oakley key exchange. In particular, this refers to the resolution draft that specifies which parts of each specification need to be implemented for IPSec use. The IKE protocol is defined in RFC 2409, RFC 2408, and RFC 2407.

integrity A security service that ensures that data modifications are detectable. Integrity services need to match application requirements. IPSec supports two forms of integrity: connectionless integrity and replay prevention. This is in contrast to connection-oriented integrity, which imposes more stringent sequencing requirements on traffic to be able to detect lost or re-ordered messages, for example. Although authentication and integrity services are often cited separately, in practice they are intimately connected and almost always offered together.

Internet Protocol (IP) The network layer for the TCP/IP protocol suite, defined in RFC 791. IP is a connectionless, best-effort packet switching protocol. It provides packet routing, fragmentation, and re-assembly through the data link layer.

IP address A 32-bit number that identifies the devices using the IP protocol. An IP address can be unicast, broadcast, or multicast. Please see RFC 791 for more information.

IP header The part of the IP packet that carries data used on packet routing. The size of this header is 20 bytes, but usually the IP options following this header are also calculated as header. The maximum length of the header is 60 bytes. The header format is defined in RFC 791.

IP packet A self-contained, independent entity of data carrying sufficient information to be routed from the source to the destination computer without reliance on earlier exchanges between this source and destination computer and the transporting network. The Internet Protocol (IP) is defined in RFC 791.

IP payload The part of the IP packet that carries upper level application data.

Internet Protocol Security (IPSec) A protocol suite for protecting IP traffic at packet level defined by the Internet Engineering Task Force (IETF). It can be used for protecting the data transmitted by any service or application that is based on IP. The IPSec protocols are defined in RFC 2401. The RFC 2411 document is a good starting point for reading about it.

IP version 4 (IPv4) This is the current version of IP.

IP version 6 (IPv6) This is a new version of the IP protocol ("next generation" IP). Among other improvements it has an extended address space and better security. It is described in RFC 2460. There is no version five.

Internet Security Association and Key Management Protocol (ISAKMP) A protocol for establishing, negotiating, modifying, and deleting SAs. ISAKMP provides a framework for authentication and key exchange but does not define them. ISAKMP is designed to be key exchange independent. That is, it is designed to support many different key exchanges. The ISAKMP/Oakley combines ISAKMP with the Oakley key exchange. Oakley describes a series of key exchanges called modes and details the services provided by each. For example, perfect forward secrecy for keys, identity protection, and authentication. ISAKMP is a part of the IKE protocol.

key enrollment Key enrollment is an action where a public key gets certified by a Certificate Authority (CA). In this action a client provides the CA with a public key and some additional data in a PKCS-10 certificate requests. The CA signs this key together with additional information with its own private key and returns the signed certificate to the client.

Layer Two Tunneling Protocol (L2TP) L2TP facilitates the tunneling of PPP packets across an intervening network in a way that is as transparent as possible to both end-users and applications. L2TP is defined in RFC 2661.

Lightweight Directory Access Protocol (LDAP) LDAP is a directory access protocol defined by RFC 2251 and RFC 1777 for accessing directories supporting the X.500 models, while not incurring the resource requirements of the X.500 Directory Access Protocol (DAP). This protocol is specifically targeted at management applications and browser applications that provide read/write interactive access to directories. The protocol is carried directly over TCP or other transport, bypassing much of the session/presentation overhead of X.500 DAP.

Message Authentication Code (MAC) A mechanism that provides message integrity by using a secret key cryptographic function.

MD5 A message-digest algorithm. It computes a secure, irreversible, cryptographically strong hash value for a document. A variant called SHA-1 is usually thought more secure. The algorithm is documented in RFC 1321.

Network Address Translation (NAT) A network address translator is a device that is connected to two networks, and translates IP addresses in packets sent across it. Typically, one of the networks will use global addresses, and the other will use local addresses. NATs are becoming increasingly common on the Internet, because the available IP address space is running scarce and NATs help large organizations to avoid renumbering their computers if they e.g. change IP providers. There are two basic types of NAT: IP address translation (described in RFC 1631), and port

translation, which may map multiple IP addresses to a single IP address but with different port numbers. Obviously, only TCP and UDP protocols work over port NAT.

node In this document, a node refers to any system implementing the TCP/IP protocol suite.

packet filtering A method for determining how passing IP packets should be handled. Packet filtering is applied to all IP packets passing the IPSec Engine. Packet filtering may modify the IP packet, pass it intact, or even drop it.

