


IBM i
Version 7.2

Security
Virtual Private Networking



IBM i
Version 7.2

Security
Virtual Private Networking



Note

Before using this information and the product it supports, read the information in “Notices” on page 83.

This document may contain references to Licensed Internal Code. Licensed Internal Code is Machine Code and is licensed to you under the terms of the IBM License Agreement for Machine Code.

© Copyright IBM Corporation 1998, 2013.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

<https://t.me/learningnets>

Contents

Virtual Private Networking 1

What's new for IBM i 7.2	1
PDF file for Virtual private network	2
VPN concepts	2
IP Security protocols	2
Authentication Header	3
Encapsulating Security Payload	5
AH and ESP combined	7
Enhanced Cryptographic Algorithms	7
Key management	9
IKE version 2	10
IKE_SA rekey	11
URL lookup of certificates	12
Layer 2 Tunnel Protocol	12
Network address translation for VPN	13
NAT compatible IPSec with UDP	14
IP Compression	15
VPN and IP filtering	16
VPN connections with no policy filters	16
Implicit IKE	16
Scenarios: VPN	17
Scenario: Basic branch office connection	17
Completing the planning worksheets	19
Configuring VPN on System A	20
Configuring VPN on System C	21
Starting VPN	21
Testing a connection	21
Scenario: Basic business to business connection	22
Completing the planning worksheets	24
Configuring VPN on System A	24
Configuring VPN on System C	25
Activating packet rules	25
Starting a connection	26
Testing a connection	26
Scenario: Protecting an L2TP voluntary tunnel with IPSec	26
Configuring VPN on System A	28
Configuring a PPP connection profile and virtual line on System A	30
Applying the 12tpocorp dynamic-key group to the toCorp PPP profile	31
Configuring VPN on System B	31
Configuring a PPP connection profile and virtual line on System B	32
Activating packet rules	32
Scenario: Firewall friendly VPN	33
Completing the planning worksheets	34
Configuring VPN on Gateway B	36
Configuring VPN on System E	37
Starting Connection	38
Testing the connection	38
Scenario: VPN connection to remote users	38
Completing planning worksheets for VPN connection from the branch office to remote sales people	39

Configuring L2TP terminator profile for System A	40
Starting receiver connection profile	40
Configuring a VPN connection on System A for remote clients	41
Activating filter rules	42
Configuring VPN on Windows client	42
Testing VPN connection between endpoints	43
Scenario: Using network address translation for VPN	43
Planning for VPN	45
VPN setup requirements	45
Determining what type of VPN to create	46
Completing VPN planning work sheets	46
Planning work sheet for dynamic connections	46
Planning work sheet for manual connections	47
Configuring VPN	49
Configuring VPN connections with the New Connection wizard	49
Configuring VPN security policies	50
Configuring an Internet Key Exchange policy	50
Configuring a data policy	51
Configuring a secure VPN connection	51
Part 1: Configure a dynamic-key group	52
Part 2: Configure a dynamic-key connection	52
Configuring a manual connection	52
Configuring a dynamic connection	53
Configuring VPN packet rules	53
Configuring the pre-IPSec filter rule	55
Configuring a policy filter rule	56
Defining an interface for the VPN filter rules	57
Activating VPN packet rules	58
Configure traffic flow confidentiality	58
Configuring extended sequence number	59
Configuring populate from packet	59
Configuring NPF VPN for VIPA	59
Configuration Differences for IKEv2	60
Starting a VPN connection	60
Managing VPN	61
Setting default attributes for your connections	61
Resetting connections in error state	61
Viewing error information	62
Viewing attributes of active connections	62
Viewing VPN server job logs	62
Viewing attributes of Security Associations	63
Stopping a VPN connection	63
Deleting VPN configuration objects	63
Troubleshooting VPN	63
Getting started with VPN troubleshooting	63
Other things to check	64
Common VPN configuration errors and how to fix them	65
VPN error message: TCP5B28	65
VPN error message: Item not found	65
VPN error message: PARAMETER PINBUF IS NOT VALID	66

VPN error message: Item not found, Remote key server...	67	Troubleshooting VPN with the QVPN journal	72
VPN error message: Unable to update the object	67	Enabling the QVPN journal	72
VPN error message: Unable to encrypt key...	67	Using the QVPN journal	73
VPN error message: CPF9821	68	QVPN journal fields	73
VPN error: All keys are blank	68	Troubleshooting VPN with the VPN job logs	74
VPN error: Connection has enabled status after you stop it	68	Common VPN Connection Manager error messages	75
VPN error: Active filter rules fail to deactivate	69	Troubleshooting VPN with the communications trace.	80
VPN error: The key connection group for a connection changes	69	Related information for VPN	81
Troubleshooting VPN with the QIPFILTER journal	69	Notices 83	
Enabling the QIPFILTER journal	70	Programming interface information	84
Using the QIPFILTER journal	70	Trademarks	85
QIPFILTER journal fields	71	Terms and conditions	85

Virtual Private Networking

A virtual private network (VPN) allows your company to securely extend its private intranet over the existing framework of a public network, such as the Internet. With VPN, your company can control network traffic while providing important security features such as authentication and data privacy.

You can configure VPN with IBM® Navigator for i, the graphical user interface (GUI) for IBM i. It allows you to create a secure end-to-end path between any combination of host and gateway. VPN uses authentication methods, encryption algorithms, and other precautions to ensure that data sent between the two endpoints of its connection remains secure.

VPN runs on the network layer of the TCP/IP layered communications stack model. Specifically, VPN uses the IP Security Architecture (IPSec) open framework. IPSec provides base security functions for the Internet, as well as furnishes flexible building blocks from which you can create robust, secure virtual private networks.

VPN also supports Layer 2 Tunnel Protocol (L2TP) VPN solutions. L2TP connections, which are also called virtual lines, provide cost-effective access for remote users by allowing a corporate network server to manage the IP addresses assigned to its remote users. Further, L2TP connections provide secure access to your system or network when you protect them with IPSec.

It is important that you understand the effect a VPN will have on your entire network. Correct planning and implementation are essential to your success. Review these topics to ensure that you know how VPNs work and how you might use them:

What's new for IBM i 7.2

Read about new or changed information for the Virtual Private Networking topic collection.

IKE Version 2

Enhancements were added to the IKE version 2 support.

- NAT compatible IpSec negotiation support in IKEv2.
- Rekey of IKE_SA.
- Authentication with Elliptical Curve Digital Signature Algorithm (ECDSA).
- URL lookup of certificates.
- Enhanced Cryptographic Algorithms for Key Exchange Policy and Data Policy security association attributes.

Filter rules and VPN commands

There are new commands to load and unload filter rules and to manage VPN connections.

- LODIPFTR - load and unload filter rules.
- STRVPNCNN - start VPN connection.
- ENDVPNCNN - end VPN connection.
- CPYVPNCFGF - XML export and import VPN configurations.

How to see what's new or changed

To help you see where technical changes were made, the information center uses:

- The  image to mark where new or changed information begins.
- The  image to mark where new or changed information ends.

In PDF files, you might see revision bars (|) in the left margin of new and changed information.

To find other information about what's new or changed this release, see the Memo to users.

PDF file for Virtual private network

You can view and print a PDF file of this information.

To view or download the PDF version of this document, select Virtual private network (VPN)  .

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF link in your browser.
2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you would like to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

You need Adobe Reader installed on your system to view or print these PDFs. You can download a copy from the Adobe Web site (www.adobe.com/products/acrobat/readstep.html) .

VPN concepts

It is important that you have at least a basic knowledge of standard VPN technologies before you implement a VPN connection.

Virtual private networking (VPN) uses several important TCP/IP protocols to protect data traffic. To better understand how any VPN connection works, familiarize yourself with these protocols and concepts and how VPN uses them:

IP Security protocols

IP Security (IPSec) provides a stable, long lasting base for providing network layer security.

IPSec supports all of the cryptographic algorithms in use today, and can also accommodate newer, more powerful algorithms as they become available. IPSec protocols address these major security issues:

Data origin authentication

Verifies that each datagram was originated by the claimed sender.

Data integrity

Verifies that the contents of a datagram were not changed in transit, either deliberately or due to random errors.

Data confidentiality

Conceals the content of a message, typically by using encryption.

Replay protection

Ensures that an attacker cannot intercept a datagram and play it back at some later time.

Automated management of cryptographic keys and security associations

Ensures that your VPN policy can be used throughout the extended network with little or no manual configuration.

VPN uses two IPSec protocols to protect data as it flows through the VPN: Authentication Header (AH) and Encapsulating Security Payload (ESP). The other part of IPSec enablement is the Internet Key Exchange (IKE) protocol, or key management. While IPSec encrypts your data, IKE supports automated negotiation of security associations (SAs), and automated generation and refreshing of cryptographic keys.

Note: Some VPN configurations could have a security vulnerability depending on how IPSec is configured. The vulnerability affects configurations where IPSec is configured to employ Encapsulating Security Payload (ESP) in tunnel mode with confidentiality (encryption), but without integrity protection (authentication) or Authentication Header (AH). The default configuration when ESP is selected always includes an authentication algorithm that provides integrity protection. Therefore, unless the authentication algorithm in the ESP transform is removed, VPN configurations will be protected from this vulnerability. The IBM Universal Connection VPN configuration is not affected by this vulnerability.

To check if your system is affected by this security vulnerability follow these steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **IP Security Policies**.
- | 2. Right-click **Data Policies** and select **Open**.
- | 3. Right-click on the data policy you want to check and select **Properties**.
- | 4. Click on the **Proposals** tab.
- | 5. Select any of the data protection proposals that are using the ESP protocol and click **Edit**.
- | 6. Click on the **Transforms** tab.
- | 7. Select any transforms from the list that use the ESP protocol and click **Edit**.
- | 8. Verify that the Authentication algorithm has any other value then **None**.

The Internet Engineering Task Force (IETF) formally defines IPSec in Request for Comment (RFC) 4301, *Security Architecture for the Internet Protocol*. You can view this RFC on the Internet at the following Web site: <http://www.rfc-editor.org>.

The principal IPSec protocols are listed below:

Related concepts:

"Key management" on page 9

A dynamic VPN provides additional security for your communications by using the Internet Key Exchange (IKE) protocol for key management. IKE allows the VPN servers on each end of the connection to negotiate new keys at specified intervals.

Related information:

 <http://www.rfc-editor.org>

Authentication Header

The Authentication Header (AH) protocol provides data origin authentication, data integrity, and replay protection. However, AH does not provide data confidentiality, which means that all of your data is sent in the clear.

AH ensures data integrity with the checksum that a message authentication code, like MD5, generates. To ensure data origin authentication, AH includes a secret shared key in the algorithm that it uses for authentication. To ensure replay protection, AH uses a sequence number field within the AH header. It is

worth noting here, that these three distinct functions are often lumped together and referred to as authentication. In the simplest terms, AH ensures that your data has not been tampered with en route to its final destination.

Although AH authenticates as much of the IP datagram as possible, the values of certain fields in the IP header cannot be predicted by the receiver. AH does not protect these fields, known as mutable fields. However, AH always protects the payload of the IP packet.

The Internet Engineering Task Force (IETF) formally defines AH in Request for Comment (RFC) 4302, *IP Authentication Header*. You can view this RFC on the Internet at the following Web site:
<http://www.rfc-editor.org>.

Ways of using AH

You can apply AH in two ways: transport mode or tunnel mode. In transport mode, the IP header of the datagram is the outermost IP header, followed by the AH header and then the payload of the datagram. AH authenticates the entire datagram, except the mutable fields. However, the information contained in the datagram is transported in the clear and is, therefore, subject to eavesdropping. Transport mode requires less processing overhead than tunnel mode, but does not provide as much security.

Tunnel mode creates a new IP header and uses it as the outermost IP header of the datagram. The AH header follows the new IP header. The original datagram (both the IP header and the original payload) comes last. AH authenticates the entire datagram, which means that the responding system can detect whether the datagram changed while in transit.

When either end of a security association is a gateway, use tunnel mode. In tunnel mode the source and destination addresses in the outermost IP header do not need to be the same as those in the original IP header. For example, two security gateways might operate an AH tunnel to authenticate all traffic between the networks they connect together. In fact, this is a very typical configuration.

The main advantage to using tunnel mode, is that tunnel mode totally protects the encapsulated IP datagram. In addition, tunnel mode makes it possible to use private addresses.

Why AH?

In many cases, your data only requires authentication. While the Encapsulating Security Payload (ESP) protocol can perform authentication, AH does not affect your system performance as does ESP. Another advantage of using AH, is that AH authenticates the entire datagram. ESP, however, does not authenticate the leading IP header or any other information that comes before the ESP header.

In addition, ESP requires strong cryptographic algorithms in order to be put into effect. Strong cryptography is restricted in some regions, while AH is not regulated and can be used freely around the world.

Using ESN with AH

If you use the AH protocol then you might want to enable Extended Sequence Number (ESN). ESN allows you to transmit large volumes of data at a high speed without re-keying. The VPN connection uses a 64-bit sequence numbers instead of 32-bit numbers over IPSec. Using 64-bit sequence numbers allows more time before re-keying, which prevents sequence number exhaustion and minimizes the use of system resources.

What algorithms does AH use to protect my information?

AH uses algorithms known as **hashed message authentication codes (HMAC)**. Specifically, VPN uses HMAC-MD5, HMAC-SHA, HMAC-SHA-256, HMAC-SHA384, HMAC-SHA512, or AES-XCBC-MAC. Each

of the algorithms take variable-length input data and a secret key to produce fixed-length output data (called a hash or MAC value). If the hashes of two messages match, then it is likely that the messages are the same.

The Internet Engineering Task Force (IETF) formally defines the algorithms in the following Request for Comments (RFC):

- HMAC-MD5 in RFC 2085, *HMAC-MD5 IP Authentication with Replay Prevention*
- HMAC-SHA in RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
- HMAC-SHA_256, HMAC-SHA-384, and HMAC-SHA-512 in RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*
- AES-XCBC-MAC in RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*

You can view these RFCs on the Internet at the following Web site: <http://www.rfc-editor.org>.

Related concepts:

“Encapsulating Security Payload”

The Encapsulating Security Payload (ESP) protocol provides data confidentiality, and also optionally provides data origin authentication, data integrity checking, and replay protection.

“Enhanced Cryptographic Algorithms” on page 7

Cryptographic algorithms supported in the VPN selection for Key Exchange Policy and Data policy security association attributes.

Related information:

 <http://www.rfc-editor.org>

Encapsulating Security Payload

The Encapsulating Security Payload (ESP) protocol provides data confidentiality, and also optionally provides data origin authentication, data integrity checking, and replay protection.

The difference between ESP and the Authentication Header (AH) protocol is that ESP provides encryption, while both protocols provide authentication, integrity checking, and replay protection. With ESP, both communicating systems use a shared key for encrypting and decrypting the data they exchange.

If you decide to use both encryption and authentication, then the responding system first authenticates the packet and then, if the first step succeeds, the system proceeds with decryption. This type of configuration reduces processing overhead, as well as reduces your vulnerability to denial-of-service attacks.

Two ways of using ESP

You can apply ESP in two ways: transport mode or tunnel mode. In transport mode, the ESP header follows the IP header of the original IP datagram. If the datagram already has an IPSec header, then the ESP header goes before it. The ESP trailer and the optional authentication data follow the payload.

Transport mode does not authenticate or encrypt the IP header, which might expose your addressing information to potential attackers while the datagram is in transit. Transport mode requires less processing overhead than tunnel mode, but does not provide as much security. In most cases, hosts use ESP in transport mode.

Tunnel mode creates a new IP header and uses it as the outermost IP header of the datagram, followed by the ESP header and then the original datagram (both the IP header and the original payload). The ESP trailer and the optional authentication data are appended to the payload. When you use both encryption and authentication, ESP completely protects the original datagram because it is now the payload data for

the new ESP packet. ESP, however, does not protect the new IP header. Gateways must use ESP in tunnel mode.

What algorithms does ESP use to protect my information?

ESP uses a symmetric key that both communicating parties use to encrypt and decrypt the data they exchange. The sender and the receiver must agree on the key before secure communication takes place between them. VPN uses Data Encryption Standard (DES), triple-DES (3DES), Advanced Encryption Standard (AES), or AES-CBC and AES-CTR for encryption.

If you choose the AES algorithm for encryption then you might want to enable Extended Sequence Number (ESN). ESN allows you to transmit large volumes of data at a high speed. The VPN connection uses a 64-bit sequence numbers instead of 32-bit numbers over IPSec. Using 64-bit sequence numbers allows more time before re-keying, which prevents sequence number exhaustion and minimizes the use of system resources.

The Internet Engineering Task Force (IETF) formally defines the algorithms in the following Request for Comments (RFC):

- DES in Request for Comment (RFC) 1829, *The ESP DES-CBC Transform*
- 3DES in RFC 1851, *The ESP Triple DES Transform*
- AES-CBC in RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
- AES-CTR in RFC 3686, *Using Advanced Encryption Standard (AES) Counter Mode with IPsec Encapsulating Security Payload (ESP)*

You can view these and other RFCs on the Internet at the following Web address: <http://www.rfc-editor.org>.

ESP uses HMAC-MD5, HMAC-SHA, HMAC-SHA-256, HMAC-SHA-384, HMAC-SHA-512, and AES-XCBC-MAC algorithms to provide authentication functions. Each of the algorithms take variable-length input data and a secret key to produce fixed-length output data (called a hash or MAC value). If the hashes of two messages match, then it is likely that the messages are the same.

The Internet Engineering Task Force (IETF) formally defines the algorithms in the following Request for Comments (RFC):

- HMAC-MD5 in RFC 2085, *HMAC-MD5 IP Authentication with Replay Prevention*
- HMAC-SHA in RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
- HMAC-SHA_256, HMAC-SHA-384, and HMAC-SHA-512 in RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*
- AES-XCBC-MAC in RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*

You can view these RFCs on the Internet at the following Web address: <http://www.rfc-editor.org>.

ESP uses AES-CCM and AES-GCM to provide encryption and authentication. An authentication algorithm cannot be selected if one of these "combined" algorithms is chosen.

- AES-CCM in RFC 4309, *Using Advanced Encryption Standard (AES) CCM mode with IPsec Encapsulating Security Payload (ESP)*
- AES-GCM in RFC 4106, *The Use of Galios/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)*

ESP uses AES-GMAC (Galios Message Authentication Code) to provide authentication, but not encryption.

- AES-GMAC in RFC 4543, *The Use of Galios Message Authentication Code (GMAC) in IPsec ESP and AH.*

Related concepts:

“Authentication Header” on page 3

The Authentication Header (AH) protocol provides data origin authentication, data integrity, and replay protection. However, AH does not provide data confidentiality, which means that all of your data is sent in the clear.

“Enhanced Cryptographic Algorithms”

- | Cryptographic algorithms supported in the VPN selection for Key Exchange Policy and Data policy security association attributes.

Related information:

 <http://www.rfc-editor.org>

AH and ESP combined

VPN allows you to combine AH and ESP for host-to-host connections in transport mode.

Combining these protocols protects the entire IP datagram. Although combining the two protocols offers more security, the processing overhead involved might outweigh the benefit.

Enhanced Cryptographic Algorithms

- | Cryptographic algorithms supported in the VPN selection for Key Exchange Policy and Data policy security association attributes.

Key Exchange Policy:

- Encryption
 - 3DES-CBC
 - | – AES-CBC (128, 192, and 256 bit)
 - | – AES-CTR (128, 192, and 256 bit)
- Hash/PRF
 - SHA
 - HMAC-SHA-256
 - | – HMAC-SHA-384
 - | – HMAC-SHA-512
 - AES-XCBC-MAC (HASH 96 bits; PRF 128 bits)
- Diffie-Hellman
 - Group 1
 - Group 2
 - Group 14
 - | – Group 19 (256 ECP)
 - | – Group 20 (384 ECP)
 - Group 24

Data Policy:

- Authentication
 - SHA
 - HMAC-SHA-256
 - | – HMAC-SHA-384
 - | – HMAC-SHA-512
 - AES-XCBC-MAC
- Diffie-Hellman for PFS
 - Group 1
 - Group 2

- Group 14
- | – Group 19 (256 bit ECP)
- | – Group 20 (384 bit ECP)
- Group 24
- | • Encryption
- | – 3DES-CBC
- | – AES-CBC (128, 192, and 256 bit)
- | – AES-CTR (128, 192, and 256 bit)
- | – AES-CCM (128, 192, and 256 bit)
- | – AES-GCM (128, 192, and 256 bit)
- | – AES-GMAC (128, 192, and 256 bit)

In addition to the enhanced cryptographic algorithms being supported, the following algorithms are de-emphasized. They are still supported, but the direction is to use them less.

- Hash
 - MD5
- Encryption
 - DES
 - RC4
 - RC5

The Internet Engineering Task Force (IETF) formally defines the algorithms in the following Request for Comments (RFC):

- AES-CBC in RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
- AES-XCBC-MAC in RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
- | • HMAC-SHA_256, HMAC-SHA-384, and HMAC-SHA-512 in RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*
- HMAC-MD5 in RFC 2085, *HMAC-MD5 IP Authentication with Replay Prevention*
- DES in Request for Comment (RFC) 1829, *The ESP DES-CBC Transform*
- | • DH groups 19 and 20 in RFC 4754, *IKE and IKEV2 Authentication Using the Elliptical Curve Digital Signature Algorithm (ECDSA)*
- | • AES-CTR in RFC 3686, *Using Advanced Encryption (AES) Counter Mode with IPsec Encapsulating Security Payload (ESP)*
- | • AES-CCM in RFC 4309, *Using Advanced Encryption Standard (AES) CCM mode with IPsec Encapsulating Security Payload (ESP)*
- | • AES-GCM in RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)*
- | • AES-GMAC in RFC 4543, *The Use of Galois Message Authentication Mode (GMAC) in IPsec ESP and AH*

You can view these RFCs on the Internet at the following Web address: <http://www.rfc-editor.org>.

Related concepts:

“Authentication Header” on page 3

The Authentication Header (AH) protocol provides data origin authentication, data integrity, and replay protection. However, AH does not provide data confidentiality, which means that all of your data is sent in the clear.

“Encapsulating Security Payload” on page 5

The Encapsulating Security Payload (ESP) protocol provides data confidentiality, and also optionally provides data origin authentication, data integrity checking, and replay protection.

Key management

A dynamic VPN provides additional security for your communications by using the Internet Key Exchange (IKE) protocol for key management. IKE allows the VPN servers on each end of the connection to negotiate new keys at specified intervals.

With each successful negotiation, the VPN servers regenerate the keys that protect a connection, thus making it more difficult for an attacker to capture information from the connection. Additionally, if you use perfect forward secrecy, attackers cannot derive future keys based on past keying information.

The VPN key manager is IBM's implementation of the Internet Key Exchange (IKE) protocol. The key manager supports the automatic negotiation of security associations (SAs), as well as the automatic generation and refresh of cryptographic keys.

A **security association (SA)** contains information that is necessary to use the IPSec protocols. For example, an SA identifies algorithm types, key lengths and lifetimes, participating parties, and encapsulation modes.

