ADDRESS TRANSLATION FUNCTIONS

Many users create their networks using private IP addresses (see "A Basic Introduction to IP Addresses" on page 101). However, these IP addresses must never appear on the Internet and if they did, they would be discarded. Therefore, networks using these addresses require their network address to be translated when they connect to the Internet.

InterJak enables three forms of Address Translation: NAPT, Exported services and NAT.

NETWORK ADDRESS AND PORT TRANSLATION

NAPT hides the local host IP addresses by translating both the source IP address and the port numbers contained in the IP header. The InterJak provides the packets with a new source IP address, which is the same as its own IP address as seen from the Internet, and a fictitious port number, which is in the range 61,000 to 65,096. Refer to Figure 17.

Figure 17 Illustrating NAPT in the InterJak.

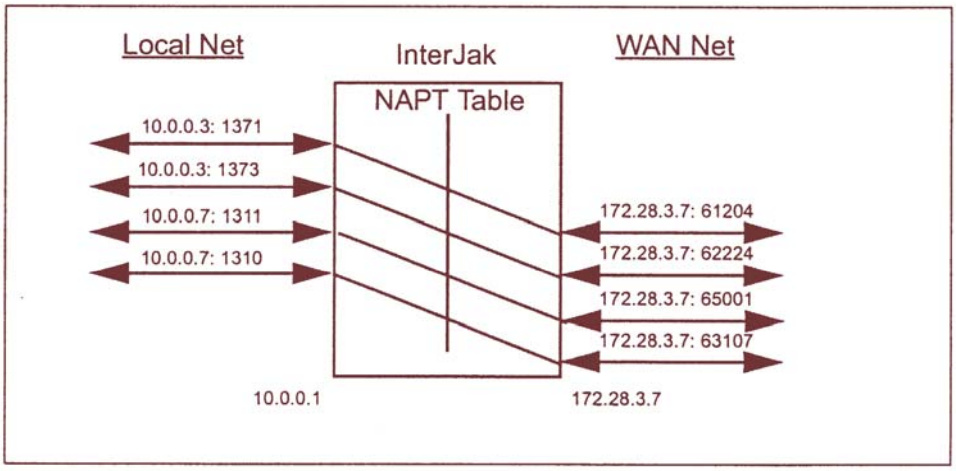

These translations are cached in the InterJak. Any replies to the hosts' enquiries contain the NAPT IP address and port numbers. These replies are received by the InterJak, which then uses the cached data to translate the addresses back to the hosts' local IP addresses and port numbers.

NOTE *In this situation, all communication must be initiated by the local hosts. If, for some reason, the cached data is lost, e.g. if the InterJak is rebooted, then the communications must be initiated again by the local hosts.*

NOTE *If a TCP connection is idle, i.e. no data is transmitted on the connection, the InterJak considers the connection dead and removes its translation after 15 minutes. This effectively breaks the connection. You can configure the length of this timeout period on the NAPT setup page.*

To configure Network Address and Port Translation:

1. Go to the **`Services:Address Translation (NAT/NAPT)`** dialog in the Web Manager.
2. Click on the **`Edit`** button in the **`Network Address and Port Translation (NAPT)`** section.
3. Click the appropriate check boxes for enabling NAPT.
   NAPT can be enabled for hosts on the LAN and on the DMZ separately. For a description of LAN and DMZ, see "Firewall" on page 121.
4. Edit the Established TCP connection timeout value, if needed.
5. Click **`Apply`** to implement changes and return to the **`Services:Address Translation (NAT/NAPT)`** dialog.

When NAPT is enabled, any workstations or servers behind the InterJak that host services for the outside become invisible.

EXPORTED SERVICES

When the InterJak has NAPT enabled, any services for the outside hosted by a local host or server behind the InterJak become invisible. However, the InterJak contains a feature that permits certain services to remain visible. This feature is called Exported Services.

In order to make a local service visible to the outside, the InterJak must be configured with the local IP address of the host providing the service. The following services can be exported:

- Web browsing (HTTP)
- E-mail retrieval (POP3)
- E-mail delivery (SMTP)
- Telnet

- FTP
- PPTP server
- H.323

If a service not listed above need to be exported, the TCP or UDP port number to be exported must be defined together with the IP addresses. The InterJak allows 10 TCP and UDP port ranges to be exported. The exported services is available on the WAN using the InterJak's WAN IP address. Refer to Figure 18.

Figure 18 Illustrating Exported Services in the InterJak.

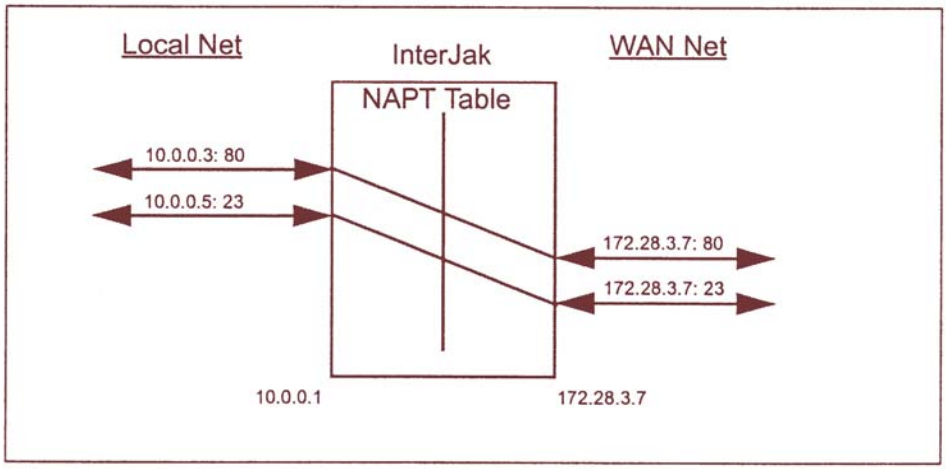

To configure Exported Services:

