
Virtual Private Networks

One of the principal features of routers is their support of virtual private networks (VPNs). This document discusses transmission security, VPNs, and how the EN-1000 sets up and uses a VPN connection.

Note: To set up VPN connections, see the following documents:

- [Configuring IPsec VPNs in the EN-1000™](#)
- [Configuring the EN-1000's VPN Firewall](#)
- [Starting and Tracking VPNs in the EN-1000](#)

A VPN is a secure encrypted transmission between two or more private endpoints over a public network. Tunneling—encapsulating data within secure packets—isolates the private data from other traffic carried by the public network, providing secure transport across the network. The public network uses the header information in the packets to deliver the packets to their destination. When the destination endpoint receives the packets, it authenticates and unpackages them, and decrypts the data.

Use of VPNs allows for dynamic, temporary connections instead of permanent physical connections. This allows an organization to build a private network over the public IP network, which reduces costs for the number of leased lines that the organization needs to maintain. In addition, connection (via VPN client software) over the internet allows business travelers to communicate with the office network from any site that has a connection to the internet.

The EN-1000 supports VPN's IP Security Protocol (IPsec, described in RFC 2401) and supports IPsec's use of the Internet Key Exchange, version 1 (IKEv1) and version 2 (IKEv2). Any EN-1000 device can use IKEv1 and IKEv2 at the same time, but not on the same port at the same time.

See the following sections:

- [Basics of Virtual Private Networks](#)
- [Developing a Virtual Private Network](#)

5.1 Basics of Virtual Private Networks

This section discusses basic principles and operations in Virtual Private Networks.

A VPN device encapsulates information into IP packets, and can perform as a VPN gateway over public networks that use IP. As a VPN gateway, a VPN device can perform IPsec tunnel initiation, IPsec tunnel termination, and IPsec passthrough. Those processes use IPsec for VPN security, performing the functions listed in [Table 5-1](#).

Table 5-1. IPsec Components Used in the EN-1000

Function	Protocols	Acronym	Standard ^a
Key Exchange	Internet Key Exchange	IKE	Version 1: RFC 2409 Version 2: RFC 5996
	Internet Security Association and Key Management Protocol	ISAKMP	RFC 2408
Encryption	Data Encryption Standard	DES	FIPS PUB 46-2
	Triple Data Encryption Standard	3DES	SP 800-67, Revision 1 (per FIPS PUB 140-2)
	Advanced Encryption Standard	AES	FIPS PUB 197
Security Protocols	Encapsulating Security Payload	ESP	RFC 2406
	Authentication Header	AH	RFC 2402
Authentication	Hashed Message Authentication Code: Message Digest 5	HMAC MD5	RFC 1321; For use of MD5 within ESP and AH: RFC 2403
	Hashed Message Authentication Code: Secure Hash Algorithm 1	HMAC SHA-1	RFC 2404
	Hashed Message Authentication Code: Secure Hash Algorithm 3	HMAC SHA-3	FIPS PUB 180-4

a. Each Request for Comments (RFC) is from the Internet Engineering Task Force (IETF). Each Federal Information Processing Standard Publication (FIPS PUB) and each Special Publication (SP) is from the National Institute of Standards and Technology (NIST).

The EN-1000 can implement IKE version 1 (IKEv1) or version 2 (IKEv2) VPN tunnels with any other IPsec-compliant VPN gateway or VPN client. The EN-2000 supports the following tunnel modes:

- **Tunnel initiation:** The device receives packets from a local user terminal. The device encapsulates the packets according to the IPsec user policy, establishes a VPN tunnel across the public network to a remote VPN gateway, and sends the packets across the VPN tunnel toward their destination.
- **Tunnel passthrough:** The device receives IPsec-encapsulated packets from a client VPN terminal, and provides transparent forwarding of the IP packets according to the IPsec user policy. The device sends the packets across the public network without repackaging them.
- **Tunnel termination:** The device terminates (accepts) an IPsec tunnel initiated by a remote VPN gateway or VPN client across the public network. The device authenticates and unpackages the tunnel's packets, and delivers

them to the destination terminal. (To perform tunnel termination, the device must maintain a table of VPN users that function as prospective tunnel initiators; see [The IP Policy Table](#).)