Cryptographic keys, as the name implies, lock, or protect, your information until it safely reaches its final destination.

Note: Securely generating your keys is the most important factor in establishing a secure and private connection. If your keys are compromised, then your authentication and encryption efforts, no matter how strong, become worthless.

Phases of key management

The VPN key manager uses two distinct phases in its implementation.

Phase 1

Phase 1 establishes a master secret from which subsequent cryptographic keys are derived in order to protect user data traffic. This is true even if no security protection yet exists between the two endpoints. VPN uses either RSA signature mode, ECDSA signature mode, or preshared keys to authenticate phase 1 negotiations, as well as to establish the keys that protect the IKE messages that flow during the subsequent phase 2 negotiations.

A *preshared key* is a nontrivial string up to 128 characters long. Both ends of a connection must agree on the preshared key. The advantage of using preshared keys is their simplicity, the disadvantage is that a shared secret must be distributed out-of-band, for example over the telephone or through registered mail, before IKE negotiations. Treat your preshared key like a password.

RSA Signature authentication provides more security than preshared keys because this mode uses digital certificates to provide authentication. You must configure your digital certificates using Digital Certificate Manager. VPN has no limit on the RSA key lengths it supports. The certificates that you use must come from certificate authorities that both key servers trust.

ECDSA Signatures are smaller than RSA Signatures of similar cryptographic strength, resulting in improved communication efficiency. ECDSA Signatures supports three key lengths, ECDSA-256, ECDSA-384, and ECDSA-521. You must configure your digital certificates using Digital Certificate Manager. The certificates that you use must come from certificate authorities that both key servers trust. ECDSA Signatures are not supported in IKEv1.

Phase 2

Phase 2, however, negotiates the security associations and keys that protect the actual application data exchanges. Remember, up to this point, no application data has actually been sent. Phase 1 protects the phase 2 IKE messages.

Once phase 2 negotiations are complete, your VPN establishes a secure, dynamic connection over the network and between the endpoints that you defined for your connection. All data that flows across the VPN is delivered with the degree of security and efficiency that was agreed on by the key servers during the phase 1 and phase 2 negotiation processes.

In general, phase 1 negotiations are negotiated once a day, while phase 2 negotiations are refreshed every 60 minutes or as often as every five minutes. Higher refresh rates increase your data security, but decrease system performance. Use short key lifetimes to protect your most sensitive data.

When you create a dynamic VPN by using IBM Navigator for i, you must define an IKE policy to enable phase 1 negotiations and a data policy to govern phase 2 negotiations. Optionally, you can use the New Connection wizard. The wizard automatically creates each of the configuration objects VPN requires to work properly, including an IKE policy, data policy.

Suggested reading

If you want to read more about the Internet Key Exchange (IKE) protocol and key management, review these Internet Engineering Task Force (IETF) Request for Comments (RFC):

- | **IKEv1**
- | • RFC 2407, *The Internet IP Security Domain of Interpretation for ISAKMP*
- | • RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
- | • RFC 2409, *The Internet Key Exchange (IKE)*
- | These RFCs are currently supported for IKEv1.

You can view these RFCs on the Internet at the following Web site: <http://www.rfc-editor.org>.

Related concepts:

“Scenario: Firewall friendly VPN” on page 33

In this scenario, a large insurance company wants to establish a VPN between a gateway in Chicago and a host in Minneapolis when both networks are behind a firewall.

“IP Security protocols” on page 2

IP Security (IPSec) provides a stable, long lasting base for providing network layer security.

“IKE version 2”

IKE version 2 is an enhancement to the Internet key exchange protocol.

Related tasks:

“Configuring an Internet Key Exchange policy” on page 50

The Internet Key Exchange (IKE) policy defines what level of authentication and encryption protection IKE uses during phase 1 negotiations.

“Configuring a data policy” on page 51

A data policy defines what level of authentication or encryption protects data as it flows through the VPN.

Related information:

 <http://www.rfc-editor.org>

IKE version 2

IKE version 2 is an enhancement to the Internet key exchange protocol.

IKE version 2 (IKEv2) was developed by the IETF with RFC4306. IKEv2 enhances the function of negotiating the dynamic key exchange and authentication of the negotiating systems for VPN.

IKEv2 also simplifies the key exchange flows and introduces measures to fix ambiguities and vulnerabilities inherent in IKEv1.

- IKEv2 provides a simpler message flow for key exchange negotiations.
- IKEv2 provides options to rekey the IKE_SA without reauthentication.
- With IKEv2, the key life times for the IKE_SA and CHILD_SA are managed independent of the peer system.
- IKEv2 is the basis for future enhancements to the key exchange protocol.

Both IKEv1 and IKEv2 protocols operate in two phases. The differences between the two protocols include:

- The first phase in IKEv2 is IKE_SA, consisting of the message pair IKE_SA_INIT. The attributes of the IKE_SA phase are defined in the Key Exchange Policy.
- The second phase in IKEv2 is CHILD_SA. The first CHILD_SA is the IKE_AUTH message pair. More CHILD_SA message pairs can be sent for rekey and informational messages. The CHILD_SA attributes are defined in the Data Policy.

IKEv2 provides a simpler and more efficient exchange.

- IKEv1 phase 1 has two possible exchanges: main mode and aggressive mode. With main mode, the phase 1 and phase 2 negotiations are in two separate phases. Phase 1 main mode uses six messages to complete; phase 2 in quick mode uses three messages.
- IKEv2 combines these modes into a four message sequence. The IKE_SA is negotiated and authenticated and then the CHILD_SA is negotiated and keys are generated in four messages. Subsequent rekeying of the CHILD_SA is accomplished in two messages.

- Despite these changes, the basic outcome of the two versions is the same. IKEv1 and IKEv2 both negotiate a security association to protect data between two endpoints.

Suggested reading

For more information about the Internet Key Exchange (IKE) protocol and key management, review these Internet Engineering Task Force (IETF) Request for Comments (RFC):

- IKEv2**
 - RFC 4306, *The Internet Key Exchange (IKEv2) Protocol*
 - RFC 5996, *The Internet key Exchange Protocol Version 2* (supported in IBM i 7.2 only).
- These RFCs are currently supported for IKEv2.

You can view these RFCs on the Internet at the following Web site: <http://www.rfc-editor.org>.

Related concepts:

“Key management” on page 9

A dynamic VPN provides additional security for your communications by using the Internet Key Exchange (IKE) protocol for key management. IKE allows the VPN servers on each end of the connection to negotiate new keys at specified intervals.

“Configuration Differences for IKEv2” on page 60

Compares configuration for IKEv1 and IKEv2.

“Troubleshooting VPN with the VPN job logs” on page 74

When you encounter problems with your VPN connections, it is always advisable to analyze the job logs. In fact, there are several job logs that contain error messages and other information related to a VPN environment.

IKE_SA rekey

- IKE_SA rekey allows the IKE_SA to be rekeyed without re-authentication.

| The IKE_SA rekey takes place as an independent exchange when one side determines the IKE_SA has expired. A new IKE_SA is rekeyed in time for the CHILD_SA rekey attempt.

| To enable IKE_SA rekeying, follow these steps:

- | • Expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
- | • Click **Properties**.
- | • Check the **Enable IKE_SA rekey for IKEv2** box to turn on IKE_SA rekeying.

| **Note:** The **Enable IKE_SA rekey for IKEv2** box is not checked by default. If it is determined that an IKE_SA should always be authenticated for security purposes, the **Enable IKE_SA rekey for IKEv2** box should remain unchecked.

| URL lookup of certificates

| The URL lookup of certificates allows IKE negotiation partners to send a URL link to a certificate being used to authenticate the exchange.

| The purpose of using URL lookup of certificates is to save bandwidth in the IKE negotiation payload.

| IKEv2 supports only peers sending URL lookup of certificates. To enable this support and require peers to always send the complete certificate during IKE_SA authentication, follow these steps:

- | • Expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
- | • Click **Properties**.
- | • Check the **Allow HTTP lookup of certificates sent by remote systems** box to turn on URL lookup of certificates.

Layer 2 Tunnel Protocol

Layer 2 Tunneling Protocol (L2TP) connections, which are also called virtual lines, provide cost-effective access for remote users by allowing a corporate network systems to manage the IP addresses assigned to its remote users. Further, L2TP connections provide secure access to your system or network when you use them in conjunction with IP Security (IPSec).

L2TP supports two tunnel modes: the voluntary tunnel and the compulsory tunnel. The major difference between these two tunnel modes is the endpoint. On the voluntary tunnel, the tunnel ends at the remote client whereas the compulsory tunnel ends at the Internet Service Provider (ISP).

With an L2TP **compulsory tunnel**, a remote host initiates a connection to its ISP. The ISP then establishes an L2TP connection between the remote user and the corporate network. Although the ISP establishes the connection, you decide how to protect the traffic by using VPN. With a compulsory tunnel, the ISP must support L2TP.

With an L2TP **voluntary tunnel**, the connection is created by the remote user, typically by using an L2TP tunneling client. As a result, the remote user sends L2TP packets to its ISP which forwards them on to the corporate network. With a voluntary tunnel, the ISP does not need to support L2TP. The scenario, Protect an L2TP voluntary tunnel with IPSec provides you with an example of how to configure a branch office system to connect to its corporate network through a gateway system with an L2TP tunnel protected by VPN.

You can view a visual presentation about the concept of L2TP voluntary tunnels protected by IPSec. This requires the Flash plug-in. Alternatively, you can use the HTML version of this presentation.

L2TP is actually a variation of an IP encapsulation protocol. The L2TP tunnel is created by encapsulating an L2TP frame inside a User Datagram Protocol (UDP) packet, which in turn is encapsulated inside an IP packet. The source and destination addresses of this IP packet define the endpoints of the connection. Because the outer encapsulating protocol is IP, you can apply IPSec protocols to the composite IP packet.

This protects the data that flows within the L2TP tunnel. You can then apply Authentication Header (AH), Encapsulated Security Payload (ESP), and the Internet Key Exchange (IKE) protocol in a straightforward way.

Related concepts:

“Scenario: Protecting an L2TP voluntary tunnel with IPSec” on page 26

In this scenario, you learn how to setup a connection between a branch office host and a corporate office that uses L2TP protected by IPSec. The branch office has a dynamically assigned IP address, while the corporate office has a static, globally routable IP address.

Network address translation for VPN

VPN provides a means for performing network address translation, called VPN NAT. VPN NAT differs from traditional NAT in that it translates addresses before applying the IKE and IPSec protocols. Refer to this topic to learn more.

Network address translation (NAT) takes your private IP addresses and translates them into public IP addresses. This helps conserve valuable public addresses while at the same time allows hosts in your network to access services and remote hosts across the Internet (or other public network).

In addition, if you use private IP addresses, they can collide with similar, incoming IP addresses. For example, you might want to communicate with another network but both networks use 10.*.* addresses, causing the addresses to collide and all packets to be dropped. Applying NAT to your outbound addresses might appear to be the answer to this problem. However, if the data traffic is protected by a VPN, conventional NAT will not work because it changes the IP addresses in the security associations (SAs) that VPN requires to function. To avoid this problem, VPN provides its own version of network address translation called VPN NAT. VPN NAT performs address translation before the SA validation by assigning an address to the connection when the connection starts. The address remains associated with the connection until you delete the connection.

Note: FTP does not support VPN NAT at this time.

How should I use VPN NAT?

There are two different types of VPN NAT that you need to consider before you begin. They are:

VPN NAT for preventing IP address conflicts

This type of VPN NAT allows you to avoid possible IP address conflicts when you configure a VPN connection between networks or systems with similar addressing schemes. A typical scenario is one where both companies want to create VPN connections by using one of the designated private IP address ranges. For example, 10.*.*. How you configure this type of VPN NAT depends on whether your system is the initiator or the responder for the VPN connection. When your system is the connection initiator, you can translate your local addresses into ones that are compatible with your VPN connection partner's address. When your system is the connection responder, you can translate your VPN partner's remote addresses into ones that are compatible with your local addressing scheme. Configure this type of address translation only for your dynamic connections.

VPN NAT for hiding local addresses

This type of VPN NAT is used primarily to hide the real IP address of your local system by translating its address to another address that you make publicly available. When you configure VPN NAT, you can specify that each publicly known IP address be translated to one of a pool of hidden addresses. This also allows you to balance the traffic load for an individual address across multiple addresses. VPN NAT for local addresses requires that your system act as the responder for its connections.

Use VPN NAT for hiding local addresses if you answer yes to these questions:

1. Do you have one or more systems that you want people to access by using a VPN?
2. Do you need to be flexible about the actual IP addresses of your systems?

3. Do you have one or more globally routable IP addresses?

The scenario, Use network address translation for VPN provides you with an example of how to configure VPN NAT to hide local addresses on your IBM i model.

For step-by-step instructions on how to set up VPN NAT on your system, use the online help available from the VPN interface in IBM Navigator for i.

Related concepts:

“Scenario: Using network address translation for VPN” on page 43

In this scenario, your company wants to exchange sensitive data with one of its business partners by using VPN. To further protect the privacy of your company's network structure, your company will also use VPN NAT to hide the private IP address of the system it uses to host the applications to which your business partner has access.

“Planning work sheet for manual connections” on page 47

Complete this work sheet before you configure a manual connection.

NAT compatible IPSec with UDP

UDP encapsulation allows IPSec traffic to pass through a conventional NAT device. Review this topic for more information about what it is and why you should use it for your VPN connections.

The problem: Conventional NAT breaks VPN

Network address translation (NAT) allows you to hide your unregistered private IP addresses behind a set of registered IP addresses. This helps to protect your internal network from outside networks. NAT also helps to alleviate the IP address depletion problem, since many private addresses can be represented by a small set of registered addresses.

Unfortunately, conventional NAT does not work on IPSec packets because when the packet goes through a NAT device, the source address in the packet changes, thereby invalidating the packet. When this happens, the receiving end of the VPN connection discards the packet and the VPN connection negotiations fail.

The solution: UDP encapsulation

In a nutshell, UDP encapsulation wraps an IPSec packet inside a new, but duplicate, IP/UDP header. The address in the new IP header gets translated when it goes through the NAT device. Then, when the packet reaches its destination, the receiving end strips off the additional header, leaving the original IPSec packet, which will now pass all other validations.

You can only apply UDP encapsulation to VPNs that will use IPSec ESP in either tunnel mode or transport mode. In addition, the system can only act as a client for UDP encapsulation. That is, it can only *initiate* UDP encapsulated traffic.

The graphics below illustrate the format of a UDP encapsulated ESP packet in tunnel mode:

Original IPv4 datagram:



After applying IPSec ESP in tunnel mode:



After applying UDP Encapsulation:




The graphics below illustrate the format of a UDP encapsulated ESP packet in transport mode:

Original IPv4 datagram:



After applying IPSec ESP in transport mode:



After applying UDP Encapsulation:



After the packet is encapsulated, the system sends the packet to its VPN partner over UDP port 4500. Typically, VPN partners perform IKE negotiations over UDP port 500. However, when IKE detects NAT during key negotiation, subsequent IKE packets are sent over source port 4500, destination port 4500. This also means that port 4500 must be unrestricted in any applicable filter rules. The receiving end of the connection can determine whether the packet is an IKE packet or a UDP encapsulated packet because the first 4 bytes of the UDP payload are set to zero on an IKE packet. For it to work properly, both ends of the connection must support UDP encapsulation.

Related concepts:

“Scenario: Firewall friendly VPN” on page 33

In this scenario, a large insurance company wants to establish a VPN between a gateway in Chicago and a host in Minneapolis when both networks are behind a firewall.

IP Compression

The IP Payload Compression protocol (IPComp) reduces the size of IP datagrams by compressing the datagrams to increase the communication performance between two partners.

The intent is to increase overall communication performance when the communication is over slow or congested links. IPComp does not provide any security and must be used along with either an AH or an ESP transform when the communication occurs over a VPN connection.

The Internet Engineering Task Force (IETF) formally defines IPComp in Request for Comments (RFC) 2393, *IP Payload compression Protocol (IPComp)*. You can view this RFC on the Internet at the following Web site: <http://www.rfc-editor.org>.

Related information:

 <http://www.rfc-editor.org>

VPN and IP filtering

IP filtering and VPN are closely related. In fact, most VPN connections require filter rules to work properly. This topic provides you information about what filters VPN requires, as well as other filtering concepts related to VPN.

Most VPN connections require filter rules to work properly. The filter rules required depend on the type of VPN connection that you are configuring as well as what type of traffic you want to control. In general, each connection will have a policy filter. The policy filter defines which addresses, protocols, and ports can use the VPN. In addition, connections that support the Internet Key Exchange (IKE) protocol typically have rules that are written explicitly to allow IKE processing over the connection. VPN can generate these rules automatically. Whenever possible, allow VPN to generate your policy filters for you. Not only will this help to eliminate errors, but it also eliminates the need for you to configure the rules as a separate step by using the Packet Rules editor in IBM Navigator for i.

There are, of course, exceptions. Review these topics to learn more about other, less common, VPN and filtering concepts and techniques that might apply to your particular situation:

Related concepts:

“Configuring VPN packet rules” on page 53

If you are creating a connection for the first time, allow VPN to automatically generate the VPN packet rules for you. You can do this by either using the New Connection wizard or the VPN properties pages to configure your connection.

VPN connections with no policy filters

If the connection endpoints of your VPN are single, specific, IP addresses and you want to start the VPN without having to write or activate filter rules on the system, you can configure a dynamic policy filter.

A policy filter rule defines which addresses, protocols, and ports can use a VPN and directs the appropriate traffic through the connection. In some cases, you might want to configure a connection that does not require a policy filter rule. For example, you might have non-VPN packet rules loaded on the interface that your VPN connection will use, so rather than deactivating the active rules on that interface, you decide to configure the VPN so that your system manages all filters dynamically for the connection. The policy filter for this type of connection is referred to as a **dynamic policy filter**. Before you can use a dynamic policy filter for your VPN connection, all of the following must be true:

- The connection can only be initiated by the local system.
- The data endpoints of the connection must be single systems. That is, they cannot be a subnet or a range of addresses.
- No policy filter rule can be loaded for the connection.

If your connection meets this criteria, then you can configure the connection so that it does not require a policy filter. When the connection starts, traffic between the data endpoints will flow across the it regardless of what other packet rules are loaded on your system.

For step-by-step instructions on how to configure a connection so that it does not require a policy filter, use the online help for VPN.

Implicit IKE

In order for Internet Key Exchange (IKE) negotiations to occur for your VPN, you need to allow UDP datagrams over port 500 for this type of IP traffic. However, if there are no filter rules on the system specifically written to permit IKE traffic, then the system will implicitly allow IKE traffic to flow.

To establish a connection, most VPNs require IKE negotiations to occur before IPSec processing can happen. IKE uses the well-known port 500, so for IKE to work properly, you need to allow UDP datagrams over port 500 for this type of IP traffic. If there are no filter rules on the system specifically

written to permit IKE traffic, then IKE traffic is implicitly allowed. However, rules written specifically for UDP port 500 traffic are handled based on what is defined in the active filter rules.

Scenarios: VPN

Review these scenarios to become familiar with the technical and configuration details involved with each of these basic connection types.

Related concepts:

QoS scenario: Secure and predictable results (VPN and QoS)

Related information:

➡ OS/400 V5R1 Virtual Private Networks: Remote Access to the IBM e(logo)server iSeries Server with Windows 2000 VPN Clients, REDP0153

➡ AS/400 Internet Security: Implementing AS/400 Virtual Private Networks, SG24-5404-00

➡ AS/400 Internet Security Scenarios: A Practical Approach, SG24-5954-00

Scenario: Basic branch office connection

In this scenario, your company wants to establish a VPN between the subnets of two remote departments through a pair of IBM i models acting as VPN gateways.

Situation

Suppose your company wants to minimize the costs incurred from communicating to and among its own branches. Today, your company uses frame relay or leased lines, but you want to explore other options for transmitting internal confidential data that are less expensive, more secure, and globally accessible. By exploiting the Internet, you can easily establish a virtual private network (VPN) to meet the needs of your company.

Your company and its branch office both require VPN protection across the Internet, but not within their respective intranets. Because you consider the intranets trusted, the best solution is to create a gateway-to-gateway VPN. In this case, both gateways are connected directly to the intervening network. In other words, they are *border* or *edge* systems, which are not protected by firewalls. This example serves as a useful introduction to the steps involved in setting up a basic VPN configuration. When this scenario refers to the term, *Internet*, it refers to the intervening network between the two VPN gateways, which might be the company's own private network or the public Internet.

Important: This scenario shows the IBM i model security gateways attached directly to the Internet. The absence of a firewall is intended to simplify the scenario. It does not imply that the use of a firewall is not necessary. In fact, consider the security risks involved any time you connect to the Internet.

Advantages

This scenario has the following advantages:

- Using the Internet or an existing intranet reduces the cost of private lines between remote subnets.
- Using the Internet or an existing intranet reduces the complexity of installing and maintaining private lines and associated equipment.
- Using the Internet allows remote locations to connect to almost anywhere in the world.
- Using VPN provides users access to all systems and resources on either side of the connection just as though they were connected using a leased line or wide area network (WAN) connection.
- Using industry standard encryption and authentication methods ensures the security of sensitive information passed from one location to another.

- Exchanging your encryption keys dynamically and regularly simplifies setup and minimizes the risk of your keys being decoded and security being breached.
- Using private IP addresses in each remote subnet makes it unnecessary to allocate valuable public IP addresses to each client.

Objectives


In this scenario, MyCo, Inc. wants to establish a VPN between the subnets of its Human Resources and Finance departments through a pair of IBM i models. Both systems will act as VPN gateways. In terms of VPN configurations, a gateway performs key management and applies IPSec to the data that flows through the tunnel. The gateways are not the data endpoints of the connection.

The objectives of this scenario are as follows:

- The VPN must protect all data traffic between the Human Resources department's subnet and the Finance department's subnet.
- Data traffic does not require VPN protection once it reaches either of the department's subnets.
- All clients and hosts on each network have full access to the other's network, including all applications.
- The gateway systems can communicate with each other and access each other's applications.

Details

The following figure illustrates the network characteristics of MyCo.



Human Resources Department

- System A acts as the Human Resources Department's VPN gateway.
- Subnet is 10.6.0.0 with mask 255.255.0.0. This subnet represents the data endpoint of the VPN tunnel at the MyCo Rochester site.
- System A connects to the Internet with IP address 204.146.18.227. This is the connection endpoint. That is, System A performs key management and applies IPSec to incoming and outgoing IP datagrams.
- System A connects to its subnet with IP address 10.6.11.1.
- System B is a production system in the Human Resources subnet that runs standard TCP/IP applications.

Finance Department

- System C acts as the Finance Department's VPN gateway.
- Subnet is 10.196.8.0 with mask 255.255.255.0. This subnet represents the data endpoint of the VPN tunnel at the MyCo Endicott site.
- System C connects to the Internet with IP address 208.222.150.250. This is the connection endpoint. That is, System C performs key management and applies IPSec to incoming and outgoing IP datagrams.
- System C connects to its subnet with IP address 10.196.8.5.