1. Go to the **`Services:Address Translation (NAT/NAPT)`** dialog in the Web Manager.
2. Click on the **`Edit`** button in the **`Network Address and Port Translation (NAPT)`** section.
3. Enter the Local IP address for the exported services.
4. If needed, enter the TCP or UDP port number together with the local IP address for exporting services not in the list.
5. Click **`Apply`** to implement changes and return to the Services:Address Translation (NAT/NAPT) dialog.

NOTE *If Exported Services is enabled for HTTP and Telnet, then management access to the InterJak from the WAN side will be lost. To prevent this from happening, alternative port numbers must be defined for these two services in the InterJak.*

This is done in the System:Management dialog and is described in "Using Non-standard TCP Ports" on page 117.

NETWORK ADDRESS TRANSLATION (NAT)

This is similar to Exported Services, where only the IP address is translated. In this situation a public IP address must be provided for each local host or server, which must be accessible from the WAN. Each public IP address must be purchased from your ISP. The outgoing services then go out using the correct port number and the new IP address. Replies and incoming requests come in to this external IP address with the correct port number. The IP address is then translated to the specified host address. This is illustrated in Figure 19

Figure 19 Illustrating NAT in the InterJak.

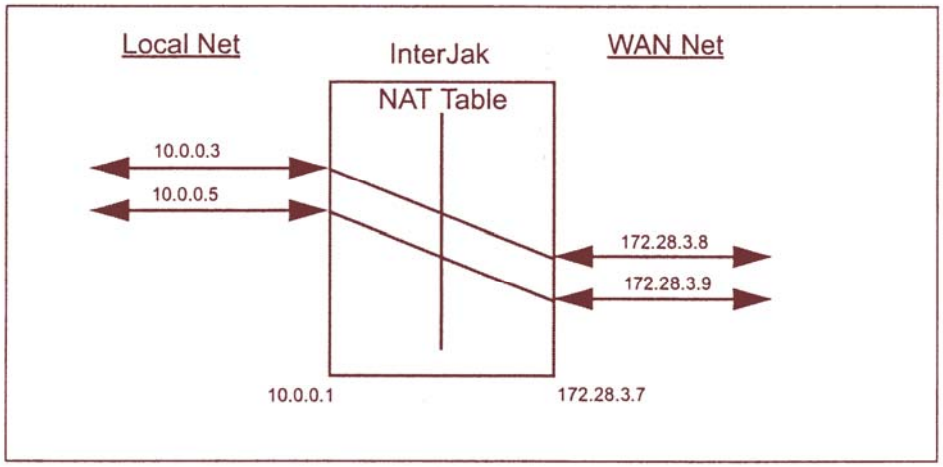

The translated host is protected by a specific firewall which must be set up when the translation is defined. Four services are specially defined:

- Web browsing (HTTP)
- E-mail retrieval (POP3)
- E-mail delivery (SMTP)
- Telnet

These services can easily be enabled using the check boxes in the Web interface. Further services can be enabled by creating a service list. This list can either be a list of services allowed, in which case services not listed are blocked. The list can also be defined as a list of services not allowed, in which case the services not

listed are allowed. Note that having an empty list of services not allowed will allow any access to the translated host.

For a description of services, see "Firewall" on page 121.

The InterJak supports up to 50 Network Address Translations. Each translated host requires an unique external IP address, different from the InterJak's IP address. The external IP address must be visible to the Internet.

To configure Network Address Translation:

1 Go to the **Services:Address Translation (NAT/NAPT)** dialog in the Web Manager.

2 To add a new NAT, click on the **Add NAT rule** button in the **Network Address Translation (NAT) Rules** section. To edit an existing NAT rule click on the Rule Name.

3 Enable the rule and enter the Public and Private IP addresses.

4 If needed, enable one or more of the known services by clicking the corresponding check box.

5 Select the type of the firewall and edit the service list as needed.

6 Click **Apply** to implement changes and return to the Services:Address Translation (NAT/NAPT) dialog.

Supported Protocols

The address translation features in the InterJak support all data protocols that do not add IP address and port number information to the data fields. In addition to this, the InterJak contains software modules that enable it to support certain protocols that do insert IP address and port numbers in the data:

- CuSeeMe
- FTP
- RealAudio

Using Non-standard TCP Ports

Although the default settings for the InterJak use the standard TCP port numbers 80 (for HTTP) and 23 (for Telnet) for communication, alternative (non-standard) TCP port numbers can be assigned by the Administrator.

These alternative TCP ports have two major functions:

- Prevents remote administrators from losing contact with the InterJak when all HTTP and Telnet traffic using standard port numbers is exported to specified hosts.
- Gives an extra level of security as the management ports are not obvious on the net. However, do not rely only on this security as the new ports can easily be found using port scanners.