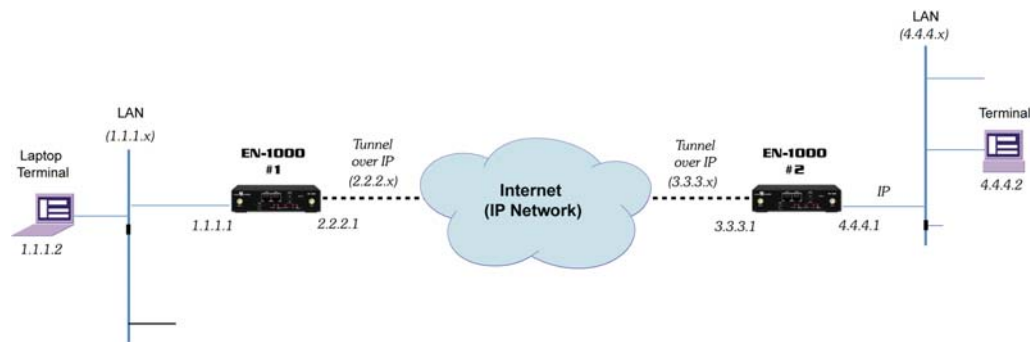
See the following:

- [A Simple Virtual Private Network](#)
- [Tunnel Modes](#)
- [Tunnel Support](#)
- [Internet Key Exchange](#)

5.1.1 A Simple Virtual Private Network

[Figure 5-1](#) illustrates two EN-1000s functioning as VPN gateways over the IP network.


Figure 5-1. EN-1000s as VPN Gateways



[Figure 5-2](#) shows a simplified example of the EN-1000's encryption and encapsulation of data.

Note: The transmission shown in [Figure 5-2](#) originates from the laptop terminal (IP address 1.1.1.2) shown in [Figure 5-1](#), and is destined for the desktop terminal (IP address 4.4.4.2) in [Figure 5-1](#).

Figure 5-2. Sample IPsec Encryption and Encapsulation



5.1.2 Tunnel Modes

The EN-1000 supports the following modes for VPN tunnels:

- [Tunnel Initiation](#)
- [Tunnel Termination](#)
- [Tunnel Passthrough](#)

5.1.2.1 Tunnel Initiation

The EN-1000 can initiate a VPN tunnel, using IKEv1 or IKEv2, to any other IPsec-compliant VPN gateway. When a local user originates packets to the EN-1000, and the packets need to travel over a VPN tunnel, the EN-1000 searches its database for an appropriate VPN policy and VPN profile.

When an appropriate VPN policy and VPN profile have been determined, the EN-1000 contacts the remote VPN gateway specified by the profile, and negotiates a security association (SA). When the gateways agree on an SA and set up a VPN tunnel, the EN-1000 encapsulates the packets according to the policy, and sends them across the public network. When the remote VPN gateway receives the packets, it forwards them to the remote destination.

Note: In order to use a VPN tunnel, the combination of origin and destination must conform to a VPN policy. Otherwise, the request will be rejected.

5.1.2.2 Tunnel Termination

When a remote VPN gateway initiates a VPN tunnel, the EN-1000 acts as a tunnel terminator. The EN-1000 looks for matches against the following items:

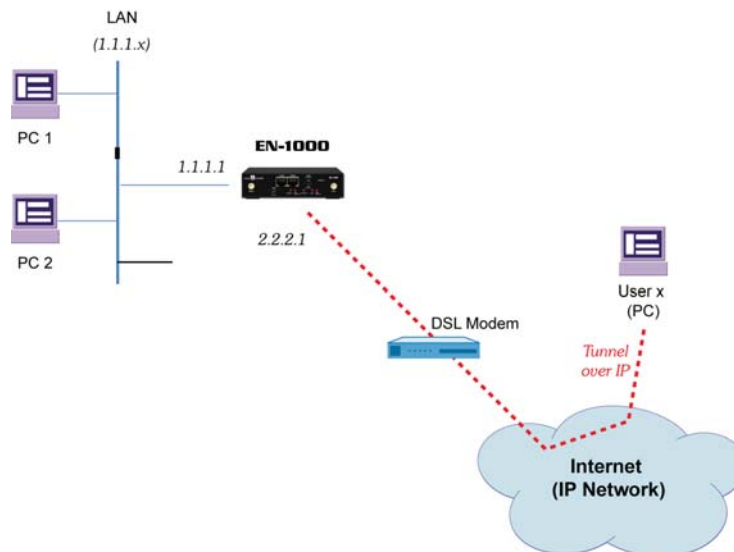
- IDs
- Preshared key
- Peer (remote) user ID. This can be a group ID or a single ID.

If the values match a VPN policy record, the EN-1000 accepts the tunnel termination. Then the EN-1000 negotiates the key, and accepts or rejects the proposals presented by the initiating VPN gateway. If the remote user's ID matches a record in the EN-1000's database, the EN-1000 agrees to terminate the tunnel.