Configuration tasks

You must complete each of these tasks to configure the branch office connection described in this scenario:

Note: Before you start these tasks verify the TCP/IP routing to ensure that the two gateway systems can communicate with each other across the Internet. This ensures that hosts on each subnet route properly to their respective gateway for access to the remote subnet.

Related concepts:

TCP/IP routing and workload balancing

Related information:

 AS/400 Internet Security Scenarios: A Practical Approach, SG24-5954-00

Completing the planning worksheets

The planning checklists illustrate the type of information you need before you begin configuring the VPN. All answers on the prerequisite checklist must be YES before you proceed with VPN setup.

Note: These worksheets apply to System A, repeat the process for System C, reversing IP addresses as necessary.

Table 1. System requirements

Prerequisite checklist	Answers
Is the Digital Certificate Manager option installed?	Yes
Is HTTP server started (to support IBM Navigator for i)?	Yes
Is IBM TCP/IP Connectivity Utilities for i installed?	Yes
Did you set the retain server security data (QRETSVRSEC *SEC) system value to 1?	Yes
Is TCP/IP configured on your system (including IP interfaces, routes, local host name, and local domain name)?	Yes
Is normal TCP/IP communication established between the required endpoints?	Yes
Have you applied the latest program temporary fixes (PTFs)?	Yes
If the VPN tunnel traverses firewalls or routers that use IP packet filtering, do the firewall or router filter rules support AH and ESP protocols?	Yes
Are the firewalls or routers configured to permit IKE (UDP port 500), AH, and ESP protocols?	Yes
Are the firewalls configured to enable IP forwarding?	Yes

Table 2. VPN configuration

You need this information to configure the VPN	Answers
What type of connection are you creating?	gateway-to-gateway
What will you name the dynamic-key group?	HRgw2FINgw

Table 2. VPN configuration (continued)

You need this information to configure the VPN	Answers
What type of security and system performance do you require to protect your keys?	balanced
Are you using certificates to authenticate the connection? If no, what is the preshared key?	No topsecretstuff
What is the identifier of the local key server?	IP address: 204.146.18.227
What is the identifier of the local data endpoint?	Subnet: 10.6.0.0 Mask: 255.255.0.0
What is the identifier of the remote key server?	IP address: 208.222.150.250
What is the identifier of the remote data endpoint?	Subnet: 10.196.8.0 Mask: 255.255.255.0
What ports and protocols do you want to allow to flow through the connection?	Any
What type of security and system performance do you require to protect your data?	balanced
To which interfaces does the connection apply?	TRLINE

Configuring VPN on System A

Complete these task to configure System A

Use the following steps and the information from your worksheets to configure VPN on System A:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
2. Click **Secure Connections** to open the **Connections** panel.
3. Click **Actions > New Connection** to start the New Connection wizard.
4. Review the **Welcome** page for information about what objects the wizard creates.
5. Click **Next** to go to the **Connection Name** page
6. In the **Name** field, enter HRgw2FINGw.
7. Optional: Specify a description for this connection group.
8. Click **Next** to go to the **Connection Scenario** page.
9. Select **Connect your gateway to another gateway**.
10. Click **Next** to go to the **Internet Key Exchange Policy** page.
11. Select **Create a new policy** and then select **Balance security and performance**.
12. Click **Next** to go to the **Certificate for Local Connection Endpoint** page.
13. Select **No** to indicate that you will not be using certificates to authenticate the connection.
14. Click **Next** to go to the **Local Key Server** page.
15. Select **Version 4 IP address** from the **Identifier type** field.
16. Select 204.146.18.227 from the **IP address** field.
17. Click **Next** to go to the **Remote Key Server** page.
18. Select **Version 4 IP address** in the **Identifier type** field.
19. Enter 208.222.150.250 in the **Identifier** field.
20. Enter topsecretstuff in the **Pre-shared key** field
21. Click **Next** to go to the **Local Data Endpoint** page.
22. Select **IP version 4 subnet** from the **Identifier type** field.
23. Enter 10.6.0.0 in the **Identifier** field.
24. Enter 255.255.0.0 in the **Subnet mask** field.
25. Click **Next** to go to the **Remote Data Endpoint** page.
26. Select **IP version 4 subnet** from the **Identifier type** field

27. Enter 10.196.8.0 in the **Identifier** field.
28. Enter 255.255.255.0 in the **Subnet mask** field.
29. Click **Next** to go to the **Data Services** page.
30. Accept the default values, and then click **Next** to go to the **Data Policy** page.
31. Select **Create a new policy** and then select **Balance security and performance**.
32. Click **Next** to go to the **Applicable Interfaces** page.
33. Select **TRLINE** from the **Line** table.
34. Click **Next** to go to the **Summary** page. Review the objects that the wizard will create to ensure they are correct.
35. Click **Finish** to complete the configuration.
36. When the **Activate Policy Filters** panel appears, select **Yes, activate the generated policy filters** then select **Permit all other traffic**.
37. Click **OK** to complete the configuration. When prompted, specify that you want to activate the rules on all interfaces.

Related tasks:

“Configuring VPN on System C”

Follow the same steps you used to configure VPN on System A , changing IP addresses as necessary. Use your planning worksheets for guidance.

Configuring VPN on System C

Follow the same steps you used to configure VPN on System A , changing IP addresses as necessary. Use your planning worksheets for guidance.

When you finish configuring the Finance Department VPN gateway, your connections will be in an *on-demand* state, which means the connection starts when IP datagrams that this VPN connection must protect are sent. The next step is to start the VPN servers, if they are not already started.

Related tasks:

“Configuring VPN on System A” on page 20

Complete these task to configure System A

Starting VPN

After you have configured your VPN connection on System A and C you need to start your VPN connection.

Follow these steps to start VPN:

1. In IBM Navigator for i, expand the **Network > All Tasks > IP Policies > Virtual Private Networking**.
2. Click **Start VPN Server**.

Testing a connection

After you finish configuring both systems and you have successfully started the VPN servers, test the connectivity to ensure that the remote subnets can communicate with each other.

To test your connection, follow these steps:

1. In IBM Navigator for i, expand **Network > All Tasks > TCP/IP Configuration**.
2. Click **Work with Ping**.
3. From the **Ping** panel, enter System C in the **IP address or host name** field.
4. Click **Ping Now** to verify connectivity from System A to System C.
5. Click **OK** when you are finished.

Scenario: Basic business to business connection

In this scenario, your company wants to establish a VPN between a client workstation in your manufacturing division and a client workstation in the supply department of your business partner.

Situation

Suppose you are a major parts supplier to a manufacturer. Since it is critical that you have the specific parts and quantities at the exact time required by the manufacturing firm, you always need to be aware of the manufacturer's inventory status and production schedules. Perhaps you handle this interaction manually today, and find it time consuming, expensive and even inaccurate at times. You want to find an easier, faster, and more effective way to communicate with your manufacturing company. However, given the confidentiality and time-sensitive nature of the information you exchange, the manufacturer does not want to publish it on its corporate Web site or distribute it monthly in an external report. By exploiting the public Internet, you can easily establish a VPN to meet the needs of both companies.

Objectives


In this scenario, MyCo wants to establish a VPN between a host in its parts division and a host in the manufacturing department of one their business partners, TheirCo.

Because the information these two companies share is highly confidential, it must be protected while it travels across the Internet. In addition, data must not flow in the clear within either company's networks because each network considers the other untrusted. In other words, both companies require end-to-end authentication, integrity, and encryption.

Important: The intent of this scenario is to introduce, by example, a simple host-to-host VPN configuration. In a typical network environment, you will also need to consider firewall configuration, IP addressing requirements, and routing, among others.

Details

The following figure illustrates the network characteristics of MyCo and TheirCo:



MyCo Supply Network

- System A has an IP address of 10.6.1.1. This is the connection endpoint, as well as the data endpoint. That is, System A performs IKE negotiations and applies IPSec to incoming and outgoing IP datagrams and is also the source and destination for data that flows through the VPN.
- System A is in subnet 10.6.0.0 with mask 255.255.0.0
- Only System A can initiate the connection with System C.

TheirCo Manufacturing Network

- System C has an IP address of 10.196.8.6. This is the connection endpoint, as well as the data endpoint. That is, System C performs IKE negotiations and applies IPSec to incoming and outgoing IP datagrams and is also the source and destination for data that flows through the VPN.
- System C is in subnet 10.196.8.0 with mask 255.255.255.0

Configuration tasks

You must complete each of these tasks to configure the business to business connection described in this scenario:

Note: Before you start these tasks verify the TCP/IP routing to ensure that the two gateway systems can communicate with each other across the Internet. This ensures that hosts on each subnet route properly to their respective gateway for access to the remote subnet.

Related concepts:

TCP/IP routing and workload balancing

Completing the planning worksheets

The planning checklists illustrate the type of information you need before you begin configuring the VPN. All answers on the prerequisite checklist must be YES before you proceed with VPN setup.

Note: These worksheets apply to System A, repeat the process for System C, reversing IP addresses as necessary.

Table 3. System requirements

Prerequisite checklist	Answers
Is the Digital Certificate Manager option installed?	Yes
Is HTTP server started (to support IBM Navigator for i)?	Yes
Is IBM TCP/IP Connectivity Utilities for i installed?	Yes
Did you set the retain server security data (QRETSVRSEC *SEC) system value to 1?	Yes
Is TCP/IP configured on your system (including IP interfaces, routes, local host name, and local domain name)?	Yes
Is normal TCP/IP communication established between the required endpoints?	Yes
Have you applied the latest program temporary fixes (PTFs)?	Yes
If the VPN tunnel traverses firewalls or routers that use IP packet filtering, do the firewall or router filter rules support AH and ESP protocols?	Yes
Are the firewalls or routers configured to permit IKE (UDP port 500), AH, and ESP protocols?	Yes
Are the firewalls configured to enable IP forwarding?	Yes

Table 4. VPN configuration

You need this information to configure the VPN	Answers
What type of connection are you creating?	gateway-to-gateway
What will you name the dynamic-key group?	HRgw2FINgw
What type of security and system performance do you require to protect your keys?	balanced
Are you using certificates to authenticate the connection? If no, what is the preshared key?	No topsecretstuff
What is the identifier of the local key server?	IP address: 204.146.18.227
What is the identifier of the local data endpoint?	Subnet: 10.6.0.0 Mask: 255.255.0.0
What is the identifier of the remote key server?	IP address: 208.222.150.250
What is the identifier of the remote data endpoint?	Subnet: 10.196.8.0 Mask: 255.255.255.0
What ports and protocols do you want to allow to flow through the connection?	Any
What type of security and system performance do you require to protect your data?	balanced
To which interfaces does the connection apply?	TRLINE

Configuring VPN on System A

Complete the following steps to configure a VPN connection on System A.

Use the information from your planning worksheets to configure VPN on System A as follows:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
2. Click **Secure Connections** to open the **Connections** panel.

3. Click **Actions > New Connections** to start the New Connection wizard.
4. Review the **Welcome** page for information about what objects the wizard creates.
5. Click **Next** to go to the **Connection Name** page.
6. In the **Name** field, enter MyCo2TheirCo.
7. Optional: Specify a description for this connection group.
8. Click **Next** to go to the **Connection Scenario** page.
9. Select **Connect your host to another host**.
10. Click **Next** to go to the **Internet Key Exchange Policy** page.
11. Select **Create a new policy** and then select **Highest security, lowest performance**.
12. Click **Next** to go to the **Certificate for Local Connection Endpoint** page.
13. Select **Yes** to indicate that you will be using certificates to authenticate the connection. Then, select the certificate that represents System A.

Note: If you want to use a certificate to authenticate the local connection endpoint, you must first create the certificate in the Digital Certificate Manager (DCM).

14. Click **Next** to go to the **Remote Key Server** page.
15. Select **Version 4 IP address** in the **Identifier type** field.
16. Enter 10.196.8.6 in the **Identifier** field.
17. Click **Next** to go to the **Data Services** page.
18. Accept the default values, and then click **Next** to go to the **Data Policy** page.
19. Select **Create a new policy** and then select **Highest security, lowest performance**.
20. Click **Next** to go to the **Applicable Interfaces** page.
21. Select **TRLIN**.
22. Click **Next** to go to the **Summary** page. Review the objects that the wizard will create to ensure they are correct.
23. Click **Finish** to complete the configuration.
24. When the **Activate Policy Filters** dialog box appears, select **No, packet rules will be activated at a later time** and then click **OK**.

The next step is to specify that only System A can initiate this connection. Do this by customizing the properties of the dynamic-key group, MyCo2TheirCo, that the wizard created:

1. Click **By Group** in the left pane of the VPN interface, the new dynamic-key group, MyCo2TheirCo, displays in the right pane. Right-click it and select **Properties**.
2. Go to the **Policy** page and select the **Local system initiates connection** option.
3. Click **OK** to save your changes.

Configuring VPN on System C

Follow the same steps you used to configure VPN on System A, changing IP addresses as necessary. Use your planning worksheets for guidance.

When you finish configuring the Finance Department VPN gateway, your connections will be in an *on-demand* state, which means the connection starts when IP datagrams that this VPN connection must protect are sent. The next step is to start the VPN servers, if they are not already started.

Activating packet rules

The VPN wizard automatically creates the packet rules that this connection requires to work properly. However, you must activate them on both systems before you can start the VPN connection.

To activate packet rules on System A, follow these steps:

1. In IBM Navigator for i, expand **Network > IP Policies**.

- | 2. Click **Packet Rules** to open the Packet Rules panel.
- | 3. Click **Actions** and then click **Activate Rules....** This opens the **Activate Packet Rules** panel.
- | 4. Select whether you want to activate only the VPN generated rules, only the selected files, or both the VPN generated rules and the selected files. You might choose the latter, for instance, if you have miscellaneous PERMIT and DENY rules that you want to enforce on the interface in addition to the VPN generated rules.
- | 5. Select the interface on which you want the rules activated. In this case, select **Activate these rules on all interfaces and all point-to-point filter identifier**.
- | 6. Click **OK** on the dialog box to confirm that you want to verify and activate the rules on the interface or interfaces you specified. After you click OK, the system checks the rules for syntax and semantic errors and reports the results in a message window at the bottom of the editor.
- | 7. Repeat these steps to activate packet rules on System C.

Starting a connection

After you have configured your VPN connection you need to start your VPN connection.

Follow these steps to start the MyCo2TheirCo connection from System A:

- | 1. In IBM Navigator for i, expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
- | 2. If the VPN server is not started, click **Start VPN Server**. This starts the VPN server.
- | 3. Expand **Network > IP Policies > Virtual Private Networking**.
- | 4. Click **Secure Connections**.
- | 5. Right-click **All Connections** and select **Open** to display a list of connections.
- | 6. Right-click **MyCo2TheirCo** and select **Start**.
- | 7. From the **View** menu, select **Refresh**. If the connection starts successfully, the status will change from *Idle* to *Enabled*. The connection might take a few minutes to start, so periodically refresh until the status changes to *Enabled*.

Testing a connection

After you finish configuring both systems and you have successfully started the VPN servers, test the connectivity to ensure that the remote subnets can communicate with each other.

To test your connection, follow these steps:

- | 1. In IBM Navigator for i, expand **Network > All Tasks > TCP/IP Configuration**.
- | 2. Click **Work with Ping**.
- | 3. From the **Ping** panel, enter System C in the **IP address or host name** field.
- | 4. Click **Ping Now** to verify connectivity from System A to System C.
- | 5. Click **OK** when you are finished.

Scenario: Protecting an L2TP voluntary tunnel with IPSec

In this scenario, you learn how to setup a connection between a branch office host and a corporate office that uses L2TP protected by IPSec. The branch office has a dynamically assigned IP address, while the corporate office has a static, globally routable IP address.

Situation

Suppose your company has a small branch office in another state. Throughout any given workday the branch office might require access to confidential information about an IBM i model within your corporate intranet. Your company currently uses an expensive leased line to provide the branch office access to the corporate network. Although your company wants to continue providing secure access to your intranet, you ultimately want to reduce the expense associated with the leased line. This can be

done by creating a Layer 2 Tunnel Protocol (L2TP) voluntary tunnel that extends your corporate network, such that the branch office appears to be part of your corporate subnet. VPN protects the data traffic over the L2TP tunnel.

With an L2TP voluntary tunnel, the remote branch office establishes a tunnel directly to the L2TP network server (LNS) of the corporate network. The functionality of the L2TP access concentrator (LAC) resides at the client. The tunnel is transparent to the remote client's Internet Service Provider (ISP), so the ISP is not required to support L2TP. If you want to read more about L2TP concepts, see Layer 2 Tunnel Protocol (L2TP).

Important: This scenario shows the security gateways attached directly to the Internet. The absence of a firewall is intended to simplify the scenario. It does not imply that the use of a firewall is not necessary. Consider the security risks involved any time you connect to the Internet.

Objectives


In this scenario, a branch office system connects to its corporate network through a gateway system with an L2TP tunnel protected by VPN.

The main objectives of this scenario are:

- The branch office system always initiates the connection to the corporate office.
- The branch office system is the only system at the branch office network that needs access to the corporate network. In other words, its role is that of a host, not a gateway, in the branch office network.
- The corporate system a host computer in the corporate office network.

Details

The following figure illustrates the network characteristics for this scenario:



System A

- Must have access to TCP/IP applications on all systems in the corporate network.
- Receives dynamically assigned IP addresses from its ISP.
- Must be configured to provide L2TP support.

System B

- Must have access to TCP/IP applications on System A.
- Subnet is 10.6.0.0 with mask 255.255.0.0. This subnet represents the data endpoint of the VPN tunnel at the corporate site.
- Connects to the Internet with IP address 205.13.237.6. This is the connection endpoint. That is, System B performs key management and applies IPSec to incoming and outgoing IP datagrams. System B connects to its subnet with IP address 10.6.11.1.

In L2TP terms, *System A* acts as the L2TP initiator, while *System B* acts as the L2TP terminator.

Configuration tasks

Assuming that TCP/IP configuration already exists and works, you must complete the following tasks:

Related concepts:

“Layer 2 Tunnel Protocol” on page 12

Layer 2 Tunneling Protocol (L2TP) connections, which are also called virtual lines, provide cost-effective access for remote users by allowing a corporate network systems to manage the IP addresses assigned to its remote users. Further, L2TP connections provide secure access to your system or network when you use them in conjunction with IP Security (IPSec).

Related information:

 AS/400 Internet Security Scenarios: A Practical Approach, SG24-5954-00

Configuring VPN on System A

Complete the following steps to configure a VPN connection on System A.

Use the information from your planning worksheets to configure VPN on System A as follows:

1. Configure the Internet Key Exchange policy

- a. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
- b. Click **IP Security Policies** to open the **IP Security Policies** panel.
- c. Right-click **Internet Key Exchange Policies** and select **New Internet Key Exchange Policy**.
- d. On the **Remote Server** page, select **Version 4 IP address** as the identifier type and then enter 205.13.237.6 in the **IP address** field.
- e. On the **Associations** page, select **Preshared key** to indicate that this connection uses a preshared key to authenticate this policy.
- f. Enter the preshared key in the **Key** field. Treat your preshared key like a password.
- g. Select **Key Identifier** for the local key server identifier type, and then enter the key identifier in the **Identifier** field. For example, thisisthekeyid. Remember that the local key server has a dynamically assigned IP address which is impossible to know in advance. System B uses this identifier to identify the System A when System A initiates a connection.
- h. On the **Transforms** page, click **Add** to add the transforms that System A proposes to System B for key protection and to specify whether the IKE policy uses identity protection when initiating phase 1 negotiations.
- i. On the **IKE Policy Transform** page, select **Preshared key** for your authentication method, **SHA** for your hash algorithm and PRF algorithm, and **3DES-CBC** for your encryption algorithm. Accept the default values for Diffie-Hellman group and Expire IKE keys after.
- j. Click **OK** to return to the **Transforms** page.
- k. Select **IKE aggressive mode negotiation (no identity protection)**.

Note: If you use preshared keys and aggressive mode negotiation together in your configuration, select obscure passwords that are unlikely to be cracked in attacks that scan the dictionary. It is also recommended you periodically change your passwords.

- l. Click **OK** to save your configurations.

2. Configure the data policy

- a. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
- b. Click **IP Security Policies** to open the **IP Security Policies** panel.
- c. Right-click **Data policies** and select **New Data Policy**
- d. On the **General** page, specify the name of the data policy. For example, **l2tpremoteuser**
- e. Go to the **Proposals** page. A proposal is a collection of protocols that the initiating and responding key servers use to establish a dynamic connection between two endpoints. You can use a single data policy in several connection objects. However, not all remote VPN key servers necessarily have the same data policy properties. Therefore, you can add several proposals to one data policy. When establishing a VPN connection to a remote key server, there must be at least one matching proposal in the data policy of the initiator and the responder.
- f. Click **Add** to add a data policy transform
- g. Select **Transport** for the encapsulation mode.
- h. Specify a key expiration date.
- i. Click **Transforms** tab.
- j. Click **Add** to add a transform. A transform defines the protocols, authentication algorithms, and encryption algorithms this data policy uses during IKE phase 2 negotiations. The initiator of a connection sends one or more data policy proposals to the responder. The responder then selects a matching proposal in its associated security policy to complete the secure connection.
- k. Click **OK** to save the transform.
- l. Click **OK** to save your new data policy.

3. Configure the dynamic-key group

- a. Click **Secure Connections** under VPN interface.
- b. Right-click **By Group** and select **New Dynamic-Key Group**.
- c. On the **General** page, specify a name for the group. For example, **l2tptocorp**.
- d. Select **Protects a locally initiated L2TP tunnel**.
- e. For system role, select **Both systems are hosts**.
- f. Go to the **Policy** page. Select the data policy you created in the step **Configure the data policy**, **l2tpremoteuser**, from the **Data policy** drop-down list.
- g. Select **Local system initiates connection** to indicate that only System A can initiate connections with System B.
- h. Go to the **Connections** page. Select **Generate the following policy filter rule for this group**. Click **Edit** to define the parameters of the policy filter.
- i. On the **Policy Filter- Local Addresses** page, select **Key Identifier** for the identifier type.
- j. For the identifier, select the key identifier, **thisisthekeyid**, that you defined in the IKE policy.
- k. Go the **Policy Filter - Remote Addresses** page. Select **IP version 4 address** from the **Identifier type** drop-down list.
- l. Enter **205.13.237.6** in the **Identifier** field.
- m. Go to the **Policy Filter - Services** page. Enter **1701** in the **Local Port** and **Remote Port** fields. Port **1701** is the well-known port for L2TP.
- n. Select **UDP** from the **Protocol** drop-down list.
- o. Click **OK** to return to the **Connections** page.
- p. Go to the **Interfaces** page. Select any line or PPP profile to which this group will apply. You have not created the PPP profile for this group yet. After you do so, you will need to edit the properties of this group so that the group applies to the PPP profile you create in the next step.
- q. Click **OK** to create the dynamic-key group, **l2tptocorp**.