To configure Non-standard TCP ports

1 Go to the **System:Management** dialog in the Web Manager.

2 Enter the port numbers to be used for HTTP or telnet.

3 Click **Apply** to implement changes and return to the **System:Basic** dialog.

APPLICATION EXAMPLES

An example of Network Address and Port translation using exported services is shown below in Figure 20:

Figure 20 Network Address and Port Translation example

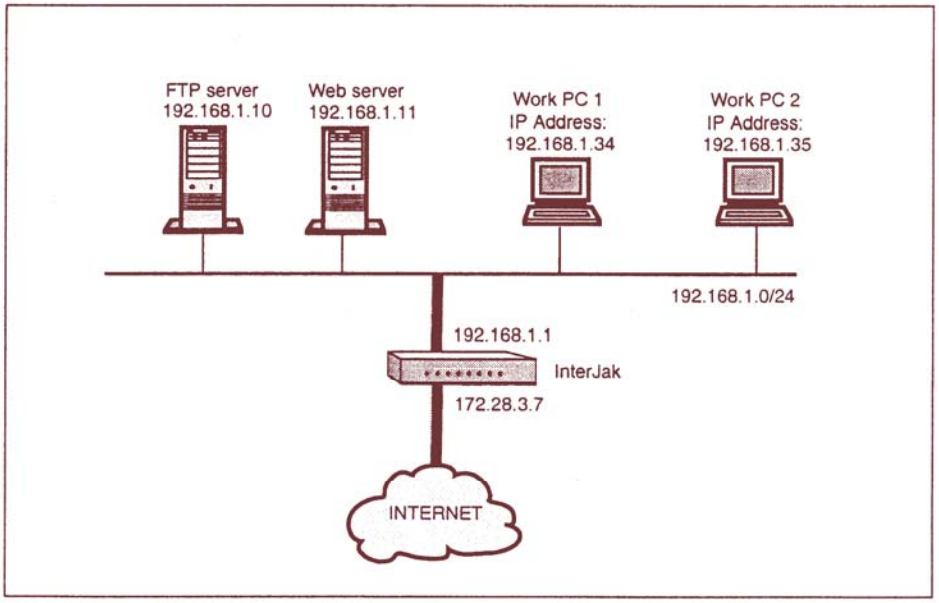

To obtain this configuration do the following:

1. Go to the **Services:Address Translation (NAT/NAPT)** dialog and click **edit** in the in the **Network Address and Port Translation (NAPT)** section.
2. Enable NAPT between LAN and WAN (it is already enabled by default).
3. Enter 192.168.1.11 in the **Web browsing (HTTP) to local host.**
4. Enter 192.168.1.10 in the **FTP to local host.**
5. Click **Apply** to implement changes and return to the **Services:Address Translation (NAT/NAPT)** dialog.

The local FTP and Web servers are now available for the Internet at the IP address 172.28.3.7. The work PC's can access the Internet using NAPT. On the Internet the work PC's will be identified using address 172.28.3.7. Note that access to the InterJak's Web management is disabled for the internet unless a non-standard TCP port has been defined. Access to the InterJak's Web management from the local PC's can be obtained using the address 192.168.1.1.

An example of Network Address Translation is shown below in Figure 21:

Figure 21 Network Address Translation example

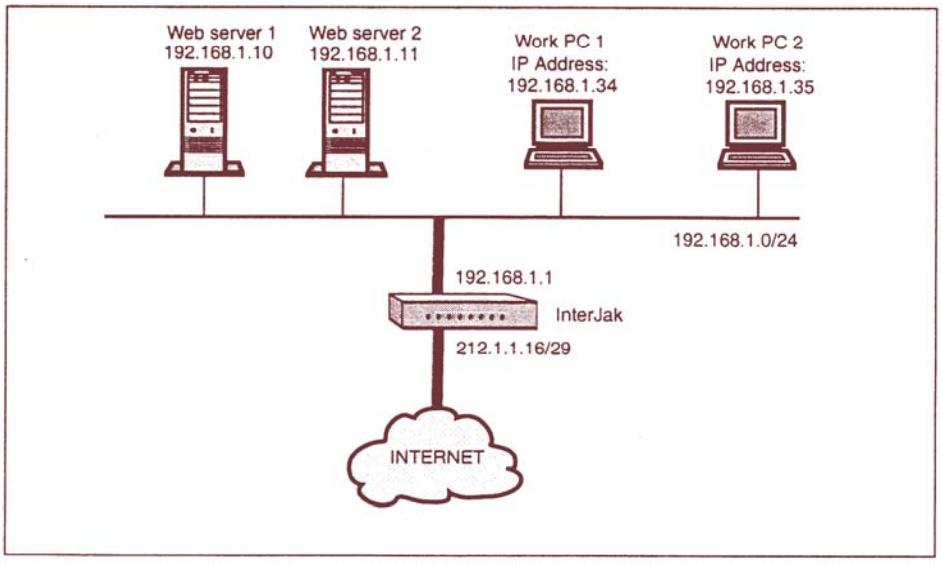

In this case, a 6 hosts subnet has been obtained from the ISP, thus the IP addresses from 212.1.1.17 to 212.1.1.22 are available. We will use 212.1.1.17 for the InterJak itself, 212.1.1.18 for Web server 1 and 212.1.1.19 for Web server 2. The

rest of the IP addresses are available for future expansion. The Work PC's should still use NAPT when accessing the internet.