In [Figure 5-3](#), a remote VPN user (User x) wishes to communicate with PC 1, so User x initiates a VPN tunnel to the EN-1000's external IP address. Because the remote user's ID matches a record in the EN-1000's database, the EN-1000 agrees to terminate the VPN tunnel. ([Figure 5-3](#) shows this VPN tunnel as a red dotted line.)

Because the remote user wishes to communicate with PC 1, the EN-1000 opens the VPN packets, decrypts them, and passes the information to PC 1.

Figure 5-3. EN-1000 Terminating Tunnel from VPN Client



[Table 5-2](#) lists sample parameters to support tunnel termination for a remote VPN user.

Table 5-2. Sample Remote User Record

Field	Sample Value
Peer ID (Remote User ID)	a1234@encorenetworks.com
Preshared Key	*****
Profile Group	1,2,4,5 Note: The choices for profile group can include up to four VPN profiles. The EN-1000 chooses the first profile that the peer ID matches. One of the profile-group choices can be a wildcard—any profile listed in the VPN profile database. You may list VPN profiles before a wildcard, but there is no need to list any profiles after a wildcard.
Certificate	*****

Note: The remote user's IP address does not need to be known in advance.

5.1.2.3 Tunnel Passthrough

Tunnel passthrough is used when a remote or local VPN user sends IPsec-encapsulated packets to the EN-1000. In passthrough mode, the originating device sets up the VPN tunnel, so the originating device is the VPN endpoint, and the EN-1000 is merely a gateway. The EN-1000 gateway uses its VPN policy to provide transparent forwarding of the VPN endpoint's IP packets.

Tunnel passthrough occurs most often when packets are received from a VPN client. If a remote user is using VPN client software, the client sets up a VPN tunnel through the EN-1000 to a remote network. In this case, the EN-1000 does not initiate a new VPN tunnel; it uses passthrough mode to carry the tunnel created by the VPN client (which is the VPN endpoint in this scenario).

5.1.3 Tunnel Support

In some situations, a single VPN tunnel can support more than one user. See the following:

- [Tunnel Sharing](#)
- [Tunnel Switching](#)
- [Split Tunneling](#)

5.1.3.1 Tunnel Sharing

More than one VPN profile can specify the same local and remote VPN gateways to reach its remote endpoint. If two such profiles are active at the same time, they are using the same tunnel between the gateways for their VPN connections to different endpoints. This is called tunnel sharing (or tunnel multiplexing).

5.1.3.2 Tunnel Switching


A remote endpoint can initiate a VPN tunnel into the network. If the remote endpoint wishes to communicate with a destination endpoint that is outside the network, the EN-1000 checks to see whether there is a VPN profile describing a tunnel to the requested destination. If so, the EN-1000 initiates a VPN tunnel to that destination, and routes the traffic from the initiating endpoint to the destination. This is called tunnel switching.

In [Figure 5-4](#) (an expansion of [Figure 5-3](#)), a remote VPN user (User x) wishes to communicate with the network's VPN host, but does not want to connect a VPN tunnel to the host itself. So User x initiates a tunnel to the EN-1000's external IP address. Because the remote user's ID matches a record in the EN-1000's database, the EN-1000 agrees to terminate the tunnel. ([Figure 5-4](#) shows this VPN tunnel as a red dotted line.)

The remote user wishes to communicate with the VPN host, so the EN-1000 accepts and decrypts the client's VPN packets. Then the EN-1000 initiates a tunnel to the VPN host, and passes the client's packets to the VPN host. ([Figure 5-4](#) shows this tunnel as a black dotted line.)

Tunnel switching also hides the VPN client's IP address.

Figure 5-4. EN-1000 Tunnel Switching between VPN Client and VPN Host



5.1.3.3 Split Tunneling

The EN-1000 can route VPN traffic through VPN tunnels and can route non-VPN traffic outside these tunnels, on the same line. The EN-1000 uses the IP policy table to determine whether to route traffic through or outside a VPN tunnel.

The EN-1000 performs split tunneling automatically and dynamically.

5.1.4 Internet Key Exchange

The EN-1000 uses the Internet Key Exchange (IKE) protocol to provide secure transmission between VPN endpoints. IKE negotiates security associations (SAs) and provides authenticated keys for these SAs. (A security association is a set of policies that establish a protected, authenticated connection for data transmission.) IKE can be used to do the following:

- Set up virtual private networks (VPNs).
- Provide a remote user secure access to a network. (The remote user's IP address does not need to be known in advance.)
- Negotiate SAs (and hide identities) for VPN client endpoints.