4. Configure the dynamic-key connection

- a. From the **Connections** panel, right-click **By Group** and select **Open**. This displays a list of all dynamic-key groups you have configured on System A.
- b. Right-click **l2tpocorp** and select **New Dynamic-Key Connection**.
- c. On the **General** page, specify an optional description for the connection.
- d. For the remote key server, select **Version 4 IP address** for the identifier type.
- e. Select 205.13.237.6 from the **IP address** drop-down list.
- f. Deselect **Start on-demand**.
- g. Go to the **Local Addresses** page. Select **Key identifier** for the identifier type and then select thisisthekeyid from the **Identifier** drop-down list.
- h. Go to the **Remote Addresses** page. Select **IP version 4 address** for the identifier type.
- i. Enter 205.13.237.6 in the **Identifier** field.
- j. Go to the **Services** page. Enter 1701 in the **Local Port** and **Remote Port** fields. Port 1701 is the well-known port for L2TP.
- k. Select **UDP** from the **Protocol** drop-down list.
- l. Click **OK** to create the dynamic-key connection.

Related tasks:

“Configuring VPN on System B” on page 31

To configure a VPN connection on System B follow the same steps you used to configure a VPN connection on System A, and change IP addresses and identifiers as necessary.

Configuring a PPP connection profile and virtual line on System A

Now that a VPN connection is configured on System A you need to create the PPP profile for System A. The PPP profile has no physical line associated with it; instead, it uses a virtual line. This is because the PPP traffic tunnels through the L2TP tunnel, while VPN protects the L2TP tunnel.

Follow these steps to create a PPP connection profile for System A:

1. In IBM Navigator for i, expand **Network > Remote Access Services**.
2. Right-click **Originator Connection Profiles** and click **Actions > New Profile**.
3. On the **Setup** page, select **PPP** for the protocol type.
4. For Mode selections, select **L2TP (virtual line)**.
5. Select **Initiator on-demand (voluntary tunnel)** from the **Operating mode** drop-down list.
6. Click **OK** to go to the PPP profiles properties pages.
7. On the **General** page, enter a name that identifies the type and the destination of the connection. In this case, enter **toCORP**. The name you specify must be 10 characters, or less.
8. Optional: Specify a description for the profile.
9. Go to the **Connection** page.
10. In the **Virtual line name** field, select **tocorp** from the drop-down list. Remember that this line has no associated physical interface. The virtual line describes various characteristics of this PPP profile; for example, the maximum frame size, authentication information, the local host name, and so on. The **L2TP Line Properties** dialog box opens.
11. On the **General** page, enter a description for the virtual line.
12. Go to the **Authentication** page.
13. In the **Local host name** field, enter the host name of the local key server, SystemA.
14. Click **OK** to save the new virtual line description and return to the **Connection** page.
15. Enter the remote tunnel endpoint address, 205.13.237.6, in the **Remote tunnel endpoint address** field.
16. Select **Requires IPSec Protection** and select the dynamic-key group you created in previous step “Configuring VPN on System A” on page 28, **l2tpocorp** from the **Connection group name** drop-down list.

17. Go to the **TCP/IP IPv4 Settings** page.
18. In the **Local IP address** section, select **Assigned by remote system**.
19. In the **Remote IP address** section, select **Use fixed IP address**. Enter 10.6.11.1, which is the remote system's IP address in its subnet.
20. In the routing section, select **Define additional static routes** and click **Routes**. If there is no routing information provided in the PPP profile, then System A is only able to reach the remote tunnel endpoint but not any other system on the 10.6.0.0 subnet.
21. Click **Add** to add a static route entry.
22. Enter the subnet, 10.6.0.0, and the subnet mask, 255.255.0.0 to route all 10.6.*.* traffic through the L2TP tunnel.
23. Click **OK** to add the static route.
24. Go to the **Authentication** page to set the user name and password for this PPP profile.
25. In the Local system identification section, select **Allow the remote system to verify the identity of this system**.
26. Under **Authentication protocol to use** select **Require encrypted password (CHAP-MD5)**. In the Local system identification section, select **Allow the remote system to verify the identity of this system**.
27. Enter the user name, SystemA, and a password.
28. Click **OK** to save the PPP profile.
29. Enter the password again to confirm the password.

Applying the 12tptocorp dynamic-key group to the toCorp PPP profile

After you have your PPP connection profile configured, you need to go back to the dynamic-key group, 12tptocorp, you created and associate it with the PPP profile.

To associate your dynamic-key group with your PPP profile follow these steps:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
2. Click **Secure Connections** to open the **Connections** panel and right-click **By Group** and then select **Open**.
3. Right-click the dynamic-key group, 12tptocorp, and select **Properties**.
4. Go to the **Interfaces** page and select **Apply this group** for the PPP profile you created in "Configuring a PPP connection profile and virtual line on System A" on page 30, toCorp.
5. Click **OK** to apply 12tptocorp to the PPP profile, toCorp.

Configuring VPN on System B

To configure a VPN connection on System B follow the same steps you used to configure a VPN connection on System A, and change IP addresses and identifiers as necessary.

Take these other points into consideration before you begin:

- The identify the remote key server by the key identifier you specified for the local key server on System A. For example, thisisthekeyid.
- Use the *exact* same preshared key.
- Make sure your transforms match the ones you configured on System A, or connections will fail.
- Do not specify **Protects a locally initiated L2TP tunnel** on the **General** page of the dynamic-key group.
- Remote system initiates the connection.
- Specify that the connection should start on-demand.

Related tasks:

"Configuring VPN on System A" on page 28

Complete the following steps to configure a VPN connection on System A.

Configuring a PPP connection profile and virtual line on System B

Now that a VPN connection is configured on System B you need to create the PPP profile for System B. The PPP profile has no physical line associated with it; instead, it uses a virtual line. This is because the PPP traffic tunnels through the L2TP tunnel, while VPN protects the L2TP tunnel.

Follow these steps to create a PPP connection profile for System B:

1. In IBM Navigator for i, expand **Network > Remote Access Services**.
2. Click **Receiver Connection Profiles** to open the **Receiver Connection Profiles** panel and click **Actions > New Profile**.
3. On the **Setup** page, select **PPP** for the protocol type.
4. For Mode selections, select **L2TP (virtual line)**.
5. Select **Terminator (network server)** from the **Operating mode** drop-down list.
6. Click **OK** to PPP profiles properties pages.
7. On the **General** page, enter a name that identifies the type and the destination of the connection. In this case, enter **tobran**. The name you specify must be 10 characters, or less.
8. Optional: Specify a description for the profile
9. Go to the **Connection** page.
10. Select the IP address of the local tunnel endpoint, 205.13.237.6.
11. In the **Virtual line name** field, select **tobran** from the drop-down list. Remember that this line has no associated physical interface. The virtual line describes various characteristics of this PPP profile; for example, the maximum frame size, authentication information, the local host name, and so on. Click **Open** next to **Virtual line name** field to open the **L2TP Line Properties** panel.
12. On the **General** page, enter a description for the virtual line.
13. Go to the **Authentication** page
14. In the **Local host name** field, enter the host name of the local key server, SystemB.
15. Click **OK** to save the new virtual line description and return to the **Connection** page.
16. Go to the **TCP/IP Settings** page.
17. In the **Local IP address** section, select the fixed IP address of the local system, 10.6.11.1.
18. In the **Remote IP address** section, select **Address pool** as the address assignment method. Enter a starting address, and then specify the number of addresses that can be assigned to the remote system.
19. Select **Allow remote system to access other networks (IP forwarding)**.
20. Go to the **Authentication** page to set the user name and password for this PPP profile.
21. In the Local system identification section, select **Allow the remote system to verify the identity of this system**. This opens the **Local System Identification** dialog box.
22. Under **Authentication protocol to use** select **Require encrypted password (CHAP-MD5)**.
23. Enter the user name, SystemB, and a password.
24. Click **OK** to save the PPP profile.

Activating packet rules

The VPN wizard automatically creates the packet rules that this connection requires to work properly. However, you must activate them on both systems before you can start the VPN connection.

To activate packet rules on System A, follow these steps:

1. In IBM Navigator for i, expand **Network > IP Policies**.
2. Click **Packet Rules** to open the **Packet Rules** panel and click **Actions > Activate Rules**. This action opens the **Activate Packet Rules** panel.

3. Select whether you want to activate only the VPN generated rules, only a selected file, or both the VPN generated rules and a selected file. You might choose the latter, for instance, if you have miscellaneous PERMIT and DENY rules that you want to enforce on the interface in addition to the VPN generated rules.
4. Select the interface on which you want the rules to be activated. In this case, select **Activate these rules on all interfaces and all point-to-point filter identifier**.
5. Click **OK** on the dialog box to confirm that you want to verify and activate the rules on the interface or interfaces you specified. After you click OK, the system checks the rules for syntax and semantic errors and reports the results in a message window at the bottom of the editor.
6. Repeat these steps to activate packet rules on System B.

Scenario: Firewall friendly VPN

In this scenario, a large insurance company wants to establish a VPN between a gateway in Chicago and a host in Minneapolis when both networks are behind a firewall.

Situation

Suppose you are a large home owner's insurance company based in Minneapolis and you just opened a new branch in Chicago. Your Chicago branch needs to access the customer database from the Minneapolis headquarters. You want to make sure the information being transferred is secure because the database contains confidential information about your customers such as, names, addresses, and phone numbers. You decide to connect both branches over the internet using a virtual private network (VPN). Both branches are behind a firewall and are using network address translation (NAT) to hide their unregistered private IP addresses behind a set of registered IP addresses. However, VPN connections have some well known incompatibilities with NAT. A VPN connection discards packets sent through a NAT device because NAT changes the IP address in the packet, thereby invalidating the packet. However, you can still use a VPN connection with NAT if you implement UDP encapsulation.

In this scenario, the private IP address from the Chicago network is put in a new IP header and gets translated when it goes through Firewall C (see following image). Then, when the packet reaches the Firewall D, it will translate the destination IP address to the IP address of System E, therefore the packet will be forwarded to System E. Finally, when the packet reaches System E it strips off the UDP header, leaving the original IPSec packet, which will now pass all validations and allow a secure VPN connection.

Objectives


In this scenario, a large insurance company wants to establish a VPN between a gateway in Chicago (Client) and a host in Minneapolis (Server) when both networks are behind a firewall.

The objectives of this scenario are as follows:

- The Chicago branch gateway always initiates the connection to the Minneapolis host.
- The VPN must protect all data traffic between the Chicago gateway and the Minneapolis host.
- Allow all users in the Chicago gateway to access a IBM i database located in the Minneapolis network over a VPN connection.

Details

The following figure illustrates the network characteristics for this scenario:



Chicago Network - Client

- Gateway B connects to the internet with IP address 214.72.189.35 and is the connection end point of the VPN tunnel. Gateway B performs IKE negotiations and applies UDP encapsulation to outgoing IP datagrams.
- Gateway B and PC A is in subnet 10.8.11.0 with mask 255.255.255.0
- PC A is the source and destination for data that flows through the VPN connection, therefore it is the data endpoint of the VPN tunnel.
- Only Gateway B can initiate the connection with System E.
- Firewall C has a Masq NAT rule with the public IP address of 129.42.105.17 that hides the IP address of Gateway B

Minneapolis Network - Server

- System E has an IP address of 56.172.1.1.
- System E is the responder for this scenario.
- Firewall D has an IP address of 146.210.18.51.
- Firewall D has a Static NAT rule that maps the public IP (146.210.18.15) to the private IP of System E (56.172.1.1). Therefore, from the clients perspective the IP address of System E is the public IP address (146.210.18.51) of Firewall D.

Configuration tasks

Related concepts:

"Key management" on page 9

A dynamic VPN provides additional security for your communications by using the Internet Key Exchange (IKE) protocol for key management. IKE allows the VPN servers on each end of the connection to negotiate new keys at specified intervals.

"NAT compatible IPSec with UDP" on page 14

UDP encapsulation allows IPSec traffic to pass through a conventional NAT device. Review this topic for more information about what it is and why you should use it for your VPN connections.

Completing the planning worksheets

The following planning checklists illustrate the type of information you need before you begin configuring the VPN. All answers on the prerequisite checklist must be YES before you proceed with VPN setup.

Note: There are separate worksheets for both Gateway B and System E.

Table 5. System requirements

Prerequisite checklist	Answers
Is the Digital Certificate Manager option installed?	Yes
Is HTTP server started (to support IBM Navigator for i)?	Yes
Is IBM TCP/IP Connectivity Utilities for i installed?	Yes
Did you set the retain server security data (QRETSVRSEC *SEC) system value to 1?	Yes
Is TCP/IP configured on your system (including IP interfaces, routes, local host name, and local domain name)?	Yes
Is normal TCP/IP communication established between the required endpoints?	Yes
Have you applied the latest program temporary fixes (PTFs)?	Yes
If the VPN tunnel traverses firewalls or routers that use IP packet filtering, do the firewall or router filter rules support AH and ESP protocols?	Yes
Are the firewalls or routers configured to permit traffic over port 4500 for key negotiations. Typically, VPN partners perform IKE negotiations over UDP port 500, when IKE detects NAT packets are sent over port 4500.	Yes
Are the firewalls configured to enable IP forwarding?	Yes

Table 6. Gateway B configuration

You need this information to configure the VPN for Gateway B	Answers
What type of connection are you creating?	gateway-to-another host
What will you name the dynamic-key group?	CHIgw2MINhost
What type of security and system performance do you require to protect your keys?	balanced
Are you using certificates to authenticate the connection? If no, what is the preshared key?	No : topsecretstuff
What is the identifier of the local key server?	IP address: 214.72.189.35
What is the identifier of the local data endpoint?	Subnet: 10.8.11.0 Mask: 255.255.255.0
What is the identifier of the remote key server?	IP address: 146.210.18.51
What is the identifier of the remote data endpoint?	IP address: 146.210.18.51
What ports and protocols do you want to allow to flow through the connection?	Any
What type of security and system performance do you require to protect your data?	balanced
To which interfaces does the connection apply?	TRLINE

Table 7. System E configuration

You need this information to configure the VPN for System E	Answers
What type of connection are you creating?	host-to-another gateway
What will you name the dynamic-key group?	CHIgw2MINhost
What type of security and system performance do you require to protect your keys?	highest
Are you using certificates to authenticate the connection? If no, what is the preshared key?	No : topsecretstuff
What is the identifier of the local key server?	IP address: 56.172.1.1

Table 7. System E configuration (continued)

You need this information to configure the VPN for System E	Answers
What is the identifier of the remote key server? Note: If the Firewall C IP address is unknown, you can use *ANYIP as the identifier for the remote key server.	IP address: 129.42.105.17
What is the identifier of the remote data endpoint?	Subnet: 10.8.11.0 Mask: 255.255.255.0
What ports and protocols do you want to allow to flow through the connection?	Any
What type of security and system performance do you require to protect your data?	highest
To which interfaces does the connection apply?	TRLINE

Configuring VPN on Gateway B

Complete the following steps to configure a VPN connection on Gateway B.

Use the information from your planning worksheets to configure VPN on Gateway B as follows:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
2. Click **Secure Connections** to open the **Connections** panel.
3. Click **Actions > New Connection** to start the Connection wizard.
4. Review the **Welcome** page for information about what objects the wizard creates.
5. Click **Next** to go to the **Connection Name** page.
6. In the **Name** field, enter CHIgw2MINhost.
7. Optional: Specify a description for this connection group.
8. Click **Next** to go to the **Connection Scenario** page.
9. Select **Connect your gateway to another host**.
10. Click **Next** to go to the **Internet Key Exchange Policy** page.
11. Select **Create a new policy** and then select **Balance security and performance**.

Note: If you get an error message stating "The certificate request could not be processed" you can ignore it because you are not using certificates for the key exchange.

12. Optional: If you have certificates installed you will see the **Certificate for Local Connection Endpoint** page. Select **No** to indicate that you will be using certificates to authenticate the connection.
13. Click **Next** to go to the **Local Key Server** page.
14. Select **IP version 4** as the **Identifier type** field.
15. Select 214.72.189.35 from the **IP address** field.
16. Click **Next** to go to the **Remote Key Server** page.
17. Select **IP version 4 address** in the **Identifier type** field.
18. Enter 146.210.18.51 in the **Identifier** field.

Note: Gateway B is initiating a connection to a Static NAT you must specify main mode key exchange in order to enter a single IP for the remote key. Main mode key exchange is selected by default when you create a connection with the VPN Connection Wizard. If aggressive mode is used in this situation, a non IPV4 type of remote identifier must be entered for remote key.

19. Enter topsecretstuff in the **Pre-shared key** field.
20. Click **Next** to go to the **Local Data Endpoint** page.
21. Select **IP version 4 subnet** from the **Identifier type** field.
22. Enter 10.8.0.0 in the **Identifier** field.
23. Enter 255.255.255.0 in the **Subnet mask** field.

24. Click **Next** to go to the **Data Services** page.
25. Accept the default values, and then click **Next** to go to the **Data Policy** page.
26. Select **Create a new policy** and then select **Balance security and performance**.
27. Click **Next** to go to the **Require Policy Filter** page.
28. Select **This connection requires a policy filter**.
29. Click **Next** to go to the **Applicable Interfaces** page.
30. Select **TRLINE** from the **Line** table.
31. Click **Next** to go to the **Summary** page.
32. Review the objects that the wizard will create to ensure they are correct.
33. Click **Finish** to complete the configuration.
34. When the **Activate Policy Filters** dialog box appears, select **Yes**, activate the generated policy filters then select **Permit all other traffic**.
35. Click **OK** to complete the configuration.

Configuring VPN on System E

Complete the following steps to configure a VPN connection on System E.

Use the information from your planning worksheets to configure VPN on System E as follows:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
 2. Click **Secure Connections** to open the **Connections** panel.
 3. Click **Actions > New Connection** to start the Connection wizard.
 4. Review the **Welcome** page for information about what objects the wizard creates.
 5. Click **Next** to go to the **Connection Name** page.
 6. In the **Name** field, enter **CHIgw2MINhost**.
 7. Optional: Specify a description for this connection group.
 8. Click **Next** to go to the **Connection Scenario** page.
 9. Select **Connect your host to another gateway**.
 10. Click **Next** to go to the **Internet Key Exchange Policy** page.
 11. Select **Create a new policy** and then select **Balanced security and performance**.
- Note:** If you get an error message stating "The certificate request could not be processed" you can ignore it because you are not using certificates for the key exchange.
12. Optional: If you have certificates installed you will see the **Certificate for Local Connection Endpoint** page. Select **No** to indicate that you will be using certificates to authenticate the connection.
 13. Click **Next** to go to the **Local Key Server** page.
 14. Select **IP version 4 address** as the **Identifier type** field.
 15. Select **56.172.1.1** from the **IP address** field.
 16. Click **Next** to go to the **Remote Key Server** page.
 17. Select **IP version 4 address** in the **Identifier type** field.
 18. Enter **129.42.105.17** in the **Identifier** field.

Note: If the Firewall C IP address is unknown, you can use ***ANYIP** as the identifier for the remote key server.

19. Enter **topsecretstuff** in the **Pre-shared key** field.
20. Click **Next** to go to the **Remote Data Endpoint** page.
21. Select **IP version 4 subnet** from the **Identifier type** field.
22. Enter **10.8.11.0** in the **Identifier** field.

23. Enter 255.255.255.0 in the **Subnet mask** field.
24. Click **Next** to go to the **Data Services** page.
25. Accept the default values, and then click **Next** to go to the Data Policy page.
26. Select **Create a new policy** and then select **Balanced security and performance**.
27. Click **Next** to go to the **Require Policy Filter** page.
28. Select **This connection requires a policy filter**.
29. Click **Next** to go to the **Applicable Interfaces** page.
30. Select **TRLINE** from the Line table.
31. Click **Next** to go to the **Summary** page.
32. Review the objects that the wizard will create to ensure they are correct.
33. Click **Finish** to complete the configuration.
34. When the **Activate Policy Filters** dialog box appears, select **Yes**, activate the generated policy filters then select **Permit all other traffic**.
35. Click **OK** to complete the configuration.

Starting Connection

After you have configured your VPN connection on System E you need to start your VPN connection.

Follow these steps to confirm that CHlgw2MINhost connection on System E is active:

1. Log in IBM Navigator for i for System E: `http://<systemE>:2001`.
2. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
3. Click **Secure Connections** to open the Connections wizard.
4. Click **Actions > All Connections**.
5. View **CHlgw2MINhost** and verify that the **Status** field is *Idle* or *On-Demand*.

Follow these steps to start the CHlgw2MINhost connection from Gateway B:

1. Log in IBM Navigator for i for Gateway B: `http://<Gateway B>:2001`.
2. In IBM Navigator for i, expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
3. If the VPN server is not started, click **Start VPN Server**.
4. Expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
5. Right-click **All Connections** to display a list of connections in the right pane.
6. Right-click **CHlgw2MINhost** and select **Start**.
7. From the **View** menu, select **Refresh**. If the connection starts successfully, the **Status** field will change from *Starting* or *On-Demand* to *Enabled*. The connection might take a short time to start, so periodically refresh until the status changes to *Enabled*.

Testing the connection

After you finish configuring both Gateway B and System E and you have successfully started the VPN servers, test the connectivity to ensure that both systems can communicate with each other.

To test your connections, follow these steps:

1. Find a system on the PC A network and open a Telnet session.
2. Specify the public IP address for System E, which is 146.210.18.51.
3. Specify any login information if it is needed. If you can view the login screen then the connection is working.

Scenario: VPN connection to remote users

The administrator needs to configure a virtual private network (VPN) connection to remote users to enable the remote connections.

The following tasks show you how the administrator configures a VPN connection to remote users.

Completing planning worksheets for VPN connection from the branch office to remote sales people

The administrator for the branch sales office uses the VPN planning worksheets to help them configure virtual private network (VPN) on their systems and remote workstations.

The VPN planning advisor is no longer supported through the Info Center. Each of the following planning worksheets can be used as a template to help plan and configure a VPN, using the VPN New Connection wizard in IBM Navigator for i.

Table 8. Planning worksheet for VPN connection between branch sales office and remote sales people

What the VPN wizard asks	Recommended values (from previous planning Advisor results)
What would you like to name this connection group?	SalestoRemote
What type of connection group would you like to create?	Select Connect your host to another host
What Internet Key Exchange policy do you want to use to protect your key?	Select Create a new policy and then select highest security, lowest performance
Are you using certificates?	Select No
Enter the identifier to represent the local key server for this connection.	Identifier type: IP version 4 address, IP address: 192.168.1.2 . For IPv6 address, identifier type: IP version 6 address, IP address: 2001:DB8::2 Note: IP addresses used in this scenario are meant for example purposes only. They do not reflect an IP addressing scheme and should not be used in any actual configuration. You should use your own IP addresses when completing these tasks.
What is the identifier of the key server that you want to connect to?	Identifier type: Any IP address, Pre-shared key: mycokey. Note: The pre-shared key is a 32-character text string that IBM i VPN uses to authenticate the connection as well as to establish the keys that protect your data. In general, you should treat a pre-shared key in the same way you treat a password.
What are the ports and protocols of the data that this connection will protect?	Local Port: 1701, Remote Port: Any port, Protocol: UDP
What data policy do you want to use to protect the data?	Select Create a new policy and then select highest security, lowest performance
Check the interfaces on the local system that this connection will be applied to.	ETHLINE (Branch sales office)

Configuring L2TP terminator profile for System A

If you want to configure the remote connections to remote workstations, you need to set up System A to accept inbound connections from these clients.