To obtain this configuration do the following:

1 Go to the **Services:Address Translation (NAT/NAPT)** dialog and click **edit** in the in the **Network Address and Port Translation (NAPT)** section.

2 Enable NAPT between LAN and WAN (it is already enabled by default)

3 Make sure that all exported services are left blank.

4 Click **Apply** to implement changes and return to the **Services:Address Translation (NAT/NAPT)** dialog.

5 Click on the **Add NAT rule** button in the **Network Address Translation (NAT) Rules** section.

6 Enable the translation

7 Enter 212.1.1.18 as **public address**.

8 Enter 192.168.1.10 as **private address**.

9 **Enable web browsing (HTTP)** by checking the check box.

10 **Select type of firewall** as **Only permit access to listed services**

11 Leave the **Service list** empty.

12 Click **Apply** to implement changes and return to the **Services:Address Translation (NAT/NAPT)** dialog.

13 Redo steps 5 to 12 for web server 2 using 212.1.1.19 as Public address and 192.168.1.11 as private address.

Web Server 1 is now available on the Internet using address 212.1.1.18 and Web Server 2 is available using address 212.1.1.19. The InterJak's Web management is available using address 212.1.1.17, provided it is enabled in the firewall.

The Work PCs can access the Internet using NAPT and will be identified on the Internet using IP address 212.1.1.17.

To allow a NAT'ed host to be accessed by NetMeeting, its firewall must be opened for all UDP and TCP ports above 1024. See the "Opening the Firewall for H.323 NetMeeting" on page 135 for further details.

2 Create a new firewall rule using the **`Services:Firewall:Add Firewall`** Rule.

   a Select the new NetMeeting service.

   b Define the source IP address as the external IP addresses allowed using NetMeeting against local hosts.

   c Define the destination IP address as the local IP addresses allowed using NetMeetings.

   d Set Action to Accept.

   e Select **`Apply`** to implement the new rule.

NOTE *This setting leaves the firewall wide open, so it is very important that the source and destination addresses above are limited as far as possible.*

CONNECTING REMOTE SITES AND CLIENTS USING VPN (IPSEC)

When confidential data needs to be transferred between sites, VPN using IPsec encryption can provide an attractive solution to an expensive leased line. A Virtual Private Network (VPN) is a "tunnel" through a public network, for example the Internet. VPN provides the following security for data passing through the tunnel:

- Confidentiality/privacy: The contents of the data packets remain confidential by using encryption, e.g. DES or triple-DES.
- Integrity: Ensures that any alteration to data does not go undetected. This is achieved using checksums.
- Data origin authentication: Uses shared secrets to ensure that traffic is indeed coming from the remote gateway you think it is.

A VPN tunnel is set up between two devices each connected to the Internet.

VPN MODELS — SITE-TO-SITE

A Site-to-Site IPsec tunnel can be configured to connect one or more IP-subnets "behind" one device with one or more subnets "behind" the other device. An IPsec tunnel is tied to the subnets it links together.

Figure 25 below illustrates an IPsec tunnel between the subnets 192.168.1.0/255.255.255.0, 192.168.2.0/255.255.255.0 and 192.168.3.0/255.255.255.0. In this setup, devices on the 192.168.1.x and 192.168.2.x subnet can access devices on the 192.168.3.x subnet and vice-versa.

But traffic from the 192.168.4.x network cannot travel through the IPsec tunnel, even though this network also resides behind an InterJak.

When traffic is "in" the tunnel it is encrypted, and when one InterJak receives an IPsec packet, it can verify that it is sent by the other tunnel end-point.

Figure 25 Site-to-Site IPsec tunnel

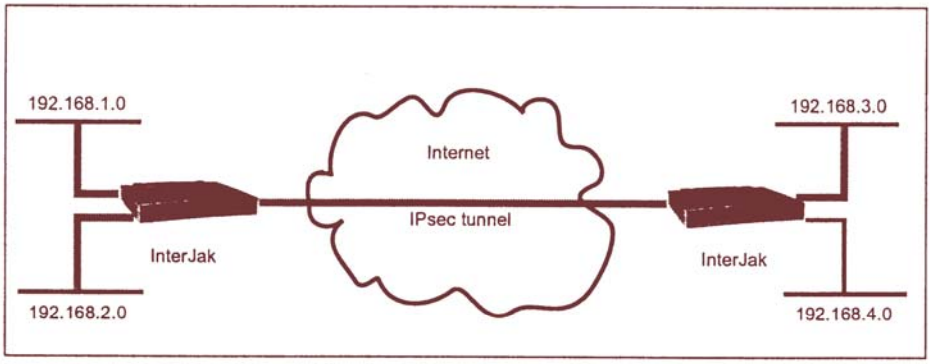

The InterJak currently supports up to four simultaneous Site-to-Site connections.