There are some differences between *IKE Version 1* (IKEv1) and *IKE Version 2* (IKEv2). The EN-1000 supports use of IKEv1 and IKEv2 at the same time, but not on the same port at the same time.

Note: All IKEv2 VPN connections work well across a device that performs network address translation (NAT). However, an IKEv1 VPN connection must use an appropriate path out of the LAN:

- When an IKEv1 VPN connection uses the Encapsulating Security Payload (ESP) protocol, the connection can cross a device that performs NAT.
- When an IKEv1 VPN connection uses the Authentication Header (AH) protocol, the connection must not cross a device that performs NAT.

5.1.4.1 Perfect Forward Secrecy

Perfect forward secrecy (PFS)—the use of uniquely derived keys to establish security associations—is an important feature of the IKE protocol. PFS comprises the following principles:

- Discovery of a key endangers only transmissions protected by that key; so
- Material used to derive one key cannot be used to derive additional keys; and
- No key can be used to derive another key.

5.1.4.2 IKE Version 1

The Internet Key Exchange protocol, version 1, has two phases:

- Phase 1 is used for key exchange. In this phase, IKE negotiates the following items to establish a Security Association for Phase 2:
 - The encryption algorithm
 - The hash algorithm
 - The authentication method
 - The Diffie–Hellman group
- Phase 2 negotiates an SA for services in the transmission. Then this phase is used for data transmission.

5.1.4.2.1 Details of IKE Version 1

IKEv1 maintains perfect forward secrecy in the way it performs the following:

- IKE uses a Diffie-Hellman (DH) exchange to set up phase 1. (A DH exchange protects the identities of the originator and the recipient.) Phase 1 can use main mode or aggressive mode (but not both).

Phase 1 establishes an SA for phase 2, as follows:

- The originator presents proposals for the SA. (The originator may send an unlimited number of proposals; the recipient can limit the number it will consider.)
 - The recipient chooses one proposal and sends its response. The recipient cannot change the proposal. If the originator notices that the proposal has changed in any way, the originator refuses the response.
 - When the originator accepts the response, the SA is set up for phase 2.
- In phase 2, IKE establishes an SA for data transmission. as follows:
 - Phase 2 negotiates for services that will be used, such as IPsec.
 - When the phase 2 SA is ready for data transmission, IKE deletes the SA that phase 1 had established.
 - In the SA for data transmission, quick mode is used for transmission. Both sides of the connection can transmit data.

Instead of extensive authentication, which consumes time and CPU resources, the SA now uses cookies for authentication. The cookie order established in phase 1 (originator vs. recipient) is always used; the cookies do not change order when the transmission direction changes.

Note: Each IKEv1 phase has a fixed lifetime. The lifetime can be defined in units of time, number of transmissions, or total amount of transmission (in kilobytes). A phase's lifetime cannot be increased after the phase has started.

5.1.4.3 IKE Version 2

The EN-1000 implements the Internet Key Exchange protocol, version 2, in conformance with IETF RFC 5996.

IKEv2 simplifies key exchanges:

- IKEv2 does not use IKEv1's main mode, aggressive mode, or quick mode. IKEv2 uses a single standard mode.
- Negotiation for set-up has reduced from a minimum of nine exchanges in IKEv1 to a minimum of four exchanges in IKEv2.
- IKEv2 uses only four types of messages, regardless of the number of exchanges.

There are some other principal changes for IKEv2:

- VPN policies are simpler in IKEv2 than in IKEv1.
- IKEv2 uses less bandwidth than IKEv1 uses.

- IKEv2 has built-in NAT traversal; IKEv1 must include ESP and exclude AH to traverse NAT.
- In IKEv2, authentication has been separated from IP policy.
- IKEv2 includes asymmetric authentication.
- IKEv2 supports authentication with EAP; IKEv1 does not support EAP.
- IKEv2 can use MOBIKE to support a traveling connection; IKEv1 does not support MOBIKE.
- IKEv2 detects whether a tunnel is live; IKEv1 does not have that capability.
- IKEv2 does not process a request until after it has determined the identity of the requester (as shown in the packet exchanges between Step 1b on page 12 and Step 2a on page 13, and as shown in the exchanges between Step 2a and Step 2b on page 13). That determination of identity reduces spoofing (and thus reduces denial-of-service attacks).