To configure a Layer Two Tunneling Protocol (L2TP) terminator profile for System A, complete the following steps:

1. In a web browser, enter `http://systemA:2001`, where `systemA` is the host name for System A.
 2. Log on to the system with your user profile and password.
 3. Expand **Network > Remote Access Services**.
 4. Right-click **Receiver Connection Profiles** to set the System A as a server that allows incoming connections from remote users, and select **New Profile**.
 5. Select the following options on the Setup page:
 - **Protocol type:** PPP
 - **Connection type:** L2TP (virtual line)
- Note:** The **Operating mode** field should automatically display **Terminator (network server)**.
- **Line service type:** Single line
 6. Click **OK**. This will launch the New Point-to-Point Profile Properties page.
 7. On the **General** tab, complete the following fields:
 - **Name:** MYCOL2TP
 - Select **Start profile with TCP** if you would like the profile to automatically start with TCP.
 8. On the **Connection** tab, select **192.168.1.2 (2001:DB8::2 in IPv6)** for the **Local tunnel endpoint IP address**.

Important: IP addresses used in this scenario are meant for example purposes only. They do not reflect an IP addressing scheme and should not be used in any actual configuration. Use your own IP addresses when completing these tasks.

9. Select **MYCOL2TP** as the **Virtual line name**. This will launch the New L2TP Properties page.
10. On the Authentication page, enter `systema` as the host name. Click **OK**. This will return you to the Connection page.
11. On the Connection page, select the following options and enter 25 as the **Maximum number of connections**.
 - a. Click the **Authentication** tab and select **Require this system to verify the identity of the remote system**.
 - b. Select **Authenticate locally with validation list**.
 - c. Enter `QL2TP` in **Validation list name** field, and click **New**.
12. On the Validation list page, select **Add**.
13. Add user names and passwords for each of your remote employees. Click **OK**.
14. On the Password confirmation page, re-enter the password for each remote employees. Click **OK**.
15. On the TCP/IP IPv4 Settings page, select `10.1.1.1` for **Local IP address**.
16. In the **IP address assignment method** field, select **Address pool**.
17. In the **Starting IP address** field, enter `10.1.1.100` and 49 for the **Number of addresses**. For IPv6 address, check **Enable IPv6** field from the TCP/IP IPv6 settings tab.
18. Select **Allow remote system to access other networks (IP forwarding)**. Click **OK**.

Starting receiver connection profile

After configuring Layer Two Tunneling Protocol (L2TP) receiver connection profile for System A, the administrator needs to start this connection so that it will listen for incoming requests from remote clients.

Note: You might receive an error message that the QUSRWRK subsystem is not started. This message occurs when attempting to start the receiver connection profile. To start the QUSRWRK subsystem, complete these steps:

1. In a character-based interface, enter strsubs.
2. On the Start Subsystem display, enter QUSRWRK in the **Subsystem description** field.

To start the receiver connection profile for remote clients, complete these tasks:

- | 1. Log in IBM Navigator for i from *System A* and expand **Network > Remote Access Services**.
- | 2. Click **Receiver Connection Profiles** to open the **Receiver Connection Profiles** panel.
- | 3. Right-click MYCOL2TP and select **Start**.
- | 4. The **Status** field is changed to **Waiting for connection requests**.

Configuring a VPN connection on System A for remote clients

After configuring and starting the Layer Two Tunneling Protocol (L2TP) receiver connection profile for System A, the administrator needs to configure a virtual private network (VPN) to protect the connection between remote clients and the network in the branch sales office.

To configure a VPN for remote clients, complete these steps:

Important: IP addresses used in this scenario are meant for example purposes only. They do not reflect an IP addressing scheme and should not be used in any actual configuration. Use your own IP addresses when completing these tasks.

- | 1. In a web browser, enter http://systemA:2001, where systemA is the host name of system A.
- | 2. Log on to the system with your user profile and password.
- | 3. Expand **Network > IP Policies > All Tasks > IP Policies > Virtual Private Networking**.
- | 4. Click **Create Connection** to start the VPN New Connection wizard. Review the Welcome page for information about what objects the wizard creates.
- | 5. Click **Next** to go to the Connection Name page.
- | 6. In the **Name** field, enter SalestoRemote.
- | 7. Optional: Specify a description for this connection group. Click **Next**.
- | 8. On the Connection Scenario page, select **Connect your host to another host**. Click **Next**.
- | 9. On the Internet Key Exchange Policy page, select **Create a new policy**, and then select **Highest security, lowest performance**. Click **Next**.
- | 10. On the Certificate for Local Connection Endpoint page, select **No**. Click **Next**.
- | 11. On the Local Key Server page, select **Version 4 IP address** as the identifier type. The associated IP address should be 192.168.1.2. Click **Next**. For IPv6 address, on the Local Key Server page, select **Version 6 IP address** as the identifier type. The associated IP address should be 2001:DB8::2. Click **Next**.
- | 12. On the Remote Key Server page, select **Any IP address** in the **Identifier type** field. In the **Pre-shared key** field, enter mycokey. Click **Next**.
- | 13. On the Data Services page, enter 1701 for the local port. Then select 1701 for the remote port and select **UDP** for the protocol. Click **Next**.
- | 14. On the Data Policy page, select **Create a new policy** and then select **Highest security, lowest performance**. Click **Next**.
- | 15. On the Applicable Interfaces page, select **ETHLINE**. Click **Next**.
- | 16. On the Summary page, review the objects that the wizard will create to ensure they are correct.
- | 17. Click **Finish** to complete the configuration. When the Activate Policy Filters window opens, select **No, packet rules will be activated at a later time**. Click **OK**.

Activating filter rules

The wizard automatically creates the packet rules that this connection requires to work properly. However, you must activate them on both systems before you can start the virtual private network (VPN) connection.

To activate filter rules on System A, follow these steps:

Important: IP addresses used in this scenario are meant for example purposes only. They do not reflect an IP addressing scheme and should not be used in any actual configuration. You should use your own IP addresses when completing these tasks.

- | 1. In a web browser, enter `http://systemA:2001`, where `systemA` is the host name of system A.
- | 2. Log on to the system with your user profile and password.
- | 3. Expand **Network > IP Policies**.
- | 4. Click **Packet Rules** and select **Activate Rules** from Packet Rules table actions.
- | 5. On the Activate Packet Rules page, select **activate only the VPN generated rules** and select **ETHLINE** as the interface on which you would like to activate these filter rules. Click **OK**.

Before remote users can configure their Windows workstations, the administrator gives them the following information so they can set up their side of the connection. For each of your remote users, give them the following information:

- Name of Pre-shared key: `mycokey`
- IP address of System A: `192.168.1.2 (2001:DB8::2 in IPv6)`
- User name and password for the connection

Note: These were created when the administrator added the user name and passwords to a validation list during the configuration of Layer Two Tunneling Protocol (L2TP) terminator profile.

Configuring VPN on Windows client

Use this procedure to configure VPN on Windows client.

Remote users at MyCo, Inc need to set up their remote Windows client by completing the following steps:

1. In the Windows **Start** menu, expand **All Programs > Accessories > Communications > New Connection Wizard**.
2. On the Welcome page, read the overview information. Click **Next**.
3. On the Network Connection Type page, select **Connect to the network at my workplace**. Click **Next**.
4. On the Network Connection page, select **Virtual Private Network connection**. Click **Next**.
5. On the Connection Name page, enter `Connection to Branch office` in the **Company Name** field. Click **Next**.
6. On the Public Network page, select **Do not dial the initial connection**. Click **Next**.
7. On the VPN Server Selection page, enter `192.168.1.2 (2001:DB8::2 in IPv6)` in the **Host name or IP address** field. Click **Next**.
8. On the Connection Availability page, select **My Use Only**. Click **Next**.
9. On the Summary page, click **Add a shortcut to this connection to my desktop**. Click **Finish**.
10. Click the **Connect Connection to MyCo** icon that has been created on your desktop.
11. On the Connect Connection to MyCo page, enter the user name and password that the administrator provided.
12. Select **Save this user name and password for the following users and Me only**. Click **Properties**.
13. On the **Security** page, ensure that the following **Security options** are selected:
 - **Typical**

- **Require secured password**
- **Require data encryption**

Click **IPSec Settings**.

14. On the IPSec Settings page, select **Use pre-shared key for authentication** and enter mycokey in the **Pre-shared key** field. Click **OK**.
15. On Networking page, select **L2TP IPSec VPN** as the **Type of VPN**. Click **OK**.
16. Sign on with user name and password and click **Connect**.

To start the virtual private network (VPN) connection on the client side, click the icon that appears on your desktop after completing the connection wizard.

Testing VPN connection between endpoints

After you finish configuring the connection between System A and remote users and have successfully started the connection, you should test the connectivity to ensure that the remote hosts can communicate with each other.

To test the connectivity, follow these steps:

1. In a web browser, enter `http://systemA:2001`, where systemA is the host name of system A.
 2. Log on to the system with your user profile and password.
 3. Expand **Network > All Tasks > TCP/IP Configuration**.
 4. Click **Work with Ping**.
 5. From the **Ping** panel, enter 10.1.1.101 (2001:DA8::1:101 in IPv6) in the **Ping** field.
- Note:** 10.1.1.101 represents the IP address dynamically assigned (to the remote sales client) from the address pool specified in the Layer Two Tunneling Protocol (L2TP) terminator profile on System A.
6. Click **Ping Now** to verify connectivity from System A to a remote workstation. Click **OK**.

To test the connection from the remote client, the remote employee completes these steps on a workstation that is running Windows:

1. From the command prompt, enter `ping 10.1.1.2` (ping 2001:DA8::2 in IPv6). This is the IP address of one of the workstations in the network of the corporate office.
2. Repeat these steps to test the connectivity from the corporate office to the branch office.

Scenario: Using network address translation for VPN

In this scenario, your company wants to exchange sensitive data with one of its business partners by using VPN. To further protect the privacy of your company's network structure, your company will also use VPN NAT to hide the private IP address of the system it uses to host the applications to which your business partner has access.

Situation

Suppose you are the network administrator for a small manufacturing company in Minneapolis. One of your business partners, a parts supplier in Chicago, wants to start doing more of their business with your company over the Internet. It is critical that your company have the specific parts and quantities at the exact time it needs them, so the supplier needs to be aware of your company's inventory status and production schedules. Currently you handle this interaction manually, but you find it time consuming, expensive and even inaccurate at times, so you are more than willing to investigate your options.

Given the confidentiality and time-sensitive nature of the information you exchange, you decide to create a VPN between your supplier's network and your company's network. To further protect the privacy of your company's network structure, you decide you will need to hide the private IP address of the system that hosts the applications to which the supplier has access.

You can use VPN's to not only create the connection definitions on the VPN gateway in your company's network, but also to provide the address translation you need to hide your local private addresses. Unlike conventional network address translation (NAT), which changes the IP addresses in the security associations (SAs) that VPN requires to function, VPN NAT performs address translation before the SA validation by assigning an address to the connection when the connection starts.


Objectives

The objectives of this scenario are to:

- allow all clients in the supplier network to access a single host system in the manufacturer's network over a gateway-to-gateway VPN connection.
- hide the private IP address of the host system in the manufacturer's network, by translating it to a public IP address by using network address translation for VPN (VPN NAT).

Details

The following diagram illustrates the network characteristics of both the supplier network and the manufacturing network:



- VPN gateway-A is configured to always initiate connections to VPN gateway-B.
- VPN gateway-A defines the destination endpoint for the connection as 204.146.18.252 (the public address assigned to System C).
- System C has a private IP address in the manufacturer's network of 10.6.100.1.
- A public address of 204.146.18.252 has been defined in the local service pool on VPN gateway-B for System C's private address, 10.6.100.1.
- VPN gateway-B translates System C's public address to its private address, 10.6.100.1, for inbound datagrams. VPN gateway-B translates returning, outbound, datagrams from 10.6.100.1 back to System C's public address, 204.146.18.252. As far as clients in the supplier network are concerned, System C has an IP address of 204.146.18.252. They will never be aware that address translation has occurred.

Configuration tasks

You must complete each of the following tasks to configure the connection described in this scenario:

1. Configure a basic gateway-to-gateway VPN between **VPN gateway-A** and **VPN gateway-B**.
2. Define a local service pool on **VPN gateway-B** to hide **System C**'s private address behind the public identifier, 204.146.18.252.
3. Configure **VPN gateway-B** to translate local addresses using local service pool addresses.

Related concepts:

“Network address translation for VPN” on page 13

VPN provides a means for performing network address translation, called VPN NAT. VPN NAT differs from traditional NAT in that it translates addresses before applying the IKE and IPSec protocols. Refer to this topic to learn more.

Planning for VPN

The first step to successfully using VPN is planning. This topic provides information about migrating from prior releases, setup requirements, and a planning worksheet that is customized to your specifications.

Planning is an essential part of your total VPN solution. There are many complex decisions you need to make to ensure that your connection works properly. Use these resources to gather all the information that you need to ensure that your VPN is a success:

- VPN setup requirements
- Determine what type of VPN to create
- Complete the VPN planning worksheets

After you come up with a plan for your VPN, you can begin configuring it.

Related tasks:

“Configuring VPN” on page 49

The VPN interface provides you with several different ways to configure your VPN connections. You can configure a manual or dynamic connection.

VPN setup requirements

For a VPN connection to function properly on your systems and with network clients, you must meet the minimum requirements

The following lists the minimum requirements to setup a VPN connection:

System requirements

- Digital Certificate Manager
- IBM Navigator for i
- Set the retain server security data (QRETSVRSEC *SEC) system value to 1
- TCP/IP must be configured, including IP interfaces, routes, local host name, and local domain name
- The IBM i 7.1 features, IKEv2 and Enhanced Cryptographic Algorithms, are only available through IBM Navigator for i

Related tasks:

“Getting started with VPN troubleshooting” on page 63

Complete this task to learn the various methods for determine any VPN problems you are having on your system.

Determining what type of VPN to create

Determining how you will use your VPN is one of the first steps in successful planning. To do this, you need to understand the role that both the local key server and the remote key server play in the connection.

For example, are the *connection* endpoints different from the *data* endpoints? Are they the same or some combination of both? Connection endpoints authenticate and encrypt (or decrypt) data traffic for the connection, and optionally provide key management with the Internet Key Exchange (IKE) protocol. Data endpoints, however, define the connection between two systems for IP traffic that flows across the VPN; for example, all TCP/IP traffic between 123.4.5.6 and 123.7.8.9. Typically, when the connection and data endpoints are different, the VPN server is a gateway. When they are the same, the VPN server is a host.

Various types of VPN implementations that are well suited to most business needs follow:

Gateway-to-gateway

The connection endpoints of both systems are different from the data endpoints. The IP Security (IPSec) protocol protects traffic as it travels between the gateways. However, IPSec does not protect data traffic on either side of the gateways within the internal networks. This is a common setup for connections between branch offices because traffic that is routed beyond the branch office gateways, into the internal networks, is often considered trusted.

Gateway-to-host

IPSec protects data traffic as it travels between your gateway and a host in a remote network. VPN does not protect data traffic in the local network because you consider it trusted.

Host-to-gateway

VPN protects data traffic as it travels between a host in the local network and a remote gateway. VPN does not protect data traffic in the remote network.

Host-to-host

The connection endpoints are the same as the data endpoints on both the local and the remote systems. VPN protects data traffic as it travels between a host in the local network and a host in the remote network. This type of VPN provides end-to-end IPSec protection.

Completing VPN planning work sheets

Use the VPN planning work sheets to gather detailed information about your VPN usage plans. You need to complete these work sheets to adequately plan your VPN strategy. You can also use this information to configure your VPN.

If you prefer, you can print and complete the planning work sheets to gather detailed information about your VPN usage plans.

Choose the work sheet for the type of connection you want to create.

- Planning work sheet for dynamic connections
- Planning work sheet for manual connections

If you will create multiple connections with similar properties, you might want to set the VPN defaults. The default values you configure seed the VPN property sheets. This means that you are not required to configure the same properties multiple times. To set the VPN defaults, select **Edit** from the VPN main menu, and then select **Defaults**.

Planning work sheet for dynamic connections

Complete this work sheet before you configure a dynamic connection.

Before you create your dynamic VPN connections, complete this work sheet. The work sheet assumes you will use the New Connection Wizard. The wizard allows you to set up a VPN based on your basic security requirements. In some cases, you might need to refine the properties that the wizard configures

for a connection. For example, you might decide that you require journaling or that you want the VPN server to start each time TCP/IP starts. If this is the case, right-click the dynamic-key group or connection that the wizard created and select **Properties**.

Answer each question before you proceed with your VPN setup.

Table 9. System requirements

Prerequisite checklist	Answers
Is the Digital Certificate Manager option installed?	Yes
Is the HTTP server started (to support IBM Navigator for i)?	Yes
Is IBM TCP/IP Connectivity Utilities for i installed?	Yes
Did you set the retain server security data (QRETSVRSEC *SEC) system value to 1?	Yes
Is TCP/IP configured on your system (including IP interfaces, routes, local host name, and local domain name)?	Yes
Is normal TCP/IP communication established between the required endpoints?	Yes
Have you applied the latest program temporary fixes (PTFs)?	Yes
If the VPN tunnel traverses firewalls or routers that use IP packet filtering, do the firewall or router filter rules support AH and ESP protocols?	Yes
Are the firewalls or routers configured to permit IKE (UDP port 500), AH, and ESP protocols?	Yes
Are the firewalls configured to enable IP forwarding?	Yes

Table 10. VPN configuration

You need this information to configure a dynamic VPN connection	Answers
What type of connection are you creating? • Gateway-to-gateway • Host-to-gateway • Gateway-to-host • Host-to-host	
What will you name the dynamic-key group?	
What type of security and system performance do you require to protect your keys? • Highest security, lowest performance • Balance security and performance • Lowest security and highest performance	
Are you using certificates to authenticate the connection? If no, what is the preshared key?	
What is the identifier of the local key server?	
What is the identifier of the local key server?	
What is the identifier of the remote key server?	
What is the identifier of the remote data endpoint?	
What type of security and system performance do you require to protect your data? • Highest security, lowest performance • Balance security and performance • Lowest security and highest performance	

Planning work sheet for manual connections

Complete this work sheet before you configure a manual connection.

Complete this work sheet to assist you in creating your virtual private network (VPN) connections that do not use IKE for key management. Answer each of these questions before you proceed with your VPN setup:

Table 11. System Requirements

Prerequisite checklist	Answers
Is the Digital Certificate Manager installed?	
Is the HTTP server started (to support IBM Navigator for i)?	
Is IBM TCP/IP Connectivity Utilities for i installed?	
Did you set the retain server security data (QRETSVRSEC *SEC) system value to 1?	
Is TCP/IP configured on your system (including IP interfaces, routes, local host name, and local domain name)?	
Is normal TCP/IP communication established between the required endpoints?	
Have you applied the latest program temporary fixes (PTFs)?	
If the VPN tunnel traverses firewalls or routers that use IP packet filtering, do the firewall or router filter rules support AH and ESP protocols?	
Are the firewalls or routers configured to permit the AH and ESP protocols?	
Are the firewalls configured to enable IP forwarding?	

Table 12. VPN configuration

You need this information to configure a manual VPN	Answers
What type of connection are you creating? • Host-to-host • Host-to-gateway • Gateway-to-host • Gateway-to-gateway	
What will you name the connection?	
What is the identifier of the local connection endpoint?	
What is the identifier of the remote connection endpoint?	
What is the identifier of the local data endpoint?	
What is the identifier of the remote data endpoint?	
What type of traffic will you allow for this connection (local port, remote port, and protocol)?	
Do you require address translation for this connection? See Network address translation for VPN for more information.	
Will you use tunnel mode or transport mode?	
Which IPSec protocol will the connection use (AH, ESP, or AH with ESP)? See IP Security (IPSec) for more information.	
Which authentication algorithm will the connection use (HMAC-MD5 or HMAC-SHA)?	
Which encryption algorithm will the connection use (DES-CBC or 3DES-CBC)? Note: You specify an encryption algorithm only if you selected ESP as your IPSec protocol.	
What is the AH inbound key? If you use MD5, the key is a 16-byte hexadecimal string. If you use SHA, the key is a 20-byte hexadecimal string.	
Your inbound key must match the outbound key of the remote server exactly.	

Table 12. VPN configuration (continued)

What is the AH outbound key? If you will use MD5, the key is a 16-byte hexadecimal string. If you will use SHA, the key is a 20-byte hexadecimal string.	
Your outbound key must match the inbound key of the remote server exactly.	
What is the ESP inbound key? If you use DES, the key is an 8-byte hexadecimal string. If you will use 3DES, the key is a 24-byte hexadecimal string.	
Your inbound key must match the outbound key of the remote server exactly.	
What is the ESP outbound key? If you use DES, the key is an 8-byte hexadecimal string. If you will use 3DES, the key is a 24-byte hexadecimal string.	
Your outbound key must match the inbound key of the remote server exactly.	
What is the inbound Security Policy Index (SPI)? The inbound SPI is a 4-byte hexadecimal string, where the first byte is set to 00.	
Your inbound SPI must match the outbound SPI of the remote server exactly.	
What is the outbound SPI? The outbound SPI is a 4-byte hexadecimal string.	
Your outbound SPI must match the inbound SPI of the remote server exactly.	

Related concepts:

“Network address translation for VPN” on page 13

VPN provides a means for performing network address translation, called VPN NAT. VPN NAT differs from traditional NAT in that it translates addresses before applying the IKE and IPSec protocols. Refer to this topic to learn more.

Configuring VPN

The VPN interface provides you with several different ways to configure your VPN connections. You can configure a manual or dynamic connection.

A dynamic connection is one that dynamically generates and negotiates the keys that secure your connection, while it is active, by using the Internet Key Exchange (IKE) protocol. Dynamic connections provide an extra level of security for the data that flows across it because the keys change, automatically, at regular intervals. Consequently, an attacker is less likely to capture a key, have time to break it, and use it to divert or capture the traffic the key protects.

A manual connection, however, does not provide support for IKE negotiations, and consequently, automatic key management. Further, both ends of the connection require you to configure several attributes that must match exactly. Manual connections use static keys that do not refresh or change while the connection is active. You must stop a manual connection to change its associated key. If you consider this a security risk, you might want to create a dynamic connection instead.

Related concepts:

“Planning for VPN” on page 45

The first step to successfully using VPN is planning. This topic provides information about migrating from prior releases, setup requirements, and a planning worksheet that is customized to your specifications.

Configuring VPN connections with the New Connection wizard

The New Connection wizard allows you to create a virtual private network (VPN) between any combination of hosts and gateways.

For example, host-to-host, gateway-to-host, host-to-gateway, or gateway-to-gateway.

The wizard automatically creates each of the configuration objects VPN requires to work properly, including the packet rules. However, if you need to add function to your VPN; for example, journaling or network address translation for VPN (VPN NAT), you might want further refine your VPN through the property sheets of the appropriate dynamic-key group or connection. To do this, you must first stop the connection if it is active. Then, right-click the dynamic-key group or connection, and select **Properties**.