VPN MODELS — CLIENT-TO-SITE

Client-to-Site VPN enables Windows PC's with an IPsec client installed to access a subnet behind the InterJak through an IPsec tunnel. The tunnel is established by the PC when the user tries to access data through it, and is removed when the user is no longer using it. Refer to Figure 26:

Figure 26 Client-to-Site IPsec tunnel

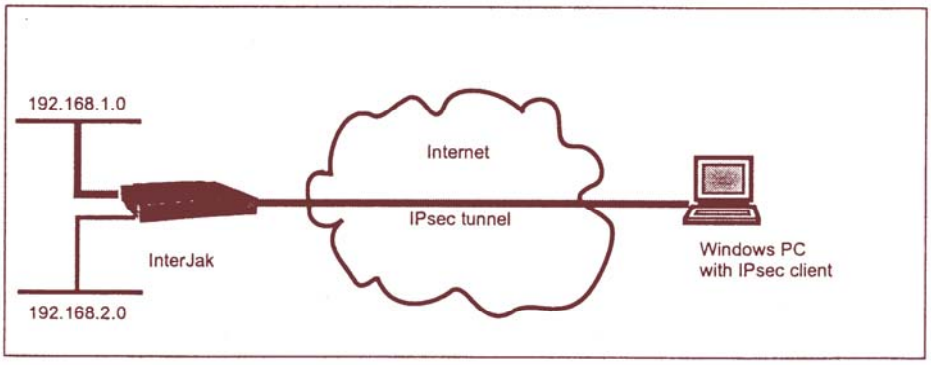

The InterJak currently supports up to four simultaneous Client-to-Site connections.

UNDERSTANDING IPSEC BASED VPN

A VPN tunnel is created using two separate negotiations. The first negotiation can be either Main Mode or Aggressive Mode. Main Mode provides identity protection for the host initiating the IPsec session, but takes slightly longer to complete. Aggressive Mode provides no identity protection, but is quicker. The second negotiation, Quick Mode, uses the results of Main Mode or Aggressive Mode to create a key for bulk data traffic. An overview of these negotiations, including the terminology used, is presented in this section.

The protocol IPsec uses the term Security Association (**SA**) as the information needed in order for one VPN gateway to send secure traffic to the other VPN gateway. In effect, an SA is the selection of encryption protocols and keys needed for secure communication. An SA typically has a limited lifetime and is one-way; for peers to communicate and create a VPN tunnel two SAs must exist. The SAs are created between the VPN gateways trying to establish a VPN tunnel. They are negotiated using the Internet Key Exchange (**IKE**) protocol (see RFC-2409). In the process of establishing a VPN tunnel, a number of SAs are created. There are SAs used by the VPN gateways to communicate securely and SAs used for the transfer of data across the VPN tunnel. For this reason, the negotiation and establishment of a VPN connection consist of two separate negotiations or phases. These are described in the following subsections.

Phase I - The VPN gateways establish SA for secure communication

The VPN gateways can negotiate the SA(s) in this phase using Main mode or using Aggressive mode. In order to explain the difference between Main and Aggressive mode, a high level overview of the negotiations of each mode is presented here.

Main mode

When doing pre-shared key authentication, Main Mode (RFC 2409, sec. 5.4) is defined as shown in Table 10 below:

Table 10

Initiator		Responder
HDR, SA	---->	
	<----	HDR, SA
HDR, KE, Ni	---->	
	<----	HDR, KE, Nr
HDR*, IDii, HASH_I	---->	
	<----	HDR*, IDir, HASH_R

As indicated in the table, Main mode consists of 3 exchanges:

- In the first exchange, the Initiator proposes one or more encryption and authentication protocols. The Responder selects one, and transmits it to the Initiator.
- In the second exchange, encryption keys are created using the Diffie-Hellman algorithm. The data required for this algorithm to work is transferred in the packets. Additional random information is sent to protect against replay attacks.
- In the third exchange, the identities of the two VPN gateways are exchanged. Since the needed information for encryption has been transferred in exchange 1 and 2, exchange 3 is encrypted. It is also authenticated using the pre-shared key. Due to the encryption of the third exchange Main Mode is providing so-called identity protection of the two VPN gateways.

Aggressive Mode
Aggressive mode with a pre-shared key is defined as shown in Table 11 below:

Table 11

Initiator		Responder
HDR, SA, KE, Ni, IDii	---->	
	<----	HDR, SA, KE, Nr, IDir, HASH_R
HDR, HASH_I	---->	

As indicated in the table, Aggressive mode only consists of 1 (and a half) exchange, in 3 packets, compared with 3 exchanges required by Main Mode.

- In Aggressive Mode, the first packet the initiator sends has the following information:
  - Encryption and authentication algorithms
  - Keying material (again for using the Diffie-Hellman algorithm)
  - Random data for replay attack prevention
  - The identity of the gateway

  Since the responder has no time to select the encryption and authentication algorithms, the initiator can only send one proposal for these. This is different from Main Mode (see Main Mode exchange, above).
- The responder replies with the following information:
  - Encryption and authentication algorithms
  - Keying material (again for using the Diffie-Hellman algorithm)
  - Random data for replay attack prevention
  - The identity of the gateway.

The second packet contains authentication data using the pre-shared key, allowing the Initiator, who receives the packet to verify that the responder has the correct key.

- The Initiator replies with a packet containing authentication data, allowing the Responder, who receives the packet to verify that the Initiator has the correct key.

Which mode to use

The default mode for Phase I negotiations is Main mode, as it provides better security. Aggressive Mode should only be used when required. Aggressive mode may be required when interoperating with 3rd party products.

Phase II - The Phase I SAs Are Used To Create New SA For Data Traffic

The second Phase II SA(s) is negotiated using Quick Mode.