Privacy-Enhanced Mail (PEM) A suite of protocols for encryption, authentication, message integrity, and key management. An IETF standard.

Perfect Forward Secrecy (PFS) Refers to the notion that any single key being compromised will permit access to only data protected by that single key. In order for PFS to exist, the key used to protect transmission of data must not be used to derive any additional keys. If the key used to protect transmission of data was derived from some other keying material, that material must not be used to derive any more keys.

PKCS The PKCS public key cryptography standards are a document series from RSA Laboratories. Some of the most important PKCS standards include PKCS-1 for RSA encryption and signature formats, PKCS-7 for cryptographic message encapsulation, PKCS-10 for certificate requests, and PKCS-11 for a cryptographic token interface commonly used with smart cards.

PKCS-1 This standard describes the usage of RSA algorithm in encryption and digital signatures. It contains explicit suggestions for encodings of keys and algorithm input formatting.

PKCS-7 This standard describes the general syntax for data that may have cryptography applied to it. This data includes digital signatures and recursive digital envelope encodings for cryptographic objects.

PKCS-10 This standard describes the certificate requests. The certificate requests are commonly used means of acquiring certificates.

PKCS-11 The standard describing CryptoKi, which is an interface for cryptographic devices. (For example, smart cards and cryptographic accelerators.)

Public Key Infrastructure (PKI) PKI consists of end entities possessing key pairs, certification authorities, certificate repositories (directories), and all the other software, components, and entities required when utilizing public key cryptography.

PKIX The IETF public key infrastructure standard based on X.509.

plaintext Text which has not been encrypted. The opposite is ciphertext.

policy The purpose of an IPSec Security Policy is to decide how an organization is going to protect itself. The policy will generally require two parts: a general policy and specific rules (e.g. system specific policy). The general policy sets the overall approach to Security. The rules define what is and what is not allowed. In this document the term Policy is used typically when referring the later. The Security Policy describes how data is protected, which traffic is allowed or denied, and who is able to use the network resources.

Point-to-Point Protocol (PPP) PPP provides a standard method for transporting multi-protocol datagrams over point-to-point links. PPP is defined in RFC 1661.

preshared secret Preshared secret is an authentication method in IKE. In this method two peers have configured a shared password that is used to authenticate the endpoints by means of encryption (A can decipher packet B has encrypted, therefore A knows B knows the same secret it knows and vice versa). This authentication method scales badly and is useable for very limited number of hosts. For large set of hosts certificate based authentication should be used.

private key In public key cryptography the private key is only known to the holder, and it can be used to sign and decrypt messages.

proxy Proxy is a cache server that acts as a firewall, protecting the local network. It allows an application inside the proxy to access resources on the global Internet.

public key In public key cryptography the public key, which is included in the certificate, can be used to verify signatures and encrypt messages.

public key cryptography In contrast to symmetric ciphers with just one cipher key, in public key cryptography each person or host has two keys. One is the private key which is used for signing outgoing messages and decrypting incoming messages, the other is the public key which is used by others to confirm the authenticity of a signed message coming from that person and for encrypting messages addressed to that person. The private key must not be available to anyone but its owner, but the public key is spread via trusted channels to anyone.

Registration Authority (RA) An entity that may perform some tasks such as key generation and certificate enrollment on behalf of the end entity.

RFC Request For Comments; a document of Internet Society under standardization.

Rijndael Rijndael is a block cipher, designed by Joan Daemen and Vincent Rijmen. The cipher has a variable block length and key length. It currently specifies how to use keys with a length of 128, 192, or 256 bits to encrypt blocks with a length of 128, 192 or 256 bits (all nine combinations of key length and block length are possible). Both block length and key length can be extended very easily to multiples of 32 bits. Rijndael can be implemented very efficiently on a wide range of processors and in hardware. The design of Rijndael was strongly influenced by the design of the block cipher Square.

router A node that forwards packets not addressed to itself. Requirements for a router are defined in RFC 1812.

RSA A public key encryption and digital signature algorithm. It was invented by Ron Rivest, Adi Shamir, and Leonard Adleman. For more information, see e.g. Bruce Schneier: "Applied Cryptography". The RSA algorithm was patented by RSA Security, but the patent expired in September 2000.

Security Association (SA) A unidirectional connection created for security purposes. All traffic traversing an SA is provided the same security processing. In the IPSec context, an SA is an internet layer abstraction implemented through the use of an AH or ESP. It contains data controlling how a transformation is applied to an IP packet. The data is determined using specially defined SA management mechanisms. The data may be a result of an automated SA and key negotiation or it may be defined manually. This term is defined in RFC 2401.