To create a VPN with the Connection wizard, follow these steps:

1. In IBM Navigator for i, expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
2. Click **Create Connection** to start the wizard.
3. Complete the wizard to create a basic VPN connection. Click **Help** if you need assistance.

Configuring VPN security policies

After you determine how you will use your VPN you must define your VPN security policies.

Note: After you configure your VPN security policies, you must then configure the secure connections.

Related tasks:

“Configuring a secure VPN connection” on page 51

After you have configured the security policies for your connection, you must then configure the secure connection.

Configuring an Internet Key Exchange policy

The Internet Key Exchange (IKE) policy defines what level of authentication and encryption protection IKE uses during phase 1 negotiations.

IKE phase 1 establishes the keys that protect the messages that flow in the subsequent phase 2 negotiations. You do not need to define an IKE policy when you create a manual connection. In addition, if you create your VPN with the New Connection wizard, the wizard can create your IKE policy for you.

- | VPN uses RSA signature mode, ECDSA, or preshared keys to authenticate phase 1 negotiations. If you plan to use digital certificates for authenticating the key servers, you must first configure them by using the Digital Certificate Manager. The IKE policy also identifies which remote key server will use this policy.

To define an IKE policy or make changes to an existing one, follow these steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**, click **IP Security Policies**.
- | 2. To create a new policy, right-click **Internet Key Exchange Policies** and select **New Internet Key Exchange Policy**. To make changes to an existing policy, right-click **Internet Key Exchange Policies** and select **Open**. Right-click the policy you want to change and select **Properties**.
- | 3. Complete each of the property sheets. Click **Help** if you have questions about how complete a page or any of its fields.
- | 4. Click **OK** to save your changes.

It is recommended that you use main mode negotiation whenever a preshared key is used for authentication. They provide a more secure exchange. If you must use preshared keys and aggressive mode negotiation, select obscure passwords that are unlikely to be cracked in attacks that scan the dictionary. It is also recommended you periodically change your passwords. To force a key exchange to use main mode negotiation, perform the following tasks:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**, click **IP Security Policies**.
- | 2. Right-click **Internet Key Exchange Policies** and select **Open**.
- | 3. Right-click a particular key exchange policy and select **Properties**.

4. On the Transforms page, click **Responding Policy**. The Responding Internet Key Exchange Policy panel appears.
5. In the Identity protection field, deselect **IKE aggressive mode negotiation (no identity protection)**.
6. Click **OK** to return to the Properties dialog.
7. Click **OK** again to save your changes.

Note: When you set the identity protection field, the change is effective for all exchanges with remote key servers, because there is only one responding IKE policy for the entire system. Main mode negotiation ensures that the initiating system can only request a main mode key policy exchange.

Related concepts:

“Key management” on page 9

A dynamic VPN provides additional security for your communications by using the Internet Key Exchange (IKE) protocol for key management. IKE allows the VPN servers on each end of the connection to negotiate new keys at specified intervals.

Related tasks:

Digital Certificate Manager

Configuring a data policy

A data policy defines what level of authentication or encryption protects data as it flows through the VPN.

The communicating systems agree on these attributes during the Internet Key Exchange (IKE) protocol phase 2 negotiations. You do not need to define a data policy when you create a manual connection. In addition, if you create your VPN with the New Connection wizard, the wizard can create a data policy for you.

To define a data policy or make changes to an existing one, follow these steps:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and then click **IP Security Policies**.
2. To create a new data policy, right-click **Data Policies** and select **New Data Policy**. To make changes to an existing data policy, right-click **Data Policies** and select **Open**. Right-click on the data policy you want to change and select **Properties**.
3. Complete each of the property sheets. Click **Help** if you have questions about how complete a page or any of its fields.
4. Click **OK** to save your changes.

Related concepts:

“Key management” on page 9

A dynamic VPN provides additional security for your communications by using the Internet Key Exchange (IKE) protocol for key management. IKE allows the VPN servers on each end of the connection to negotiate new keys at specified intervals.

Configuring a secure VPN connection

After you have configured the security policies for your connection, you must then configure the secure connection.

For dynamic connections, the secure connection object includes a dynamic-key group and a dynamic-key connection.

The **dynamic-key group** defines the common characteristics of one or more VPN connections. Configuring a dynamic-key group allows you to use the same policies, but different data endpoints for each connection within the group. Dynamic-key groups also allow you to successfully negotiate with remote initiators when the data endpoints proposed by the remote system are not specifically known

ahead of time. It does this by associating the policy information in the dynamic-key group with a policy filter rule with an IPSEC action type. If the specific data endpoints offered by the remote initiator fall within the range specified in the IPSEC filter rule, they can be subjected to the policy defined in the dynamic-key group.

The **dynamic-key connection** defines the characteristics of individual data connections between pairs of endpoints. The dynamic-key connection exists within the dynamic-key group. After you configure a dynamic-key group to describe what policies connections in the group use, you need to create individual dynamic-key connections for connections that you initiate locally.

To configure the secure connection object, complete both the Part 1 and Part 2 tasks:

Related concepts:

“Configuring VPN security policies” on page 50

After you determine how you will use your VPN you must define your VPN security policies.

“Configuring VPN packet rules” on page 53

If you are creating a connection for the first time, allow VPN to automatically generate the VPN packet rules for you. You can do this by either using the New Connection wizard or the VPN properties pages to configure your connection.

Related tasks:

“Activating VPN packet rules” on page 58

You must activate the VPN packet rules before you can start your VPN connections.

Part 1: Configure a dynamic-key group

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and then click **Secure Connections**.
2. Right-click **By Group** and select **New Dynamic-Key Group**.
3. Click **Help** if you have questions about how complete a page or any of its fields.
4. Click **OK** to save your changes.

Part 2: Configure a dynamic-key connection

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and then click **Secure Connections**.
2. Right-click **By Group** and select **Open**.
3. Right-click the dynamic-key group you created in part one and select **New Dynamic-Key Connection**.
4. Click **Help** if you have questions about how complete a page or any of its fields.
5. Click **OK** to save your changes.

After you complete these steps, you need to activate the packet rules that the connection requires to work properly.

Note: In most cases, allow the VPN interface to generate your VPN packet rules automatically by selecting the **Generate the following policy filter for this group** option on the **Dynamic-Key Group - Connections** page. However, if you select the **The policy filter rule will be defined in Packet Rules** option, you must then configure VPN packet rules by using the Packet Rules editor and then activate them.

Configuring a manual connection

A manual connection is one where you must configure all of your VPN properties without using any wizards.

Further, both ends of the connection require you to configure several elements that must match *exactly*. For example, your inbound keys must match the remote system's outbound keys or the connection will fail.

Manual connections use static keys that are not refreshed or changed while the connection is active. You must stop a manual connection in order to change its associated key. If you consider this a security risk, and both ends of the connection support the Internet Key Exchange (IKE) protocol, you might want to consider setting up a dynamic connection instead.

To define the properties for your manual connection, follow these steps:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
2. Right-click **All Connections** and select **New Manual Connection**.
3. Complete each of the property sheets. Click **Help** if you have questions about how complete a page or any of its fields.
4. Click **OK** to save your changes.

Note: In most cases, allow the VPN interface to generate your VPN packet rules automatically by selecting the **Generate a policy filter that matches the data endpoints** option on the **Manual Connection - Connection** page. However, if you select the **The policy filter rule will be defined in Packet Rules** option, you must then configure a policy filter rule by hand and then activate them.

Related tasks:

“Configuring a policy filter rule” on page 56

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

Configuring a dynamic connection

A dynamic connection dynamically generates and negotiates the keys that secure your connection, while it is active, by using the Internet Key Exchange (IKE) protocol.

Complete the New Dynamic-Key Connection wizard to configure a dynamic connection by following these steps:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and then click **Secure Connections**.
2. Right-click **By Group** and select **Open**.
3. Right-click on the specific dynamic-key group and select **New Dynamic-Key Connection**.
4. Complete each of the property sheets. Click **Help** if you have questions about how complete a page or any of its fields.
5. Click **OK** to save your changes.

Configuring VPN packet rules

If you are creating a connection for the first time, allow VPN to automatically generate the VPN packet rules for you. You can do this by either using the New Connection wizard or the VPN properties pages to configure your connection.

If you decide to create your VPN packet rules by using the Packet Rules editor in IBM Navigator for i, create any additional rules this way as well. Conversely, if you have VPN generate your policy filter rules, create all additional policy filter rules this way.

In general, VPNs require two types of filter rules: Pre-IPSec filter rules and policy filter rules. Review the topics below to learn how to configure these rules by using the Packet Rules editor in IBM Navigator for i. If you want to read about other VPN and filtering options, see the VPN and IP filtering section of the VPN concepts topic.

- **Configuring the pre-IPSec filter rule**

The pre-IPSec rules are any rules on your system that come before rules with an IPSEC action type. This topic only discusses the pre-IPSec rules that VPN requires to work properly. In this case, the pre-IPSec rules are a pair of rules that allow IKE processing over the connection. IKE allows dynamic key generation and negotiations to occur for your connection. You might need to add other pre-IPSec rules depending on your particular network environment and security policy.

Note: You only need to configure this type of pre-IPSec rule if you already have other rules that permit IKE for specific systems. If there are no filter rules on the system specifically written to permit IKE traffic, then IKE traffic is implicitly allowed.

- **Configure a policy filter rule**

The policy filter rule defines the traffic that can use the VPN and what data protection policy to apply to that traffic.

Things to consider before you begin

When you add filter rules to an interface, the system automatically adds a default DENY rule for that interface. This means that any traffic not explicitly permitted is denied. You cannot see or change this rule. As a result, you might find that traffic that previously worked mysteriously fails after you activate your VPN filter rules. If you want to allow traffic other than VPN on the interface, you must add explicit PERMIT rules to do so.

After you configure the appropriate filter rules, you must define the interface to which they apply, and then activate them.

It is essential that you configure your filter rules properly. If you do not, the filter rules can block all IP traffic coming into and going out of your system. This includes your connection to IBM Navigator for i, which you use to configure the filter rules.

If the filter rules do not permit IBM i traffic, IBM Navigator for i cannot communicate with your system. If you find yourself in this situation, you must log on to your system using an interface that still has connectivity, such as the operations console. Use the RMVTCPTBL command to remove all filters on this system. This command also ends the *VPN servers and then restarts them. Then, configure your filters and reactivate them.

Related concepts:

“VPN and IP filtering” on page 16

IP filtering and VPN are closely related. In fact, most VPN connections require filter rules to work properly. This topic provides you information about what filters VPN requires, as well as other filtering concepts related to VPN.

Related tasks:

“Configuring a secure VPN connection” on page 51

After you have configured the security policies for your connection, you must then configure the secure connection.

“Configuring the pre-IPSec filter rule” on page 55

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

“Configuring a policy filter rule” on page 56

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

“Defining an interface for the VPN filter rules” on page 57

After you configure your VPN packet rules and any other rules that you need to enable your VPN connection, you must define the interface to which they apply.

“Activating VPN packet rules” on page 58

You must activate the VPN packet rules before you can start your VPN connections.

Configuring the pre-IPSec filter rule

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

A pair of Internet Key Exchange (IKE) servers dynamically negotiate and refresh keys. IKE uses the well-known port, 500. For IKE to work properly, you need to allow UDP datagrams over port 500 for this IP traffic. To do this, you will create a pair of filter rules; one for inbound traffic and one for outbound traffic, so that your connection can dynamically negotiate keys to protect the connection:

1. In IBM Navigator for i, expand **Network > IP Policies** and click **Packet Rules**.
2. On the **Packet Rules** panel, select **Actions > Rules Editor**. This opens the Packet Rules editor, which allows you to create or edit filter and NAT rules for your system.
3. On the Welcome panel, select **Create a new packet rules file** and click **OK**.
4. If the **Getting started** panel shows up, read the user guide on the **Getting started** panel and click **OK**.
5. From the Packet Rules editor select **Insert > Filter**.
6. On the **General** page, specify a set name for your VPN filter rules. It is recommended that you create at least three different sets: one for your pre-IPSec filter rules, one for your policy filter rules, and one for miscellaneous PERMIT and DENY filter rules. Name the set that contains your pre-IPSec filter rules with a prefix of *preipsec*. For example, *preipsecfilters*.
7. In the **Action** field, select **PERMIT** from the drop-down list.
8. In the **Direction** field, select **OUTBOUND** from the drop-down list.
9. In the **Source address name** field, select = from the first drop-down list and then enter the IP address of the local key server in the second field. You specified the IP address of the local key server in the IKE policy.
10. In the **Destination address name** field, select = from the first drop-down list and then enter the IP address of the remote key server in the second field. You also specified the IP address of the remote key server in the IKE policy.
11. On the **Services** page, select **Service**. This enables the **Protocol**, **Source port**, and **Destination port** fields.
12. In the **Protocol** field, select **UDP** from the drop-down list.
13. For **Source port**, select = in the first field, then enter 500 in the second field.
14. Repeat the previous step for **Destination port**.
15. Click **OK**.
16. Repeat these steps to configure the INBOUND filter. Use the same set name and reverse addresses as necessary.

Note: A less secure, but easier option for permitting IKE traffic through the connection, is to configure only one pre-IPSec filter, and use wildcard values (*) in the **Direction**, **Source address name**, and **Destination address name** fields.

The next step is to configure a policy filter rule to define what IP traffic the VPN connection protects.

Related concepts:

“Configuring VPN packet rules” on page 53

If you are creating a connection for the first time, allow VPN to automatically generate the VPN packet rules for you. You can do this by either using the New Connection wizard or the VPN properties pages

to configure your connection.

Related tasks:

“Configuring a policy filter rule”

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

Configuring a policy filter rule

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

The policy filter rule (a rule where action=IPSEC) defines which addresses, protocols, and ports can use the VPN. It also identifies the policy that will be applied to traffic in the VPN connection. To configure a policy filter rule, follow these steps:

Note: If you just configured the pre-IPSec rule (for dynamic connections, only) the Packet Rules editor will still be open; go to step 4.

1. In IBM Navigator for i, expand **Network > IP Policies** and click **Packet Rules**.
2. Click **Actions > Rules Editor**. This opens the Packet Rules editor, which allows you to create or edit filter and NAT rules for your system.
3. On the Welcome panel, select **Create a new packet rules file** and click **OK**.
4. From the Packet Rules editor select **Insert > Filter**.
5. On the **General** page, specify a set name for your VPN filter rules. It is recommended that you create at least three different sets: one for your pre-IPSec filter rules, one for your policy filter rules, and one for miscellaneous PERMIT and DENY filter rules. For example, `policyfilters`
6. In the **Action** field, select **IPSEC** from the drop-down list. The **Direction** field defaults to **OUTBOUND** and you cannot change it. The **Action** field is specified as **OUTBOUND** since that is the only direction that IPSec will be applied to packets that are not already in ESP protocol.
7. For **Source address name**, select **=** in the first field, and then enter the IP address of the local data endpoint in the second field. You can also specify a range of IP addresses or an IP address plus a subnet mask after you define them by using the **Define Addresses** function.
8. For **Destination address name**, select **=** in the first field, and then enter the IP address of the remote data endpoint in the second field. You can also specify a range of IP addresses or an IP address plus a subnet mask after you define them by using the **Define Addresses** function.
9. In the **Journaling** field, specify what level of journaling you require.
10. In the **Connection group** field, select the connection definition to which these filter rules apply.
11. (optional) Enter a description.
12. On the **Services** page, select **Service**. This enables the **Protocol**, **Source port**, and **Destination port** fields.
13. In the **Protocol** field, **Source port**, and **Destination port** fields, select the appropriate value for the traffic. Or, you can select the asterisk (*) from the drop-down list. This allows any protocol using any port to use the VPN.
14. Click **OK**.

The next step is to define the interface to which these filter rules apply.

Note: When you add filter rules for an interface, the system automatically adds a default DENY rule for that interface. This means that any traffic not explicitly permitted is denied. You cannot see or change this rule. As a result, you might find that connections that previously worked mysteriously fail after you activate your VPN packet rules. If you want to allow traffic other than VPN on the interface, you must add explicit PERMIT rules to do so.

Related concepts:

“Configuring VPN packet rules” on page 53

If you are creating a connection for the first time, allow VPN to automatically generate the VPN packet rules for you. You can do this by either using the New Connection wizard or the VPN properties pages to configure your connection.

Related tasks:

“Configuring a manual connection” on page 52

A manual connection is one where you must configure all of your VPN properties without using any wizards.

“Configuring the pre-IPSec filter rule” on page 55

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

“Defining an interface for the VPN filter rules”

After you configure your VPN packet rules and any other rules that you need to enable your VPN connection, you must define the interface to which they apply.

Defining an interface for the VPN filter rules

After you configure your VPN packet rules and any other rules that you need to enable your VPN connection, you must define the interface to which they apply.

To define an interface to which to apply your VPN filter rules, follow these steps:

Note: If you just configured the VPN packet rules, the Packet Rules interface will still be open; go to step four.

1. In IBM Navigator for i, expand **Network > IP Policies** and then click **Packet Rules**.
2. Click **Actions > Rules Editor**. This opens the Packet Rules editor, which allows you to create or edit filter and NAT rules for your system.
3. On the Welcome panel, select **Create a new packet rules file** and click **OK**.
4. From the Packet Rules editor select **Insert > Filter Interface**.
5. On the **General** page, select **Line name**, and then select the line description to which your VPN packet rules apply from the drop-down list.
6. (optional) Enter a description.
7. On the **Filter Sets** page, click **Add** to add each set name for the filters you just configured.
8. Click **OK**.
9. Save your rules file. The file is saved into the integrated file system on your system with an extension of .i3p.

Note: Do not save your file into the following directory:

/QIBM/UserData/OS400/TCPIP/RULEGEN

This directory is for system use only. If you ever need to use the RMVTCPTBL *ALL command to deactivate packet rules, the command will delete all files within this directory.

After you define an interface for your filter rules, you must activate them before you can start the VPN.

Related concepts:

“Configuring VPN packet rules” on page 53

If you are creating a connection for the first time, allow VPN to automatically generate the VPN packet rules for you. You can do this by either using the New Connection wizard or the VPN properties pages to configure your connection.

Related tasks:

“Configuring a policy filter rule” on page 56

Complete this task only if you have specified that you do not want VPN to generate your policy filter rule automatically.

“Activating VPN packet rules”

You must activate the VPN packet rules before you can start your VPN connections.

Activating VPN packet rules

You must activate the VPN packet rules before you can start your VPN connections.

You cannot activate (or deactivate) packet rules when you have VPN connections running on your system. So, before you activate your VPN filter rules, ensure that there are no active connections associated with them.

If you created your VPN connections with the New Connection wizard, you can choose to have the associated rules activated, automatically, for you. Be aware that, if there are other packet rules active on any of the interfaces you specify, the VPN policy filter rules will replace them.

If you choose to activate your VPN generated rules by using the Packet Rules Editor, follow these steps:

1. In IBM Navigator for i, expand your **Network > IP Policies** and then click **Packet Rules**.
2. Click **Actions > Activate Rules**. This opens the **Activate Packet Rules** panel.
3. Select whether you want to activate only the VPN generated rules, only a selected file, or both the VPN generated rules and a selected file. You might choose the latter, for instance, if you have miscellaneous PERMIT and DENY rules that you want to enforce on the interface in addition to the VPN generated rules.
4. Select the interface on which you want the rules activated. You can choose to activate on a specific interface, on a point-to-point identifier, or on all interfaces and all point-to-point identifiers.
5. Click **OK** on the dialog box to confirm that you want to verify and activate the rules on the interface or interfaces you specified. After you click OK, the system checks the rules for syntax and semantic errors and reports the results in a message window at the bottom of the editor.

After you activate your filter rules, you can start your VPN connection.

Related concepts:

“Configuring VPN packet rules” on page 53

If you are creating a connection for the first time, allow VPN to automatically generate the VPN packet rules for you. You can do this by either using the New Connection wizard or the VPN properties pages to configure your connection.

Related tasks:

“Configuring a secure VPN connection” on page 51

After you have configured the security policies for your connection, you must then configure the secure connection.

“Defining an interface for the VPN filter rules” on page 57

After you configure your VPN packet rules and any other rules that you need to enable your VPN connection, you must define the interface to which to they apply.

“Starting a VPN connection” on page 60

Complete this task to start connections you will initiate locally.

Configure traffic flow confidentiality

If your data policy is configured for tunnel mode you can use traffic flow confidentiality (TFC) to conceal the actual length of the data packets transferred over a VPN connection.

TFC adds extra padding to the packets being sent and sends dummy packets with different lengths at random intervals to conceal the actual length of the packets. Use TFC for extra security against attackers who might guess the type of data being sent from the length of the packet. When you enable TFC you gain more security, but at the cost of system performance. Therefore, you should test your systems performance before and after you enable TFC on a VPN connection. TFC is not negotiated by IKE, and user should only enable TFC when both systems support it.

To enable TFC on a VPN connection follow these steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**.
- | 2. Click **Secure Connections**, right-click **All Connections**, and select **Open**.
- | 3. Right-click the connection you want to enable TFC and select **Properties**.
- | 4. On the **General** tab select **Use Traffic Flow Confidentiality (TFC) when in Tunnel Mode**.

Configuring extended sequence number

You can use extended sequence number (ESN) to increase the data transfer rate for a VPN connection.

ESN allows you to transmit large volumes of data at a high speed with out re-keying. The VPN connection uses a 64-bit sequence numbers instead of 32-bit numbers over IPSec. Using 64-bit sequence numbers allows more time before re-keying, which prevents sequence number exhaustion and minimizes the use of system resources.

To enable ESN on a VPN connection follow these steps:

- | 1. In IBM Navigator for i, expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
- | 2. Click **Properties**.
- | 3. On the **General** tab select **Use Extended Sequence Number (ESN)**.

| Configuring populate from packet

| A single host to host connection to a single destination system can be automatically generated based on packet information.

| Populate from packet allows individual VPN connections to be negotiated using endpoints based on the information found in an outbound packet. The endpoints must fall within the predefine range of allowed endpoints for the VPN tunnel.

| To configure a connection to enable populate from packet, follow these steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking**, and click **Secure Connections**.
- | 2. Right-click **All Connections** and select **Open**.
- | 3. Right-click on the connection that will be the *local system*, and select **Properties**.
- | 4. On the **General** tab, select **Start on-demand**. Click **OK**.
- | 5. Right-click on the connection again, and select **Group Properties**.
- | 6. On the **General** tab, select **Local system is a host, and the remote system is a gateway**.
- | 7. On the **Connection** tab, select **Customize match to policy filter**. Click **Customize**.
- | 8. Select **Single value** for the **Remote address**. Click **OK**.

Configuring NPF VPN for VIPA

You can allow No Policy Filter (NPF) VPN connections to be applied automatically as preferred interfaces for Virtual Ip addresses (VIPA).

Prior to IBM i 7.2, the NPF connection would only be loaded against the one line description currently indicated as the local choice for the next hop given a route to the desired destination. This causes problems if the Virtual IP addresses are configured, which could allow for multiple paths to a particular destination.