Quick Mode

The Quick Mode negotiation phase is protected within an **IKE SA** negotiated in Phase I. This mode negotiates the SA for the data encryption and manages the key exchange for that IPsec SA. The Quick Mode negotiation is presented in Table 12 below:

Table 12

Initiator		Responder
HDR*, HASH(1), SA, Ni, [, KE] [, IDci, IDcr]	---->	
	<----	HDR*, HASH(2), SA, Nr, [, KE] [, IDci, IDcr]
HDR*, HASH(3)	---->	

The SA(s) used for the bulk traffic (i.e. traffic traveling between the two VPN subnets) is negotiated using Quick Mode. A Quick Mode negotiation consists of 3 packets, all of which are encrypted and authenticated using the SAs established in Phase I (Main or Aggressive) negotiation.

- In the first packet the Initiator sends
  - Proposal for protocol (ESP/AH), encryption and authentication algorithms
  - Random data for replay attack prevention
  - Additional keying material (only if Perfect Forward Secrecy (or PFS) is enabled)
  - The subnets that this tunnel is between.
- The responder replies with similar data to accept the proposal.

- The Initiator responds again with data that proves to the Responder that the first message was not a playback of an earlier exchange.

UNDERSTANDING IPSEC PROTOCOLS

The transfer of bulk data is done by transferring the IP-packets that are to use the VPN tunnel "inside" packets that belong to the two IPsec protocols: ESP or AH. That is, the traffic packets are "tunneled" in ESP or AH packets.

- ESP (Encapsulating Security Payload): Provides both data origin authentication and encryption. This is done by encrypting the data packets and adding additional headers and trailers. (See Figure 27)
- AH (Authentication Header): Provides authentication only. This is done by adding additional headers to the data packets. (See Figure 28)

Figure 27 IPsec ESP data encapsulation

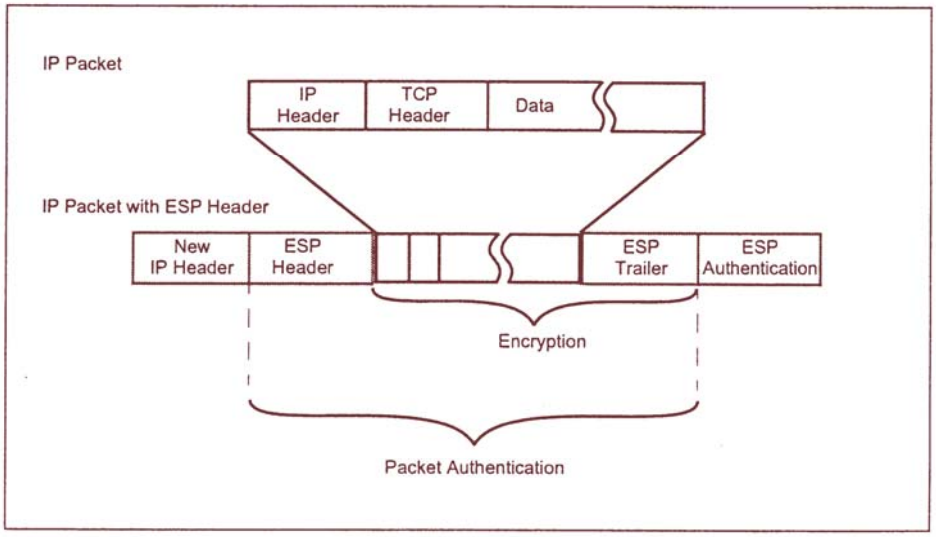

Figure 28 IPsec AH data encapsulation

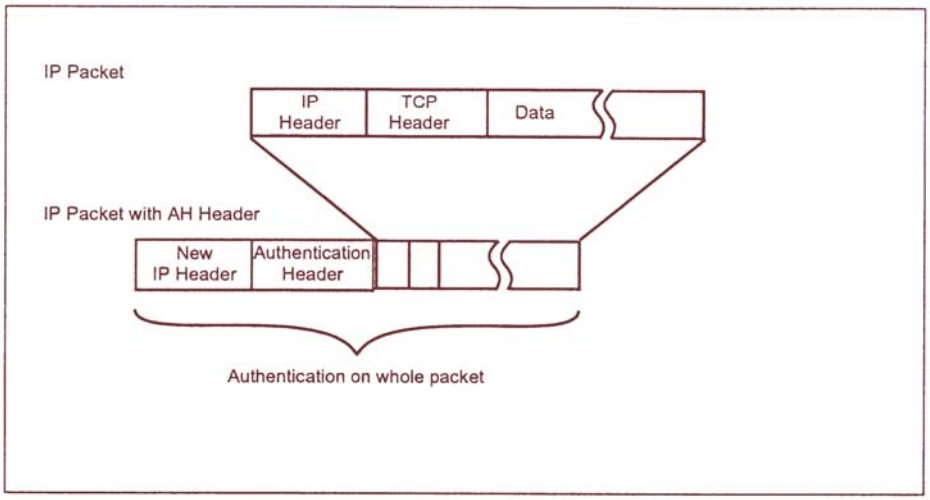

The headers added by ESP or AH contain the exterior IP address of the IPsec gateways, which keeps the inside network address of the packets invisible. This means that private IP addresses (as described in "Private networks" on page 103) can be used for the subnets that are tied together by the VPN tunnel.

IPSEC VPN RESILIENCE

At times it is necessary to have redundancy in a VPN solution. VPN Resilience provides redundancy by providing a backup to the VPN gateway at the remote end-point. This method of VPN redundancy and the methods implemented to detect a failure in the remote VPN gateway are presented in the following subsections.