5.1.4.3.1 Extensible Authentication Protocol

The Extensible Authentication Protocol (EAP) defines message formats used in IKEv2. EAP for IKEv2 is described in IETF RFC 5106.

Note: EAP is a format protocol; it is not a transmission protocol. IKEv2 defines the way that it transmits the EAP format.

5.1.4.3.1.1 EAP Authentication

IKEv2 uses EAP in providing authentication and establishing keys. EAP supports authentication for the following:

- Password: A simple character string, shared by both sides of the connection. Passwords are usually generated by humans, sometimes via algorithms, sometimes not.
- Symmetric key: A complex character string, shared by both sides of the connection. Symmetric keys are usually generated by computer-driven algorithms.
- Asymmetric key pair: A complex set of keys. The public key is available to anyone; the private key is known only to the key owner (generally the recipient in any single exchange). An asymmetric key pair is generated by the key owner, using a hash-encryption technique.

5.1.4.3.1.2 EAP Exchanges

An exchange pair may use a different authentication mechanism in each direction. [Table 5-3](#) lists the expected combinations for IKEv2 EAP key exchanges.

Table 5-3. Standard EAP Combinations for IKEv2 Authentication

IKEv2 Initiator Sends:	IKEv2 Responder Sends:
one part of asymmetric key pair	counterpart in asymmetric key pair, or password, or symmetric key
symmetric key	symmetric key

5.1.4.3.2 MOBIKE

The EN-1000 includes support for the IKEv2 Mobility and Multihoming Protocol (MOBIKE, described in IETF RFC 4555). MOBIKE permits IP addresses to change dynamically—for example, as a mobile client moves during an IKEv2 VPN connection.

When one side of the VPN connection moves, the endpoints use INFORMATIONAL exchange pairs to indicate a change in that side's IP address, without interrupting the connection. MOBIKE supports this by updating the information for the tunnel headers, and leaves the information inside the tunnel unchanged.

The parties in the connection usually do not experience any interruption in service—and, if there is an interruption, the connection is re-established quickly; the call is not dropped.

- ! **Caution:** MOBIKE does not currently support simultaneous movement of both parties in the connection. It is advised that only one party in the VPN connection be mobile, and that the other party remain in a fixed location (such as an office).

The EN-1000 does not travel, so MOBIKE is used only for the remote device in the connection.

5.1.4.3.3 Sample IKEv2 Exchanges

This section provides examples of basic exchanges in IKEv2.

Note: Each request requires a response; so exchanges are always in pairs.

See the following:

- [Overview of IKEv2 Exchanges](#)
- [Details of IKEv2 Exchanges](#)

5.1.4.3.3.1 Overview of IKEv2 Exchanges

Phase 1 does the following:

- 1 The first exchange pair (IKE_SA_INIT) sets up the Security Association (IKE_SA).
- 2 The second exchange pair (IKE_AUTH) sets up the following:
 - The authentication method
 - The Security Association for the next phase (CHILD_SA, for data exchange)

Phase 2 (CHILD_SA) does the following:

- 3 The VPN endpoints generate security keys.
- 4 The VPN endpoints exchange one or more INFORMATIONAL pairs.

Note: Phase 2 can establish a new CHILD_SA whenever a new Security Association is needed.
- 5 When the communication has ended, the VPN endpoints close the VPN tunnel.

5.1.4.3.3.2 Details of IKEv2 Exchanges

Details of IKEv2 Phase 1:

- 1 The first exchange pair (IKE_SA_INIT) negotiates security parameters for the IKE_SA, sets Diffie–Hellman values, and sends generated nonces, as shown in [Substep a](#) and [Substep b](#).

a The Initiator sends the following information:

Header ¹	Initiator's 1st Security Association ²	Initiator's Key Exchange ³	Initiator's Nonce ^{4,5}
---------------------	---	---------------------------------------	----------------------------------

1. Including Security Parameter Indexes, version numbers, and flags
2. Including supported cryptographic algorithms offered for the security association
3. Including the Initiator's Diffie–Hellman value
4. A nonce provides additional security. It is randomly generated for a single transmission. If a nonce is repeated in a subsequent transmission, that transmission is rejected. Nonces might also be used to derive some keys.
5. We advise the use of separate random generators for keys and nonces, to avoid the possibility that a nonce might compromise a key.

b The Responder replies with the following information and request:

Header ¹	Responder's 1st Security Association ²	Responder's Key Exchange ³	Responder's Nonce ^{4,5}	[OPTIONAL] Request for Certificates
---------------------	---	---------------------------------------	----------------------------------	-------------------------------------

1. Including Security Parameter Indexes, version numbers, and flags
2. Including the cryptographic algorithm to use (chosen from those offered by the Initiator)
3. Completing the Diffie–Hellman exchange
4. A nonce provides additional security. It is randomly generated for a single transmission. If a nonce is repeated in a subsequent transmission, that transmission is rejected. Nonces might also be used to derive some keys.
5. We advise the use of separate random generators for keys and nonces, to avoid the possibility that a nonce might compromise a key.