If a NPF connection is being loaded and the local IP address is determined to be a VIPA address, the connection will be loaded against any additional line description that is included in the IP address preferred local interface list.

If the VPN configured for NPF does not work with a particular VIPA address, perform one of the follow 3 options:

NOTE: the following example is specifically for VPN configurations for IBM Universal Care connectivity. In this example the Ip address to the IBM Gateway is 129.42.160.16.

1. Define a host route for the two IBM VPN gateways that bind the route to one particular local interface `ADDTCPRTE RTEDEST('129.42.160.16') SUBNETMASK(*HOST) NEXTHOP('192.168.20.86') BINDIFC('192.168.20.103')`
2. Disable all but one local interface that uses the VIPA.
3. Alter the VPN connection group to predefine filter rules for the VPN connection (for eCare the group names are QIBM01VPN1 and QIBM01VPN2). To do this, perform the following.
 - a. In IBM Navigator for i, expand **Network > IP Policies**. Click **Secure Connections**.
 - b. Right-click **All connections**, and select **Open**.
 - c. Right-click on the profile with the modified IP address, select **Group Properties**.
 - d. In the Connections tab, verify **Generate the following policy filter for this group** is selected. Click **Edit**.
 - e. In the Local Addresses tab, select **Any IP Version 4 Address** for the **Identifier type**.
 - f. In the Remote Addresses tab, enter the VPN endpoint for the **Identifier**.
 - g. In the Services tab, enter 1701 for the **Local port** and the **Remote port**. Select UDP for the **Protocol**. Click **OK**.
 - h. In the Interfaces tab, select the interfaces that will be communicating to the IBM VPN Gateway address. Click **OK**.
 - i. Activate the rule by expanding **Network > IP Policies**
 - j. Click **Packet Rules** to open the **Packet Rules** panel and click **Actions > Activate Rules**. This opens the **Activate Packet Rules** panel.
 - k. Select **Activate only the VPN generated rules**, and select **Activate these rules on all interfaces and all point-to-point filter identifier**. Click **OK**.

Configuration Differences for IKEv2

Compares configuration for IKEv1 and IKEv2.

Existing objects were used as much as possible to enable either IKEv1 or IKEv2 exchanges. This design was used to minimize the impacts to the current GUI interface and VPN configuration objects when IKE version 2 is enabled.

- To enable IKEv2, an IKE Version setting is provided on the Dynamic Connection definition.
- Key Exchange policies can be used for either IKEv1 or IKEv2.
- There are no more differences to other attributes such as Key Exchange policy identifiers (all are still supported) and transforms.
- Main mode/aggressive mode settings are ignored if the Key Exchange Policy is used for IKEv2.
- IBM i 7.1 and IBM Navigator for i are required to configure an IKEv2 connection.

Related concepts:

"IKE version 2" on page 10

IKE version 2 is an enhancement to the Internet key exchange protocol.

Starting a VPN connection

Complete this task to start connections you will initiate locally.

These instructions assume you have properly configured your VPN connection. Follow these steps to start your VPN connection:

- | 1. In IBM Navigator for i, expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
- | 2. If the VPN server is not started, click **Start VPN Server**. This starts the VPN server.
- | 3. Ensure your packet rules are activated.
- | 4. Expand **Network > IP Policies > Virtual Private Networking**.
- | 5. Click **Secure Connections**.
- | 6. Right-click **All Connections** and select **Open** to display a list of connections.
- | 7. Right-click the connection you want to start and select **Start**. To start multiple connections, select each connection you want to start, right-click and select **Start**.
- | The STRVPNCNN command allows you to start a VPN connection from the command line or programmatically.

Related tasks:

“Activating VPN packet rules” on page 58

You must activate the VPN packet rules before you can start your VPN connections.

“Getting started with VPN troubleshooting” on page 63

Complete this task to learn the various methods for determine any VPN problems you are having on your system.

Managing VPN

You can use the VPN interface in IBM Navigator for i to handle all of your VPN management tasks such as, stopping a connection and viewing connection attributes.

Use the VPN interface in IBM Navigator for i to handle all of your management tasks, including:

Setting default attributes for your connections

The default values seed the panels that you use to create new policies and connections. You can set defaults for security levels, key session management, key lifetimes, and connection lifetimes.

The default security values seed various fields when you initially create new VPN objects.

To set default security values for your VPN connections, follow these steps:

- | 1. In IBM Navigator for i, expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
- | 2. Click **Defaults**.
- | 3. Click **Help** if you have questions about how complete a page or any of its fields.
- | 4. Click **OK** after you have completed each of the property sheets.

Resetting connections in error state

Resetting connections in error returns them to the idle state.

To refresh a connection that is in the error state, follow these steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and then click **Secure Connections**.
- | 2. Right-click **All Connections** and select **Open** to display a list of connections in the right pane.
- | 3. Right-click the connection you want to reset and select **Reset**. This resets the connection to the idle state. To reset multiple connections that are in the error state, select each connection you want to reset, right-click and select **Reset**.

Viewing error information

Complete this task to help you determine why your connection is causing an error.

To view information about connections in error, follow these steps:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
2. Right-click **All Connections** to display a list of connections in the right pane.
3. Right-click the connection in error that you want to view and select **Error Information**.

Related tasks:

“Getting started with VPN troubleshooting” on page 63

Complete this task to learn the various methods for determine any VPN problems you are having on your system.

Viewing attributes of active connections

Complete this task to check the status and other attributes of your active connections.

To view the current attributes of an active or on-demand connection, follow these steps:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
2. Right-click **All Connections** and select **Open** to display a list of connections in the right pane.
3. Right-click the active or on-demand connection that you want to view and select **Properties**.
4. Go to the **Current Attributes** page to view the attributes of the connection.

You can also view the attributes of all connections from the IBM Navigator for i window. By default, the only attributes that display are Status, Description, and Connection Type. You can change what data displays by following these steps:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
2. Right-click **All Connections** and select **Open** to display a list of connections in the right pane.
3. From the **Objects** menu, select **Columns**. This opens a dialog box which allows you to select which attributes you want to display in the IBM Navigator for i panel.

Be aware that when you change the columns to view, the changes are not specific to a particular user or PC, but rather, are system-wide.

Related concepts:

“Common VPN Connection Manager error messages” on page 75

The VPN Connection Manager logs two messages in the QTOVMAN job log when an error occurs with a VPN connection.

Viewing VPN server job logs

Follow these instructions to view the job logs for the VPN Key Manager and the VPN Connection Manager.

To view the current job logs of either the VPN Key Manager or the VPN Connection Manager, follow these steps:

1. In IBM Navigator for i, expand **Network > All Tasks > IP Policies > Virtual Private Networking**.
2. Click **Display Server Jobs** to view the job logs.

Viewing attributes of Security Associations

Complete this task to display the attributes of the Security Associations (SAs) that are associated with an enabled connection.

To view the attributes of the security associations (SAs) that are associated with an enabled connection. To do so, follow these steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
- | 2. Right-click **All Connections** and select **Open** to display a list of connections in the right pane.
- | 3. Right-click the appropriate active connection and select **Security Associations**. The resulting panel allows you to view the properties for each of the SAs associated with a specific connection.

Stopping a VPN connection

Complete this task to stop active connections.

To stop an active or on-demand connection, follow these steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
- | 2. Right-click **All Connections** to display a list of connections in the right pane.
- | 3. Right-click the connection you want to stop and select **Stop**. To stop multiple connections, select each connection you want to stop, right-click and select **Stop**.
- | The ENDVPNCNN command allows you to end a VPN connection from the command line or programmatically.

Deleting VPN configuration objects

Before you delete a VPN configuration object from the VPN policy database, make sure you understand how it affects other VPN connections and connection groups.

If you are certain you need to delete a VPN connection from the VPN policy database, perform the following steps:

- | 1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and then click **Secure Connections**.
- | 2. Right-click **All Connections** to display a list of connections in the right pane.
- | 3. Right-click the connection you want to delete and select **Delete**.

Troubleshooting VPN

Use the following troubleshooting methods to solve some of the basic problems you might experience while configuring a VPN connection.

VPN is a complex and rapidly changing technology that requires at least a basic knowledge of standard IPSec technologies. You must also be familiar with IP packet rules because VPN requires several filter rules in order to work properly. Because of this complexity, you might, from time to time, experience trouble with your VPN connections. Troubleshooting your VPN is not always an easy task. You must understand your system and your network environments, as well as the components you use to manage them. The following topics provide you with hints on how to troubleshoot the various problems you might encounter while using VPN:

Getting started with VPN troubleshooting

Complete this task to learn the various methods for determine any VPN problems you are having on your system.

There are several ways to begin analyzing VPN problems:

1. Always make sure that you have applied the latest Program Temporary Fixes (PTFs).
2. Ensure that you meet the minimum VPN setup requirements.
3. Review any error messages that are found in the Error Information window or in the VPN server job logs for both the local and the remote systems. In fact, when you are troubleshooting VPN connection problems it is often necessary to look at both ends of the connection. Further, you need to take into account that there are four addresses you must check: The local and remote connection endpoints, which are the addresses where IPsec is applied to the IP packets, and the local and remote data endpoints, which are the source and destination addresses of the IP packets.
4. If the error messages you find do not provide enough information to solve the problem, check the IP filter journal.
5. The communication trace on the system offers you a another place to find general information about whether the local system receives or sends connection requests.
6. The Trace TCP Application (TRCTCPAPP) command provides yet another way to isolate problems. Typically, IBM Service uses TRCTCPAPP to obtain trace output in order to analyze connection problems.

Related concepts:

“VPN setup requirements” on page 45

For a VPN connection to function properly on your systems and with network clients, you must meet the minimum requirements

“Troubleshooting VPN with the VPN job logs” on page 74

When you encounter problems with your VPN connections, it is always advisable to analyze the job logs. In fact, there are several job logs that contain error messages and other information related to a VPN environment.

“Troubleshooting VPN with the communications trace” on page 80

IBM i provides the capability to trace data on a communications line, such as a local area network (LAN) or wide area network (WAN) interface. The average user might not understand the entire contents of the trace data. However, you can use the trace entries to determine whether a data exchange between the local and the remote systems took place.

Related tasks:

“Viewing error information” on page 62

Complete this task to help you determine why your connection is causing an error.

“Troubleshooting VPN with the QIPFILTER journal” on page 69

View this information to learn about VPN filter rules.

“Starting a VPN connection” on page 60

Complete this task to start connections you will initiate locally.

Other things to check

If an error occurs after you set up a connection, and you are not sure where in the network the error occurred, try reducing the complexity of your environment. For example, instead of investigating all parts of a VPN connection at one time, start with the IP connection itself. The following list gives you some basic guidelines on how to start VPN problem analysis, from the simplest IP connection to the more complex VPN connection:

1. Start with an IP configuration between the local and remote host. Remove any IP filters on the interface that both the local and remote system use for communicating. Can you PING from the local to the remote host?

Note: Remember to prompt on the PING command; enter the remote system address and use PF10 for additional parameters, then enter the local IP address. This is particularly important when you have multiple physical or logical interfaces. It ensures that the right addresses are placed in the PING packets.

If you answer **yes**, then proceed to step 2. If you answer **no**, then check your IP configuration, interface status, and routing entries. If the configuration is correct, use a communication trace to check, for example, that a PING request leaves the system. If you send a PING request but you receive no response, the problem is most likely the network or remote system.

Note: There might be intermediate routers or firewall that do IP packet filtering and might be filtering the PING packets. PING is typically based on the ICMP protocol. If the PING is successful, you know you have connectivity. If the PING is unsuccessful, you only know the PING failed. You might want to try other IP protocols between the two systems, such as Telnet or FTP to verify connectivity.

2. Check the filter rules for VPN and ensure that they are activated. Does filtering start successfully? If you answer **yes**, then proceed to step 3. If you answer **no**, then check for error messages in the Packet Rules panel in IBM Navigator for i. Ensure that the filter rules do not specify Network Address Translation (NAT) for any VPN traffic.
3. Start your VPN connection. Does the connection start successfully? If you answer **yes**, then proceed to step 4. If you answer **no**, then check the QTOVMAN, QTOKVPNIKE, and QTOKVPNIK2 job logs for errors. When you use VPN, your Internet Service Provider (ISP) and every security gateway in your network must support the Authentication Header (AH) and Encapsulated Security Payload (ESP) protocols. Whether you choose to use AH or ESP depends on the proposals you define for your VPN connection.
4. Are you able to activate a user session over the VPN connection? If you answer **yes**, then the VPN connection works as required. If you answer **no**, then check the packet rules and the VPN dynamic-key groups and connections for filter definitions that do not allow the user traffic you want.

Common VPN configuration errors and how to fix them

Use this information to review common VPN error messages and learn their possible resolutions.

Note: When you configure VPN, you are actually creating several different configuration objects, each of which VPN requires to enable a connection. In terms of the VPN GUI, these objects are: The IP Security Policies and the Secure Connections. So, when this information refers to an object, it is referring to one or more of these parts of the VPN.

VPN error message: TCP5B28

When you attempt to activate filter rules on an interface, you get this message: TCP5B28 CONNECTION_DEFINITION order violation

Symptom:

When you attempt to activate filter rules on a specific interface, you receive this error message:
TCP5B28: CONNECTION_DEFINITION order violation

Possible resolution:

The filter rules you were attempting to activate contained connection definitions that were ordered differently than in a previously activated set of rules. The easiest way to resolve this error is to activate the rules file on **all interfaces** instead of on a specific interface.

VPN error message: Item not found

When you right-click a VPN object and select either **Properties** or **Delete**, you get a message that says, **Item not found**.

Symptom:

When you right-click an object in the Virtual Private Networking window and either select **Properties** or **Delete**, the following message appears:



Possible resolution:


- You might have deleted the object or renamed it, and have not refreshed the window yet. Consequently, the object still appears in the Virtual Private Networking window. To verify that this is the case, from the **View** menu, select **Refresh**. If the object still appears in the Virtual Private Networking window, continue to the next item in this list.
- When you configured the properties for the object, a communication error might have occurred between the VPN server and your system. Many of the objects that appear in the VPN window relate to more than one object in the VPN policy database. This means that communication errors might cause some of the objects in the database to continue to be related to an object in the VPN. Whenever you create or update an object, an error will occur when the loss of synchronization actually happens. The only way to fix the problem is to select **OK** on the error window. This launches the property sheet for the object in error. Only the name field on the property sheet has a value in it. Everything else is blank (or contains default values). Enter the correct attributes of the object, and select **OK** to save your changes.
- A similar error occurs when you try to delete the object. To fix this problem, complete the blank property sheet that opens when you click **OK** on the error message. This updates any links to the VPN policy database that were lost. You can now delete the object.

VPN error message: **PARAMETER PINBUF IS NOT VALID**

When you attempt to start a connection, you get a message that says, **PARAMETER PINBUF IS NOT VALID...**

Symptom:

When you attempt to start a connection, a message similar to the following appears:



Possible resolution:

This happens when your system is set to use certain locales to which lowercase letters do not map correctly. To fix this error, either make sure that all objects use only uppercase letters, or change the locale of the system.

VPN error message: Item not found, Remote key server...

When you select **Properties** for a dynamic-key connection, you get an error that says that the server cannot find the remote key server you specified.

Symptom:

When you select **Properties** for a dynamic-key connection, a message like the following appears:



Possible resolution:

This happens when you create a connection with a particular remote key server identifier, and then the remote key server is removed from its dynamic-key group. To fix this error, click **OK** on the error message. This opens the property sheet for the dynamic-key connection that is in error. From here, you can either add the remote key server back into the dynamic-key group, or select another remote key server identifier. Click **OK** on the property sheet to save your changes.

VPN error message: Unable to update the object

When you select **OK** on the property sheet for a dynamic-key group or manual connection, you get a message that tells you the system cannot update the object.

Symptom:

When you select **OK** on the property sheet for a dynamic-key group or manual connection, the following message appears:



Possible resolution:

This error happens when an active connection is using the object you are trying to change. You cannot make changes to an object within an active connection. To make changes to an object, identify the appropriate active connection and then, right-click it and select **Stop** from the resulting context menu.

VPN error message: Unable to encrypt key...

You get a message that says that the system cannot encrypt your keys because the QRETSVRSEC value must be set to 1.

Symptom:

The following error message appears:



Possible resolution:

QRETSVRSEC is a system value that indicates whether your system can store encrypted keys on it. If this value is set to 0, then preshared keys and the keys for the algorithms in a manual connection cannot be stored in the VPN policy database. To fix this problem, use a 5250 emulation session to your system. Type `wrksysval` at the command line and press **Enter**. Look for QRETSVRSEC in the list and type 2 (change) next to it. On the next panel, type 1 and press **Enter**.

Related concepts:

“VPN error: All keys are blank”

When you view the properties of a manual connection, all preshared keys and the algorithm keys for the connection are blank.

VPN error message: CPF9821

When you try to expand or open the IP Policies container in IBM Navigator for i, the CPF9821- Not authorized to program QTFRPRS in QSYS library message appears.

Symptom:

When you try to expand the IP Policies container in IBM Navigator for i, the CPF9821- Not authorized to program QTFRPRS in QSYS library message appears.

Possible resolution:

You might not have the required authority to retrieve the current status of Packet Rules or the VPN connection manager. Ensure that you have *IOSYSCFG authority to get access to the Packet Rules functions in IBM Navigator for i.

VPN error: All keys are blank

When you view the properties of a manual connection, all preshared keys and the algorithm keys for the connection are blank.

Symptom:

All preshared keys and the algorithm keys for manual connections are blank.

Possible resolution:

This happens whenever the system value QRETSVRSEC is set back to 0. Setting this system value to 0 erases all of the keys in the VPN policy database. To fix this problem, you must set the system value to 1 and then reenter all of the keys. Refer to Error message: Unable to encrypt keys for more information about how to do this.

Related concepts:

“VPN error message: Unable to encrypt key...” on page 67

You get a message that says that the system cannot encrypt your keys because the QRETSVRSEC value must be set to 1.

VPN error: Connection has enabled status after you stop it

- | After you stop a connection, the IBM Navigator for i panel indicates that the connection is still enabled.

Symptom:

After you stop a connection, the IBM Navigator for i panel indicates that the connection is still enabled.

Possible resolution:

This typically happens because you have not refreshed the connection list panel on IBM Navigator for i yet. As such, the connection list contains outdated information. To fix this, click the **Refresh** button on the list menu.

VPN error: Active filter rules fail to deactivate

When you try to deactivate the current set of filter rules, the message, The active rules failed to be deactivated appears in the results window.

Symptom:

When you try to deactivate the current set of filter rules, the message, The active rules failed to be deactivated appears in the results window.

Possible resolution:

Typically, this error message means that there is at least one active VPN connection. You must stop each of the connections that have a status of enabled. To do this, right-click each of the active connections and select **Stop**. You can now deactivate the filter rules.

VPN error: The key connection group for a connection changes

When you create a dynamic-key connection, you specify a dynamic-key group and an identifier for the remote key server. Later, when you view the properties of the related connection object, the General page of the property sheet displays the same remote key server identifier, but a different dynamic-key group.

Symptom:

When you create a dynamic-key connection, you specify a dynamic-key group and an identifier for the remote key server. Later, when you select **Properties** on the related connection object, the **General** page of the property sheet displays the same remote key server identifier, but a different dynamic-key group.

Possible resolution:

The identifier is the only information stored in the VPN policy database that refers to the remote key server of the dynamic-key connection. When VPN looks up a policy for a remote key server, it looks for the first dynamic-key group that has that remote key server identifier in it. So, when you view the properties for one these connections, it uses the same dynamic-key group that VPN found. If you do not want to associate the dynamic-key group with that remote key server, you can do one of the following:

1. Remove the remote key server from the dynamic-key group.
2. Expand **By Groups** in the left pane of the VPN interface, and select and drag the dynamic-key group you want to the top of the table in the right pane. This ensures that VPN checks this dynamic-key group first for the remote key server.

Troubleshooting VPN with the QIPFILTER journal

View this information to learn about VPN filter rules.

The QIPFILTER journal is located in the QUSRSYS library and contains information about filter rule sets, as well as information about whether an IP datagram was permitted or denied. Logging is performed based on the journaling option you specify in your filter rules.

Related tasks:

“Getting started with VPN troubleshooting” on page 63

Complete this task to learn the various methods for determine any VPN problems you are having on your system.

Enabling the QIPFILTER journal

Use the Packet Rules editor in IBM Navigator for i to activate the QIPFILTER journal.

You must enable the logging function for each individual filter rule. There is no function that allows logging for all IP datagrams going into or out of the system.

Note: To enable the QIPFILTER journal, your filters must be deactivated.

The following steps describe how to enable journaling for a particular filter rule:

1. In IBM Navigator for i, expand **Network > IP Policies** and click **Packet Rules**.
2. Click **Actions > Rules Editor**.
3. Open an existing filter rule file in the **rules Editor**.
4. Select the filter rule you want to journal, and then select **Actions > Properties**.
5. On the **General** page, select **FULL** in the **Journaling** field as in the dialog box shown above. This enables logging for this particular filter rule.
6. Click **OK**.
7. Save and activate the changed filter rule file.

If an IP datagram matches the definitions of the filter rule, an entry is made in the QIPFILTER journal.

Using the QIPFILTER journal

IBM i automatically creates the journal the first time you activate IP packet filtering.

To view the entry-specific details in the journal, you can display the journal entries on the screen or you can use an output file. By copying the journal entries to the output file, you can easily view the entries using query utilities such as Query/400 or SQL. You can also write your own HLL programs to process the entries in the output files.

The following is an example of the Display Journal (DSPJRN) command:

```
DSPJRN JRN(QIPFILTER) JRNCDE((M)) ENTYP((TF)) OUTPUT(*OUTFILE)
      OUTFILFMT(*TYPE4) OUTFILE(mylib/myfile) ENTDTALEN(*VARLEN *CALC)
```

Use the following steps to copy the QIPFILTER journal entries to the output file:

1. Create a copy of the system-supplied output file QSYS/QATOFIPF into a user library by using the Create Duplicate Object (CRTDUPOBJ) command. The following is an example of the CRTDUPOBJ command:

```
CRTDUPOBJ OBJ(QATOFIPF) FROMLIB(QSYS) OBJTYPE(*FILE) TOLIB(mylib)
      NEWOBJ(myfile)
```
2. Use the Display Journal (DSPJRN) command to copy the entries from the QUSRSYS/QIPFILTER journal to the output file you created in the previous step.

If you copy the DSPJRN into an output file that does not exist, the system creates a file for you, but this file does not contain the correct field descriptions.

Note: The QIPFILTER journal only contains permit or deny entries for filter rules where the journaling option is set to FULL. For example, if you set up only PERMIT filter rules, IP datagrams that are not explicitly permitted are denied. For those denied datagrams, no entry is added to the journal. For problem analysis you might add a filter rule that explicitly denies all other traffic and performs FULL journaling. Then, you will get DENY entries in the journal for all IP datagrams that are denied. Due to performance reasons, it is not recommended that you enable journaling for all filter rules. Once your filter sets are tested, reduce the journaling to a useful subset of entries.