Failures at a remote VPN gateway can be caused by the following:

- Network problems between the VPN gateways.
- Reboot(s) of the remote gateway (this terminates active VPN tunnels).
- Remote VPN gateway taken down for maintenance.
- Problems with the remote gateway, either with the hardware or software.

Remote VPN Gateway Failure Detection

The InterJak software uses a method for detecting failures of the remote VPN gateway called Dead Peer Detection (**DPD**). It works as follows:

- Each time TCP traffic travels through a VPN tunnel, it is expected that TCP traffic will be returned as a response to the traffic sent.
- Every 20 seconds after the detection of TCP traffic, returned TCP traffic is checked to determine if the VPN tunnel is operating correctly.
- If traffic has been sent and two consecutive 20 second periods have resulted in no returned TCP traffic, the InterJak attempts to renegotiate the tunnel.

This method works but has the disadvantage that if return traffic is unable to return through the tunnel due to routing problems behind the VPN gateway, the tunnel will be constantly renegotiated even though no VPN problems exist. A better method has been implemented with IKE keep-alives.

IKE keep-alives

This feature allows a VPN gateway to enquire whether the other VPN gateway is still alive, and has a specific Phase I SA active. This enquiry is done by sending an "Are U There" packet to the other end. The packet is encrypted using the Phase I SA. The other end responds with an "Are U There Ack" packet, encrypted with the SA. For example, if the other gateway has been rebooted (and thereby lost its active SAs) no response will be sent to the "Are U There" enquiry.

The use of IKE keep-alives is negotiated during the bringing up of the tunnel, and must be enabled for the InterJak to attempt to negotiate it.

The definition of IKE keep-alives leaves it up to the implementation when to send a keep-alive packet. Some devices send keep-alives every 10 to 20 seconds when no traffic is coming through the tunnel. The InterJak does it in a slightly different way, which in essence is an addition to the DPD.

- Each time TCP traffic travels through a VPN tunnel, it is expected that TCP traffic will be returned as a response to the traffic sent.
- Every 20 seconds after the detection of TCP traffic, returned TCP traffic is checked to determine if the VPN tunnel is operating correctly.
- If traffic has been sent but not received in one 20 second period, a keep-alive is sent.
- A total of 3 keep-alives are sent with a 5 second interval; if the remote VPN gateway does not respond to any of these keep-alives, the InterJak attempts to renegotiate the tunnel.

Thus if the problem lies behind the other VPN gateway the renegotiation problem with DPD is solved by a relatively cheap keep-alive packet, which is sent every 20 seconds, instead of an expensive renegotiation every 40 seconds.

Switching to the fail-over Remote VPN end-point

Once a failure in the remote VPN end-point is detected, the next step to take in the implementation of a Resilience VPN tunnel is to establish the VPN tunnel with the fail-over VPN gateway.

The Interjak can be configured to use a fail-over remote VPN gateway. Upon detection of a field VPN end-point, the InterJak tries to negotiate the VPN connections to the primary end-point. If this negotiation fails, the InterJak tries the backup end-point.

Once the VPN tunnel has been established to the fail-over VPN end-point, the InterJak tries to fall back to the primary end-point. The InterJak attempts to fall back to the primary VPN end-point at each renegotiation of the Phase II SA. The default configuration for this negotiation is every 8 hours.

Prevention of backup tunnel negotiation

In the default setting, the InterJak allows the backup tunnel to be negotiated from the network containing the duplicated VPN devices. This will cause a switch over to the backup tunnel, which in some setups might disrupt the connections through the VPN.

For this reason, there is an option to prevent negotiation of the backup tunnel from the outside. The option is **`Remote Initiation of Backup Tunnel`** and is placed in the Advanced Settings section of the VPN Site-to-Site configuration page of the web manager. If this option is unchecked, the InterJak will not allow the backup tunnel to be negotiated from the outside, but only if the primary tunnel is up.

Routing behind redundant gateways

Making VPN connections to redundant gateways present some routing problems at the site with the redundant gateways. Consider the following setup:

Figure 29

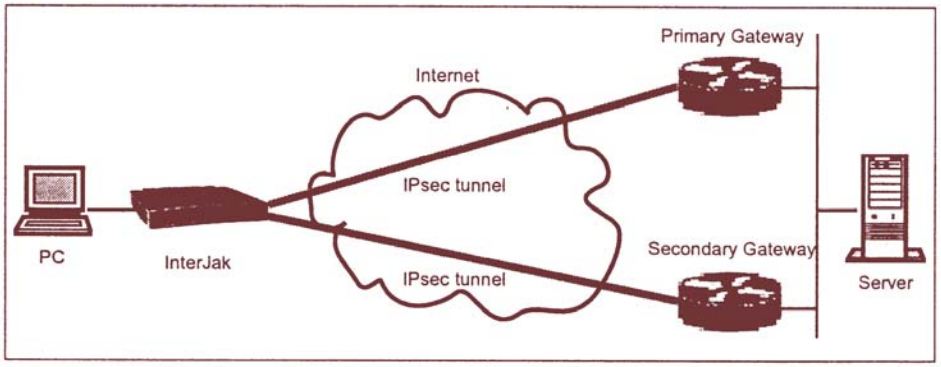

A VPN connection exists between the InterJak and the Primary Gateway with the Secondary Gateway as backup. The Server must have some idea of how to route traffic back to the PC. The route the PC uses must be dependent on which Gateway is currently active. This problem must be solved at the remote end, as Filanet does not currently offer a solution. VPN devices supporting e.g. Virtual Router Redundancy Protocol (VRRP) provide such functionality.