2 The second exchange pair (IKE_AUTH) transmits identities, demonstrates knowledge of secrets corresponding to identities, and establishes the security association (SA, using AH or ESP) for the first CHILD_SA, as shown in [Substep a](#) and [Substep b](#).

a The Initiator replies to the Responder's first exchange, and sends the following information and request:

Header ¹	Initiator's Identity	[OPTIONAL] ² One or more Certificates ³	[OPTIONAL] Request for Certificates
---------------------	----------------------	--	--

[OPTIONAL] Responder's Identity ⁴	Authentication	Initiator's 2nd Security Association ⁵	Initiator's Traffic Selection	Responder's Traffic Selection
---	----------------	---	-------------------------------	-------------------------------

1. Including Security Parameter Indexes, version numbers, and flags
2. If the Responder requested certificates, they must be provided.
3. The first certificate must include the public key used to verify the Authentication in this packet.
4. If the Responder supports more than one identity, this item selects one of those identities for this connection.
5. Beginning negotiation for the CHILD_SA

b The Responder replies with the following information:

Header ¹	Responder's Identity	[OPTIONAL] ² One or more Certificates ³	Authentication
---------------------	----------------------	--	----------------

Responder's 2nd Security Association ⁴	Initiator's Traffic Selection	Responder's Traffic Selection
---	-------------------------------	-------------------------------

1. Including Security Parameter Indexes, version numbers, and flags
2. If the Initiator requested certificates, they must be provided.
3. The first certificate must include the public key used to verify the Authentication in this packet.
4. Concluding negotiation for the CHILD_SA

❖ The CHILD_SA is established. Data can now be exchanged.

Note: If establishment of the CHILD_SA fails, the exchange establishes a generic IKE_SA.

Details of IKEv2 Phase 2:

- 3 After the CHILD_SA has been established, each side generates a security key seed (SKEYSEED).

Note: The SKEYSEED is used to derive encryption and authentication keys to protect the established (current) security association. In addition, the SKEYSEED might be used to develop keys for subsequent CHILD_SAs in this connection.

- 4 One or more INFORMATIONAL pairs are exchanged.

Note: Subsequent exchange pairs in this IKEv2 connection can be:

- INFORMATIONAL: to exchange data, to report a change of IP address (as in [MOBIKE](#)), to report errors, to delete the Security Association, or to perform other administrative tasks.

Either side can initiate an INFORMATIONAL exchange.

- CREATE_CHILD_SA: to create a new CHILD_SA.

Either side can initiate an exchange to CREATE_CHILD_SA.

- 5 When all INFORMATIONAL transactions have been completed, either side initiates an INFORMATIONAL pair to end the VPN connection.

❖ The VPN tunnel closes.

5.2 Developing a Virtual Private Network

[Table 5-4](#) lists the information needed to set up a VPN tunnel.

Table 5-4. Information Required to Configure the EN-1000 for VPNs

Item	Central Site	Remote Site
WAN IP Address	Usually a public IP address supplied by your internet service provider, e.g., 68.x.x.34	Supplied by the service provider, e.g., 65.x.x.72 . (If the WAN IP address is issued by the remote server, select Dynamic .)
WAN Subnet Mask	Subnet mask for the address above, e.g., 255.255.255.252	Subnet mask for the address above, e.g., 255.255.255.240 (not applicable if WAN IP address is dynamic)
WAN Default Router (a.k.a. Default Gateway)	The next hop router for the WAN IP address, e.g., 68.x.x.33	IP address of the remote modem, e.g., 65.x.x.65 (not applicable if WAN IP address is dynamic)
VPN Gateway IP Address or DNS Name (required only for end that initiates tunnel)	Required only if the central site is the initiator. This will be the WAN IP address of the remote site unit or the DNS name of the remote site unit—e.g., 65.x.x.72 or en1000@remote.com .	Required only if the remote site is the initiator. This will be the WAN IP address of the central site unit or the DNS name of the central site unit—e.g., 68.x.x.34 or en1000@central.com .
LAN IP Address	Fixed address on the LAN segment to be assigned to the router LAN port, e.g., 10.10.10.1	Fixed address on the LAN segment to be assigned to the router LAN port, e.g., 192.168.1.1
LAN Subnet Mask	Subnet mask for the address above, e.g., 255.255.255.0	Subnet mask for the address above, e.g., 255.255.255.0
DHCP IP Address Pool (Range, from low to high)	If the router is to issue IP addresses via DHCP on the LAN side, enter the address range here, e.g., 10.10.10.2 to 10.10.10.24 .	If the router is to issue IP addresses via DHCP on the LAN side, enter the address range here, e.g., 192.168.1.2 to 192.168.1.10 .
<i>Additional Security Information Required . . .</i>	<i>. . . When Running IPsec Encryption</i>	<i>. . . If the EN-1000 is to Provide DNS Information</i>
User ID	Must be the same at both ends, e.g., user1@site2 or test123	Primary DNS address
Preshared Key	Maximum of 18 characters. Must be the same at both ends, e.g., e2we36TJK@s8h12Q	Secondary DNS address