Simple Certificate Enrollment Protocol (SCEP) SCEP is developed by Cisco Systems and VeriSign. It is an enrollment protocol supported by Cisco's routers.

security gateway An intermediate system that acts as the communications interface between two networks. The internal subnetworks and hosts served by a security gateway are presumed to be trusted because of shared local security administration. See also trusted subnetwork. The set of hosts and networks on the external side of the security gateway is viewed as not trusted or less trusted. In the IPSec context, a security gateway is the point at which an AH or ESP is implemented in order to serve a set of internal hosts. A security gateway provides security services for these hosts when they communicate with external hosts also employing IPSec either directly or via another security gateway. The term is defined in RFC 2401.

Secure Hash Algorithm (SHA) A United States government standard for a cryptographically strong hash algorithm. See MD5. The SHA-1 algorithm is defined in FIPS PUB 180-1.

SHA-1 The Secure Hash Algorithm version one. The algorithm was designed by NSA, and is part of the U.S. Digital Signature Standard (DSS). This algorithm is considered very good.

smart card A smart card, or an integrated circuit card, is a device for secure identification of users of information systems. Typically smart cards contain a processor that can do a private key operation using a private key on the card, some kind of a file system that can hold certificates, public keys, or other data relevant for the use of the card.

Simple Network Management Protocol (SNMP) A protocol that is commonly used to monitor the status of routers and other network elements. The protocol is defined in RFC 1157.

Security Parameters Index (SPI) An arbitrary value used in combination with a destination address and a security protocol to uniquely identify an SA. The SPI is carried in AH and ESP protocols to enable the receiving system to select the SA under which a received IP packet will be processed. An SPI has only local significance as it is defined by the creator of the SA, which is usually the receiver of the IP packet carrying the SPI. Thus an SPI is generally viewed as an opaque bit string. However, the creator of an SA may choose to interpret the bits in an SPI to facilitate local processing. This term is defined in RFC 2401.

stream cipher A representative of symmetric (secret key) encryption algorithms that encrypt a single bit at a time. With a stream cipher, the same plaintext bit or byte will encrypt to a different bit or byte every time it is encrypted.

Transmission Control Protocol (TCP) A widely used connection-oriented, reliable (but insecure) communications protocol. This is the standard transport protocol used on the Internet. It is defined in RFC 793.

Transport Layer Security (TLS) Transport Layer Security is a protocol providing confidentiality, authentication, and integrity for stream-like connections. It is typically used to secure HTTP connections. The protocol is being standardized by a working group of the IETF.

traffic analysis The analysis of network traffic flow for the purpose of deducing information that is useful to an adversary. For example, frequency of transmission, the identities of the conversing parties, sizes of IP packets, and flow identifiers.

transformation A particular type of change applied to an IP packet. For example, ESP encryption, AH integrity service, and payload compression are transformation types. An SA supplies the keys and other association-specific data to a transformation. The IPSEC transformations are defined in RFC 2401, RFC 2402, RFC 2403, RFC 2404, RFC 2406, and RFC 2405.

trusted subnetwork A subnetwork of hosts and routers that can trust each other not to engage in active or passive attacks. It is also assumed that the underlying communications channel such as a LAN or CAN is not being attacked by any other means.

Twofish A strong and fast block cipher designed by Bruce Schneier. Twofish uses a block size of 128 bits and a key length of up to 256 bits.

User Datagram Protocol (UDP) A datagram-oriented unreliable communications protocol widely used on the Internet. It is a layer over the IP protocol. It is defined in RFC 768.

Universal Resource Identifier (URI) URIs are supposed to identify resources or objects in the world or on the Internet. They are defined in RFC 2396. The most commonly used form of an URI is an URL.

Universal Resource Locators (URL) URLs are used to describe the location of web pages, and are also used in many other contexts. An example of an URL is `http://www.ssh.com/ipsec/index.html`. They are defined in RFC 1738 and RFC 1808. URLs are a special case of URIs.

Virtual Private Network (VPN) The use of encryption in the lower protocol layers to provide a secure connection through an otherwise insecure network, typically the Internet. The encryption may be performed by firewall software or possibly by routers.

X.500 The family of standards defining the X.500 directory. The directory can be used for many applications, such as storing certificates, or information about people. LDAP is often used to access the X.509 directory.

X.509 The ISO/OSI X.509 standard specifies certificate and CRL formats. It is widely used in a variety of applications. Version 3 of the standard (X.509v3) added support for IP addresses and domain names in certificates. It is closely related to the PKIX standard.

Chapter 9 **Index**