Related concepts:

“QIPFILTER journal fields”

Review the following table that describes the fields in the QIPFILTER output file

QIPFILTER journal fields

Review the following table that describes the fields in the QIPFILTER output file

Field Name	Field Length	Numeric	Description	Comments
TFENTL	5	Y	Length of entry	
TFSEQN	10	Y	Sequence number	
TFCODE	1	N	Journal code	Always M
TFENTT	2	N	Entry type	Always TF
TFTIME	26	N	SAA timestamp	
TFJOB	10	N	Job name	
TFUSER	10	N	User profile	
TFNBR	6	Y	Job number	
TFPGM	10	N	Program name	
TFRES1	51	N	Reserved	
TFUSPF	10	N	User	
TFSYMN	8	N	System name	
TFRES2	20	N	Reserved	
TFRESA	50	N	Reserved	
TFLINE	10	N	Line description	*ALL if TFREVT is U* , Blank if TFREVT is L* , Line name if TFREVT is L
TFREVT	2	N	Rule event	L* or L when rules are loaded. U* when rules are unloaded, A when filter action
TFPDIR	1	N	IP Packet direction	O is outbound, I is inbound
TFRNUM	5	N	Rule number	Applies to the rule number in the active rules file
TFACT	6	N	Filter action taken	PERMIT, DENY, or IPSEC
TFPROT	4	N	Transport protocol	1 is ICMP 6 is TCP 17 is UDP 50 is ESP 51 is AH
TFSRCA	15	N	Source IP address	
TFSRCP	5	N	Source port	Garbage if TFPROT= 1 (ICMP)
TFDSTA	15	N	Destination IP address	

Field Name	Field Length	Numeric	Description	Comments
TFDSTP	5	N	Destination port	Garbage if TFPROT= 1 (ICMP)
TFTEXT	76	N	Additional text	Contains description if TFREVT= L* or U*

Related tasks:

“Using the QIPFILTER journal” on page 70

IBM i automatically creates the journal the first time you activate IP packet filtering.

Troubleshooting VPN with the QVPN journal

Provides information about IP traffic and connections.

VPN uses a separate journal to log information about the IP traffic and connections called the QVPN journal. The QVPN is stored in the QUSRSYS library. The journal code is M and the journal type is TS. You will rarely use journal entries on a daily basis. Instead, you might find them useful for troubleshooting and verifying that your system, keys, and connections are functioning in the manner that you specified. For example, journal entries help you understand what happens to your data packets. They also keep you informed of your current VPN status.

Enabling the QVPN journal

Use the virtual private networking interface in IBM Navigator for i to activate the VPN journal.

There is no function that allows logging for all VPN connections. Therefore, you must enable the logging function for each individual dynamic-key group or manual connection.

The following steps describe how to enable the journal function for a particular dynamic-key group or manual connection:

1. In IBM Navigator for i, expand **Network > IP Policies > Virtual Private Networking** and click **Secure Connections**.
2. For dynamic-key groups, right-click **By Group** and select **Open** to see the groups. Right-click the dynamic key group for which you want to enable journaling and select **Properties**.
3. For manual connections, right-click **Manual Connections** and select **Open** to show **All Connections**. Right-click the manual connection for which you want to enable journaling.
4. On the **General** page, select the level of journaling you require. You have the choice between four options. These are:
 - None** No journaling occurs for this connection group.
 - All** Journaling occurs for all connection activities, such as starting or stopping a connection, or key refreshes, as well as IP traffic information.
 - Connection Activity** Journaling occurs for such connection activity as starting or stopping a connection.
 - IP traffic** Journaling occurs for all of the VPN traffic that is associated with this connection. A log entry is made every time a filter rule is invoked. The system records IP traffic information in the journal QIPFILTER, which is located in the QUSRSYS library.
5. Click **OK**.
6. Start the connection to activate journaling.

Note: Before you can stop journaling, make sure that the connection is inactive. To change the journaling status of a connection group, make sure that no active connections are associated with that particular group.

Using the QVPN journal

To view the entry-specific details in the VPN journal, you can display the entries on the screen or you can use the output file.

By copying the journal entries to the output file, you can easily view the entries by using query utilities such as Query/400 or SQL. You can also write your own HLL programs to process the entries in the output files. The following is an example of the Display Journal (DSPJRN) command:

```
DSPJRN JRN(QVPN) JRNCDE((M)) ENTYP((TS)) OUTPUT(*OUTFILE) OUTFILMT(*TYPE4)
      OUTFILE(mylib/myfile) ENTDTALEN(*VARLEN *CALC)
```

Use the following steps to copy the VPN journal entries to the output file:

1. Create a copy of the system-supplied output file QSYS/QATOVSOFF into a user library. You can do this by using the Create Duplicate Object (CRTDUPOBJ) command. The following is an example of the CRTDUPOBJ command:

```
CRTDUPOBJ OBJ(QATOVSOFF) FROMLIB(QSYS) OBJTYPE(*FILE) TOLIB(mylib)
      NEWOBJ(myfile)
```
2. Use the Display Journal (DSPJRN) command to copy the entries from the QUSRSYS/QVPN journal to the output file created in the previous step. If you attempt to copy the DSPJRN into an output file that does not exist, the system creates a file for you, but this file does not contain the correct field descriptions.

Related concepts:

“QVPN journal fields”

Review the following table that describes the fields in the QVPN output file.

QVPN journal fields

Review the following table that describes the fields in the QVPN output file.

Field Name	Field Length	Numeric	Description	Comments
TSENTL	5	Y	Length of entry	
TSSEQN	10	Y	Sequence number	
TSCODE	1	N	Journal code	Always M
TSENTT	2	N	Entry type	Always TS
TSTIME	26	N	SAA entry timestamp	
TSJOB	10	N	Name of job	
TSUSER	10	N	User of job	
TSNBR	6	Y	Number of job	
TSPGM	10	N	Name of program	
TSRES1	51	N	Not used	
TSUSPF	10	N	User profile name	
TSSYNM	8	N	System name	
TSRES2	20	N	Not used	
TSRESA	50	N	Not used	
TSESDL	4	Y	Length of specific data	
TSCMPN	10	N	VPN component	
TSCONM	40	N	Connection name	
TSCOTY	10	N	Connection type	

Field Name	Field Length	Numeric	Description	Comments
TSCOS	10	N	Connection state	
TSCOSD	8	N	Start date	
TSCOST	6	N	Start time	
TSCOED	8	N	End date	
TSCOET	6	N	End time	
TSTRPR	10	N	Transport protocol	
TSLCAD	43	N	Local client address	
TSLCPR	11	N	Local ports	
TSRCAD	43	N	Remote client address	
TSCPR	11	N	Remote ports	
TSLEP	43	N	Local endpoint	
TSREP	43	N	Remote endpoint	
TSCORF	6	N	Times refreshed	
TSRFDA	8	N	Date of next refresh	
TSRFTI	6	N	Time of next refresh	
TSRFLS	8	N	Refresh life-size	
TSSAPH	1	N	SA Phase	
TSAUTH	10	N	Authentication type	
TSENCR	10	N	Encryption type	
TSDHGR	2	N	Diffie-Hellman group	
TSERRC	8	N	Error code	

Related tasks:

“Using the QVPN journal” on page 73

To view the entry-specific details in the VPN journal, you can display the entries on the screen or you can use the output file.

Troubleshooting VPN with the VPN job logs

When you encounter problems with your VPN connections, it is always advisable to analyze the job logs. In fact, there are several job logs that contain error messages and other information related to a VPN environment.

It is important that you analyze job logs on both sides of the connection if both sides are IBM i models. When a dynamic connection fails to start, it is helpful if you understand what is happening on the remote system.

The VPN jobs, QTOVMAN, QTOKVPNIKE, and QTOKVPNIK2, run in the subsystem QSYSWRK. You can view their respective job logs from IBM Navigator for i.

This section introduces the most important jobs for a VPN environment. The following list shows the job names with a brief explanation of what the job is used for:

QTCPIP

This job is the base job that starts all the TCP/IP interfaces. If you have fundamental problems with TCP/IP in general, analyze the QTCPIP job log.

QTOKVPNIKE

The QTOKVPNIKE job is the VPN key manager job. The VPN key manager listens to UDP port 500 to perform the Internet Key Exchange (IKE) protocol processing.

QTOKVPNIK2

The QTOKVPNIK2 job is the VPN key manager job for IKEv2. The VPN key manager listens to UDP port 500 to perform the Internet Key Exchange version 2 (IKEv2) protocol processing.

QTOVMAN

This job is the connection manager for VPN connections. The related job log contains messages for every connection attempt that fails.

QTPPANStxx

This job is used for PPP dial-up connections. It answers to connection attempts where *ANS is defined in a PPP profile.

QTPPPCTL

This is a PPP job for dial-out connections.

QTPPPL2TP

This is the Layer Two Tunneling Protocol (L2TP) manager job. If you have problems setting up an L2TP tunnel, look for messages in this job log.

Related concepts:

“IKE version 2” on page 10

IKE version 2 is an enhancement to the Internet key exchange protocol.

Related tasks:

“Getting started with VPN troubleshooting” on page 63

Complete this task to learn the various methods for determine any VPN problems you are having on your system.

Common VPN Connection Manager error messages

The VPN Connection Manager logs two messages in the QTOVMAN job log when an error occurs with a VPN connection.

The first message provides details regarding the error. You can view information about these errors in IBM Navigator for i by right-clicking the connection in error and selecting **Error Information**.

The second message describes the action you were attempting to perform on the connection when the error occurred. For example, starting or stopping it. Messages TCP8601, TCP8602, and TCP860A, described below, are typical examples of these second messages.

VPN Connection Manager error messages

Message	Cause	Recovery
TCP8601 Could not start VPN connection [<i>connection name</i>]	Could not start this VPN connection due to one of these reason codes: 0 - A previous message in the job log with the same VPN connection name has more detailed information. 1 - VPN policy configuration. 2 - Communications network failure. 3 - VPN Key Manager failed to negotiate a new security association. 4 - The remote endpoint for this connection is not configured properly. 5 - VPN Key Manager failed to respond to VPN Connection Manager. 6 - IP Security Component VPN connection load failure. 7 - PPP Component failure.	1. Check the job logs for additional messages. 2. Correct the errors and try the request again. 3. Use IBM Navigator for i to view the connection status. Connections that could not start will be in error state.

VPN Connection Manager error messages

Message	Cause	Recovery
TCP8602 Error occurred stopping VPN connection [<i>connection name</i>]	The specified VPN connection was requested to be stopped, however, it did not stop or stopped in error due to Reason Code: 0 - A previous message in the job log with the same VPN connection name has more detailed information. 1 - The VPN connection does not exist. 2 - Internal communications failure with VPN Key Manager. 3 - Internal communications failure with IPSec component. 4 - Communication failure with VPN connection remote endpoint.	1. Check the job logs for additional messages. 2. Correct the errors and try the request again. 3. Use IBM Navigator for i to view the connection status. Connections that could not start will be in error state.
TCP8604 Start of VPN connection [<i>connection name</i>] failed	A start of this VPN connection failed due to one of these reason codes: 1 - Could not translate the remote host name to an IP address. 2 - Could not translate the local host name to an IP address. 3 - VPN policy filter rule associated with this VPN connection is not loaded. 4 - A user-specified key value is not valid for its associated algorithm. 5 - The initiation value for the VP connection does not allow the specified action. 6 - A system role for the VPN connection is inconsistent with information from the connection group. 7 - Reserved. 8 - Data endpoints (local and remote addresses and services) of this VPN connection are inconsistent with information from the connection group. 9 - Identifier type not valid.	1. Check the job logs for additional messages. 2. Correct the errors and try the request again. 3. Use IBM Navigator for i to check or correct the VPN policy configuration. Ensure that the dynamic-key group associated with this connection has acceptable values configured.
TCP8605 VPN Connection Manager could not communicate with VPN Key Manager	The VPN Connection Manager requires the services of the VPN Key Manager to establish security associations for dynamic VPN connections. The VPN Connection Manager could not communicate with the VPN Key Manager.	1. Check the job logs for additional messages. 2. Verify the *LOOPBACK interface is active by using the NETSTAT OPTION(*IFC) command. 3. End the VPN server by using the ENDTCPSPVR SERVER(*VPN) command. Then restart the VPN server by using the STRTCPSRV SERVER(*VPN) command. Note: This causes all current VPN connections to end.

VPN Connection Manager error messages

Message	Cause	Recovery
TCP8606 The VPN Key Manager could not establish the requested security association for connection, [<i>connection name</i>]	The VPN Key Manager could not establish the requested security association due to one of these reason codes: 24 - VPN Key Manager key connection authentication failed. 8300 - Failure occurred during VPN Key Manager key connection negotiations. 8306 - No local preshared key found. 8307 - No remote IKE phase 1 policy found. 8308 - No remote preshared key found. 8327 - VPN Key Manager key connection negotiations timed out. 8400 - Failure occurred during VPN Key Manager VPN connection negotiations. 8407 - No remote IKE phase 2 policy found. 8408 - VPN Key Manager VPN connection negotiations timed out. 8500 or 8509 - VPN Key Manager network error has occurred.	1. Check the job logs for additional messages. 2. Correct the errors and try the request again. 3. Use IBM Navigator for i to check or correct the VPN policy configuration. Ensure that the dynamic-key group associated with this connection has acceptable values configured.
TCP8608 VPN connection, [<i>connection name</i>], could not obtain a NAT address	This dynamic-key group or data connection specified that network address translation (NAT) be done on one or more addresses, and that failed due to one of these likely reason codes: 1 - Address to apply NAT to is not a single IP address. 2 - All available addresses have been used.	1. Check the job logs for additional messages. 2. Correct the errors and try the request again. 3. Use IBM Navigator for i to check or correct the VPN policy. Ensure that the dynamic-key group associated with this connection has acceptable values for addresses configured.
TCP8620 Local connection endpoint not available	Could not enable this VPN connections because the local connection endpoint was not available.	1. Check the job logs for additional messages pertaining to this connection. 2. Make sure the local connection endpoint is defined and started by using the NETSTAT OPTION(*IFC) command. 3. Correct any errors and try the request again.
TCP8621 Local data endpoint to available	Could not enable this VPN connection because the local data endpoint was not available.	1. Check the job logs for additional messages pertaining to this connection. 2. Make sure the local connection endpoint is defined and started by using the NETSTAT OPTION(*IFC) command. 3. Correct any errors and try the request again.

VPN Connection Manager error messages

Message	Cause	Recovery
TCP8622 Transport encapsulation not permitted with a gateway	Could not enable this VPN connection because the negotiated policy specified transport encapsulation mode and this connection is defined as a security gateway.	1. Check the job logs for additional messages pertaining to this connection. 2. Use IBM Navigator for i to change the VPN policy associated with this VPN connection. 3. Correct any errors and try the request again.
TCP8623 VPN connection overlaps with an existing one	Could not enable this VPN connection because an existing VPN connection is already enabled. This connection has a local data endpoint of, <i>[local data endpoint value]</i> and a remote data endpoint of, <i>[remote data endpoint value]</i> .	1. Check the job logs for additional messages pertaining to this connection. 2. Use IBM Navigator for i to view all enabled connections that have local data endpoints and remote data endpoints overlapping the connection. Change the policy of the existing connection if both connections are required. 3. Correct any errors and try the request again.
TCP8624 VPN connection not within scope of associated policy filter rule	Could not enable this VPN connection because the data endpoints are not within the defined policy filter rule.	1. Check the job logs for additional messages pertaining to this connection. 2. Use IBM Navigator for i to display the data endpoint restrictions for this connection or dynamic-key group. If Subset of policy filter or Customize to match policy filter is selected, then check the data endpoints of the connection. These must fit within the active filter rule that has an IPSEC action and a VPN connection name associated with this connection. Change the existing connection's policy or the filter rule to enable this connection. 3. Correct any errors and try the request again.
TCP8625 VPN connection failed an ESP algorithm check	Could not enable this VPN connection because the secret key associated with the connection was insufficient.	1. Check the job logs for additional messages pertaining to this connection. 2. Use IBM Navigator for i to display the policy associated with this connection and enter a different secret key. 3. Correct any errors and try the request again.

VPN Connection Manager error messages

Message	Cause	Recovery
TCP8626 VPN connection endpoint is not the same as the data endpoint	Could not enable this VPN connection because the policy specifies that it is a host, and the VPN connection endpoint is not the same as the data endpoint.	1. Check the job logs for additional messages pertaining to this connection. 2. Use IBM Navigator for i to display the data endpoint restrictions for this connection or dynamic-key group. If Subset of policy filter or Customize to match policy filter is selected, then check the data endpoints of the connection. These must fit within the active filter rule that has an IPSEC action and a VPN connection name associated with this connection. Change the existing connection's policy or the filter rule to enable this connection. 3. Correct any errors and try the request again.
TCP8628 Policy filter rule not loaded	The policy filter rule for this connection is not active.	1. Check the job logs for additional messages pertaining to this connection. 2. Use IBM Navigator for i to display the active policy filters. Check the policy filter rule for this connection. 3. Correct any errors and try the request again.
TCP8629 IP packet dropped for VPN connection	This VPN connection has VPN NAT configured and the required set of NAT addresses has exceeded the available NAT addresses.	1. Check the job logs for additional messages pertaining to this connection. 2. Use IBM Navigator for i to increase the number of NAT addresses assigned for this VPN connection. 3. Correct any errors and try the request again.
TCP862A PPP connection failed to start	This VPN connection was associated with a PPP profile. When it was started, an attempt was made to start the PPP profile, but a failure occurred.	1. Check the job logs for additional messages pertaining to this connection. 2. Check the job log associated with the PPP connection. 3. Correct any errors and try the request again.

Related tasks:

"Viewing attributes of active connections" on page 62

Complete this task to check the status and other attributes of your active connections.

Troubleshooting VPN with the communications trace

IBM i provides the capability to trace data on a communications line, such as a local area network (LAN) or wide area network (WAN) interface. The average user might not understand the entire contents of the trace data. However, you can use the trace entries to determine whether a data exchange between the local and the remote systems took place.

Starting the communications trace

Use the Start Communications Trace (STRCMNTRC) command to start the communications trace on your system. The following is an example of the STRCMNTRC command:

```
STRCMNTRC CFGOBJ(TRNLINE) CFGTYPE(*LIN) MAXSTG(2048) TEXT('VPN Problems')
```

The command parameters are explained in the following list:

CFGOBJ (Configuration object)

The name of the configuration object to trace. The object is either a line description, a network interface description, or a network server description.

CFGTYPE (Configuration type)

Whether a line (*LIN), a network interface (*NWI), or a network server (*NWS) is being traced.

MAXSTG (Buffer size)

The buffer size for the trace. The default value is set to 128 KB. The range goes from 128 KB to 64 MB. The actual maximum system-wide buffer size is defined within the System Service Tools (SST). Therefore, you might receive an error message when using a larger buffer size on the STRCMNTRC command than defined in the SST. Keep in mind that the sum of buffer sizes specified on all started communications traces must not exceed the maximum buffer size defined in the SST.

DTADIR (Data direction)

The direction of data traffic to be traced. The direction can be outbound traffic only (*SND), inbound traffic only (*RCV), or both directions (*BOTH).

TRCFULL (Trace full)

What occurs when the trace buffer is full. This parameter has two possible values. The default value is *WRAP, which means, when the trace buffer is full, the trace wraps to the beginning. The oldest trace records are written over by new ones as they are collected.

The second value *STOPTRC let the trace stop when the trace buffer, specified in the MAXSTG parameter is full of trace records. As general rule, always define the buffer size to be large enough to store all the trace records. If the trace wraps, you might lose important trace information. If you experience a highly intermittent problem, define the trace buffer to be large enough that a wrap of the buffer will not discard any important information.

USRDTA (Number of user bytes to trace)

Defines the number of data to be traced in the user data part of the data frames. By default only the first 100 bytes of user data are captured for LAN interfaces. For all other interfaces, all user data is captured. Make sure you specify *MAX if you suspect problems in the user data of a frame.

TEXT (Trace description)

Provides a meaningful description of the trace.

Stopping the communications trace

If you do not otherwise specify, the trace typically stops as soon as the condition for which you are tracing occurs. Use the End Communications Trace (ENDCMNTRC) command to stop the trace. The following command is an example of the ENDCMNTRC command:

```
ENDCMNTRC CFGOBJ(TRNLINE) CFGTYPE(*LIN)
```

The command has two parameters:

CFGOBJ (Configuration object)

The name of the configuration object for which the trace is running. The object is either a line description, a network interface description, or a network server description.

CFGTYPE (Configuration type)

Whether a line (*LIN), a network interface (*NWI), or a network server (*NWS) is being traced.

Printing the trace data

After you stop the communications trace, you need to print the trace data. Use the Print Communications Trace (PRTCMNTRC) command to perform this task. Since all line traffic is captured during the trace period, you have multiple filter options for output generation. Try to keep the spooled file as small as possible. This makes the analysis faster and more efficient. In the case of a VPN problem, filter on IP traffic only and, if possible, on a specific IP address. You also have the option of filtering on a specific IP port number. The following is an example of the PRTCMNTRC command:

```
PRTCMNTRC CFGOBJ(TRNLINE) CFGTYPE(*LIN) FMTTCP(*YES) TCPIPADR('10.50.21.1)
          SLTPORT(500) FMTBCD(*NO)
```

In this example, the trace is formatted for IP traffic and contains only data for the IP address, where the source or destination address is 10.50.21.1 and the source or destination IP port number is 500.

Only the most important command parameters for analyzing VPN problems, are explained below:

CFGOBJ (Configuration object)

The name of the configuration object for which the trace is running. The object is either a line description, a network interface description, or a network server description.

CFGTYPE (Configuration type)

Whether a line (*LIN), a network interface (*NWI), or a network server (*NWS) is being traced.

FMTTCP (Format TCP/IP data)

Whether to format the trace for TCP/IP and UDP/IP data. Specify *YES to format the trace for IP data.

TCPIPADR (Format TCP/IP data by address)

This parameter consists of two elements. If you specify IP addresses on both elements, only IP traffic between those addresses will print.

SLTPORT (IP port number)

The IP port number to filter.

FMTBCD (Format broadcast data)

Whether all broadcast frames are printed. Yes is the default. If you do not want; for example, Address Resolution Protocol (ARP) requests, specify *NO; otherwise you might be overwhelmed with broadcast messages.

Related tasks:

“Getting started with VPN troubleshooting” on page 63

Complete this task to learn the various methods for determine any VPN problems you are having on your system.

Related information for VPN

IBM Redbooks publications and Web sites contain information that relates to the Virtual private networking topic collection. You can view or print any of the PDF files.

IBM Redbooks®

- IBM System i® Security Guide for IBM i5/OS Version 5 Release 4 
- AS/400 Internet Security: Implementing AS/400 Virtual Private Networks
- AS/400 Internet Security Scenarios: A Practical Approach® 
- OS/400® V5R2 Virtual Private Networks: Remote Access to the IBM eServer™ iSeries Server with Windows 2000 VPN Clients 

Web sites

- | • TCP/IP on IBM i: RFC Documents 

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department YBWA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Programming interface information

This Virtual private networking publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Product Number: 5770-SS1

Printed in USA