Note: To set up IP routes and policies, see the document [Configuring the EN-1000's VPN Firewall](#).

5.2.1 VPN Configuration Plan

The following tables provide an example of planning a configuration for your virtual private network users.

5.2.1.1 The IP Policy Table

IP Policy Tables are used to establish processes and types of connections.

[Table 5-5](#) shows a sample IP Policy Table.

Each policy includes the VPN profile that the connection must use; the user must also be authorized to use the specified profile. Your organization's IP Policy Table may include additional fields.

The EN-1000's IP Policy Table is set up in the document [Configuring the EN-1000's VPN Firewall](#).

Note: In [Table 5-5](#), the column for **Record 2** provides an example of tunnel termination: If a record's **Direction** is "incoming," then the record's **Source IP Addresses** (in the range from **Low** to **High**) indicate one or more remote devices. If the **Action** is "tunnel termination," a device with an IP address in the source range can initiate a tunnel that the local device will accept.)

The IP Policy Table must include a field naming the profile used in the policy. In [Table 5-5](#), this is the field **VPN Profile Used**. The value in this field cross-references the profile's configuration (in a VPN Profile Table).

Table 5-5. Sample IP Policy Table

Field	Value for Record 1	Value for Record 2	Records 3, 4, 5, . . .
Low IP Address for Source	1.1.1.1	4.4.4.1	. . .
High IP Address for Source	1.1.1.255	4.4.4.255	. . .
Low IP Address for Destination	4.4.4.1	1.1.1.1	. . .
High IP Address for Destination	4.4.4.255	1.1.1.255	. . .
Global Path	LAN	LAN	. . .
Direction	Outgoing	Incoming	. . .
Action	Tunnel Initiation	Tunnel Termination	. . .
Description	Tunnel A	Terminate P27	. . .
VPN Profile Used	Profile 1	Profile 7	. . .

5.2.1.2 The VPN Profile Table

[Table 5-6](#) shows a sample VPN profile table, with the field **VPN Profile Name** cross-referenced against profiles listed in the IP Policy Table. (Your VPN Profile Table may show additional fields.)

Table 5-6. Sample VPN Profile Table

Field ^a	Value for Record 1	Value for Record 2	Records 3, 4, 5, . . .
VPN Profile Name	Profile 1	John's VPN Connection	. . .
Local ID (User ID)	1.2.1.12	Set_1@encore-net-works.com	. . .
Remote VPN Gateway Address	3.43.3.12	3.43.3.12	. . .
Keying	Manual ^b	Auto-Key	. . .
Security Protocol	ESP	----	. . .
Local SPI	1ffff	----	. . .
Remote SPI	1000	----	. . .
Authentication Mode	----	Main mode, Aggressive mode	. . .
Authentication Protocol	HMAC-SHA1	----	. . .
Authentication Key	48454C4C4F000000 0000000000000000	----	. . .
Preshared key	----	*****	. . .
Encryption	3DES	----	. . .
Encryption Key	48454C4C4F000000	----	. . .
Phase 1, Proposal 1 ^c	----	PRE-G2-DES-MD5	. . .
Phase 1, Proposal 2 ^c	----	VSA-G2-3DES-SHA	. . .
Phase 2, Proposal 1 ^c	----	STD-G2-3DES-MD5	. . .
Phase 2, Proposal 2 ^c	----	PFS-G2-3DES-SHA	. . .
Replay Protection	----	enabled	. . .
User ID Verification	----	enabled	. . .
Password Verification	----	disabled	. . .
Timeout	----	30 minutes	. . .

a. A VPN Profile Table includes all records. When the user specifies the type of keying the profile will use, the EN-1000 presents for configuration only the fields that apply to the specified keying. ([Table 5-7](#) presents parameters for autokeying.)

b. The EN-1000 does not use manual keying in normal operation. The EN-1000 normally performs only automatic keying. If you wish to use manual keying, contact your Encore Networks representative.

c. Used only in IKEv1.