CONFIGURATION OPTIONS FOR VPN CONNECTIONS

As described in the previous sections, the VPN tunnel setup is a combination of many different options. This section describes the various options for a VPN tunnel.

Necessary options (without defaults)

- **Name**. A descriptive name for the VPN connection. Names of VPN connections should be unique.
- **Remote end-point**. The IP address or host name of the IPsec device, which defines the other end of the tunnel. It is not possible to define multiple tunnels to the same remote end-point.
- **Interface**. Specifies the interface on the InterJak on which the VPN tunnel starts. This interface must be of type WAN or DMZ.
- **Enabled**. Determines whether the VPN tunnel is currently enabled.

- **Local subnet**. Defines the local subnets that can go through the VPN tunnel. The subnets are defined using IP addresses and netmasks.
- **Remote subnet**. Defines the remote subnets that can go through the VPN tunnel. The subnets are defined using IP addresses and netmasks.
- **Shared secret**. The shared secret, which provides authentication of the two VPN end-points.

Advanced options (with defaults)

- **Backup Remote end-point**. Specifies the IP address or hostname of a backup remote VPN gateway. If a backup remote gateway is present and the system loses connection to the primary gateway, a VPN connection will be negotiated to the backup gateway instead. By default, this field is left blank.
- **Remote Initiation of Backup Tunnel.** Determines whether it is allowed for the backup tunnel to be initiated from the other end, if the primary tunnel is up. By default, this option is selected and the backup tunnel negotiation is allowed.
- **Negotiate keepalive**. Determines whether the connection should support IKE keepalives, which is a mechanism for faster detection of problems with the remote gateway. By default, this is selected.
- **Phase I negotiation mode**. Selects the negotiation mode; either Main Mode or Aggressive Mode. The default (and recommended) mode is Main.
- **Phase I options**. Due to the limitation that Aggressive mode only allows a single proposal for the encryption and authentication, it is possible to configure the next three options as a single combination.
  - **Encryption**. DES, 3DES
  - **Authentication**. MD5, SHA1
  - **Diffie-Hellman group**. Group 1 (768-bit), Group 2 (1024-bit), Group 5 (1536-bit).

The Default configuration for Phase I Options is 3DES encryption, MD5 authentication and Diffie-Hellman group 2. It is important that the chosen combination matches the remote VPN configuration.

- **Protocol**: It is possible to select protocol and encryption level.
  The allowed values for ESP are:
  - Triple DES encryption, MD5 signatures (default)
  - Triple DES encryption, SHA-1 signatures
  - Single DES encryption, MD5 signatures

- Single DES encryption, SHA-1 signatures
- No encryption, MD5 signatures
- No encryption, SHA-1 signatures

The allowed values for AH are:

- SHA-1 signatures
- MD5 signatures

- **Perfect Forward Secrecy (PFS)** Group Configuration.

In quick mode, it is possible to refresh the keying material at each renegotiation of keys. This is known as Perfect Forward Secrecy (PFS). PFS uses Diffie-Hellman and therefore a DH group can be configured with the following options

- No --- Does not use PFS
- Group 1 (768-bit)
- Group 2 (1024-bit) --- Default
- Group 5 (1536-bit)

The settings for PFS must match each other on the peer VPN gateways.

- **ISAKMP security association lifetime**: This is the lifetime of the Main mode encryption key. Valid values are: 10 mins - 8 hours. Default: 1 hour.
- **IPsec security association lifetime**: This is the lifetime of the Quick mode encryption key. Valid values are: 10 mins - 24 hours. Default: 8 hours.

In "Phase I - The VPN gateways establish SA for secure communication" on page 138 it is shown that Phase I negotiations exchange the identity of the VPN gateways. The identity is normally the IP address of the VPN gateway. However in some situations it is necessary to specify the identity explicitly.

The InterJak supports the following types of identities:

- IP addresses, e.g. 10.1.2.35
- Email addresses, e.g. user@company.com
- DNS names, e.g. interjak.company.com

You can configure the identities of the VPN gateways in the **Identities** section in the **Services:VPN Site-toSite/Client:Add Connection** page of the web manager. These identities are exchanged in the Phase I negotiations.

The identities configured on the peer VPN gateways must match in order for the VPN connection to be negotiated. The identity fields on the **Services:VPN Site-toSite/Client:Add Connection** web page are:

- **Local identity**

  Sets the local identity. Some VPN devices require the identity to be something other than the IP address, when a device is connecting from a non-static IP address.

- **Remote identity**

  Sets the remote identity. For example if the remote VPN gateway is behind a device using NAT, then the Remote Identity will not be equal to the IP address from which the IPsec negotiation seems to come from. In this case the local IP address of the remote VPN gateway can be inserted here.

NOTE *For a standard VPN connection the identity fields can be left blank.*

COMPATIBILITY

The InterJak is designed to be used with a broad spectrum of 3rd party VPN devices. Please see the Filanet web site (http://www.filanet.com) for the latest information about tested and compatible devices.

The following devices have, however, been tested and were found to be compatible with the InterJak:

- Cisco PIX
- Firewall-1 4.1
- Windows 2000
- FreeS/Wan for Linux
- Netscreen 5XP
- Nortel Contivity