5.2.2 Automatic Keying

In automatic keying (autokeying), keys are dynamic, always changing. Special keys are exchanged at the beginning of the connection, and the VPN gateways negotiate other keys for the connection. If desired, keys can time out and new keys can be negotiated for subsequent parts of the connection.

The EN-1000 uses the Internet Key Exchange (IKE) protocol for automatic generation of keys in VPN connections. When an EN-1000 uses the automatic keying feature, an IKE tunnel is set up for key exchange. That IKE tunnel sets up keys for a subsequent data tunnel (if a subsequent tunnel is needed). The data tunnel is used for data exchange. See [Section 5.1.4, Internet Key Exchange](#).

[Table 5-7](#) shows sample parameters to set up automatic keying for a VPN connection.

Table 5-7. Sample VPN Profile, Automatic Keying

Sample Fields	Sample Values
Authentication Mode	Main mode (also known as ID Protection), Aggressive mode
Local ID (User ID) ^a	1.1.1.1
Remote Gateway IP Address ^b	3.3.3.1
Preshared Key ^c	*****
Phase 1, Proposal 1 ^{d, e}	PRE-G2-DES-MD5
Phase 1, Proposal 2 ^{d, e}	VSA-G2-3DES-SHA
Phase 2, Proposal 1 ^{d, e}	STD-G2-3DES-MD5
Phase 2, Proposal 2 ^{d, e}	PFS-G2-3DES-SHA
Replay Protection	Enable/Disable

a. There are three formats for the local ID:

- E-mail format: ascii-format@ascii-format
- IP address format: x.x.x.x
- Perfect domain name format: hostdomain.net

b. There are two kinds of remote IP addresses: static and dynamic.

c. The preshared key is used to establish the IKE tunnel. This preshared key must be protected as a super-password. The preshared key uses Diffie–Hellman Exchange 2 (DH2).

d. Used only in IKEv1.

e. The initiator may provide up to four proposals per phase. The recipient must choose at least one proposal for each phase.

[Table 5-8](#) and [Table 5-9](#) illustrate sample proposal combinations for IKEv1 phase 1 and phase 2, respectively.

Table 5-8. Sample IKEv1 Phase 1 Proposal

Sample Fields	Sample Values ^a
Authentication mode	preshared
Diffie–Hellman (DH) group	group 2
Encryption	DES, 3DES
Authentication	HMAC-MD5, HMAC-SHA1
Lifetime ^b	1–100 units
Lifetime units ^b	seconds, minutes, hours, days

a. This sample proposal is tunnel-specific, not session-specific.

b. When the lifetime is reached for the indicated unit, a new key is exchanged.

Table 5-9. Sample IKEv1 Phase 2 Proposal

Sample Fields	Sample Values
Perfect forward secrecy (PFS)	none DH2 (Diffie–Hellman 2)
Security protocol	ESP AH
Encryption	3DES DES
Authentication	HMAC-MD5 HMAC-SHA1
Lifetime ^a	1–100 units
Lifetime unit ^a	number of seconds number of minutes number of hours number of days kilobytes of data sent through the tunnel

a. When the lifetime is reached for the unit indicated, a new key is exchanged.

5.2.3 Sample Configuration for a Remote User

In [Figure 5-3](#), one connection showed a VPN remote user tunneling to an EN-1000 VPN gateway. The EN-1000, in turn, created a tunnel to a VPN host at another site. [Table 5-10](#) lists a sample set of values for the connection between the EN-1000 and the remote user.

Table 5-10. Sample Tunnel User Table

Fields	Values
Profile Name	profile 2
Authentication Mode	aggressive
Keying	auto-IKE
Local User ID	1.1.1.1
Gateway	3.3.3.1
Preshared Key	*****
Phase 1, Proposal 1 ^a	PRE-G2-DES-MD5
Phase 1, Proposal 2 ^a	VSA-G2-3DES-SHA
Phase 2, Proposal 1 ^a	STD-G2-3DES-MD5
Phase 2, Proposal 2 ^a	PFS-G2-3DES-SHA
Replay Protection	enable

a. Used only in IKEv1.