Active Directory in Networks Segmented by Firewalls

Microsoft Corporation

Published: July 2002

Updated: October 2004

Abstract

Microsoft® Active Directory® service domain controllers are increasingly being deployed on networks segmented by firewalls. Three common scenarios are: (1) domain controllers separated from clients in a perimeter network (also known as DMZ, demilitarized zone, and screened subnet), (2) domain controllers in a perimeter network separated from other domain controllers on the network, and (3) networks divided into segments, each containing clients and domain controllers. This white paper describes best practices for deploying domain controllers in segmented networks in a manner that supports client authentication, secure resource access by clients, and replication traffic between domain controllers on opposite sides of a firewall. This paper also provides detailed procedures for configuring IPSec policies to protect Active Directory traffic between domain controllers on opposite sides of a firewall and recommended practices for managing IPSec policies that are assigned to domain controllers.

Contents

Introduction

Early adoption of the Active Directory® directory service demonstrated customer interest in deploying it in environments in which the domain controllers and domain members are separated by one or more firewalls. This paper describes how to configure the domain controllers and firewalls to enable Active Directory functionality in these scenarios; it describes the functionality that is unavailable in such scenarios; and it identifies security considerations in each scenario. This paper considers the three most common scenarios:

- A domain member server residing in the perimeter network is separated from a domain controller for a domain residing in the corporate environment.
- Two forests deployed on opposite sides of a firewall—one in the perimeter network and one in an internal network with an explicit trust established between the domains.
- A single forest is deployed in an internal network with different portions of the forest separated by firewall(s).

The scenarios and recommendations in this paper apply to systems using Microsoft® Windows® 2000. Information about additional scenarios or features included with Microsoft® Windows Server™ 2003 will be provided in a future paper. Appendix A includes detailed procedures for using IPSec to secure traffic between domain controllers through the firewall. The procedures described in Appendix A can be used in a Windows Server 2003 environment. However, in a Windows Server 2003 environment, there might be small differences in the procedures required to create the IPSec policy.

Note

Active Directory functionality is not supported over a router that has Network Address Translation (NAT) enabled. The configuration recommendations in this paper apply only to non-NAT environments.

Operational Building Blocks

Each network scenario can be broken down into a set of operations that a particular client is trying to achieve. These operations are the building blocks for other network scenarios. This section describes each operation individually; you can use these descriptions to create customized scenarios that are not covered in this paper. For a list of commonly used ports referenced in the following operations, see Appendix C.

User Login and Authentication

A user network logon across a firewall uses the following:

- Microsoft-DS traffic (445/tcp, 445/udp)
- Kerberos authentication protocol (88/tcp, 88/udp)
- Lightweight Directory Access Protocol (LDAP) ping (389/udp)
- Domain Name System (DNS) (53/tcp, 53/udp)

Computer Login and Authentication

A computer logon to a domain controller uses the following:

- Microsoft-DS traffic (445/tcp, 445/udp)
- Kerberos authentication protocol (88/tcp, 88/udp)
- LDAP ping (389/udp)
- DNS (53/tcp, 53/udp)

Establishing an Explicit Trust Between Domains

When establishing a trust between domain controllers in different domains, the domain controllers communicate with each other by means of the following:

- Microsoft-DS traffic (445/tcp, 445/udp)
- LDAP (389/tcp) or 636/tcp if using Secure Sockets Layer (SSL))
- LDAP ping (389/udp)
- Kerberos authentication protocol (88/tcp, 88/udp)
- DNS (53/tcp, 53/udp)

Validating and Authenticating a Trust

Trust validation between two domain controllers in different domains uses the following:

- Microsoft-DS traffic (445/tcp, 445/udp)
- LDAP (389/tcp or 636/tcp if using SSL)
- LDAP ping (389/udp)
- Kerberos (88/tcp, 88/udp)

- DNS (53/tcp, 53/udp)
- Net Logon service

Because the Net Logon service cannot be locked down to a single RPC port, the RPC endpoint mapper (135/tcp and 135/udp) needs to be open, as does a small range of dynamic RPC ports for the mapper to use. For information about how to limit the range of dynamic RPC ports, see Appendix E.

Access File Resource

File access uses SMB over IP (445/tcp, 445/udp).

Perform a DNS Lookup

To perform a DNS lookup across a firewall ports 53/tcp and 53/udp must be open. DNS is used for name resolution and supports other services such as the domain controller locator.

Perform Active Directory Replication

The type of network traffic that is required for replication differs based on whether the replication is between domain controllers of one or more domains. Both types of replication require the following:

- Directory service RPC traffic (configurable directory service RPC port)
- LDAP (389/tcp or 636/tcp if using SSL)
- LDAP ping (389/udp)
- Kerberos (88/tcp, 88/udp)
- DNS (53/tcp, 53/udp)
- SMB over IP traffic (445/tcp, 445/udp)

Replication within a domain also requires File Replication service (FRS) using a dynamic RPC port. Replication traffic and configuration is further described in “Domain Controller Replication Across a Firewall” later in this paper. For instructions for configuring a static directory service RPC port, see Appendix D. For the procedure to limit the range of dynamic RPC ports, see Appendix E.

Common Scenarios

The following section describes the most common customer scenarios in which Active Directory domain controllers and domain members are separated by one or more firewalls. Each scenario requires special configuration of the firewall to enable computers to authenticate (and authenticate to) other resources residing on the other side of the firewall, including one that requires enabling Active Directory replication across the firewall. This section makes appropriate recommendations regarding domain controller and firewall configurations, depending on whether Active Directory replication needs to be enabled.

The following represent some common examples; other, more complex scenarios can be supported by identifying the appropriate operations as described in Operational Building Blocks and configuring the firewalls to allow propagation of network traffic as required by the identified operations.

Member Servers Separated from a Domain Controller

The following scenario has firewalls separating one or more member servers from one or more domain controllers. A common example of this scenario is an application server or a member server that is running Microsoft® Exchange Server and resides in the perimeter network accessing data from an internal resource, such as a file share for scripts, and authenticating to a domain controller within the intranet.

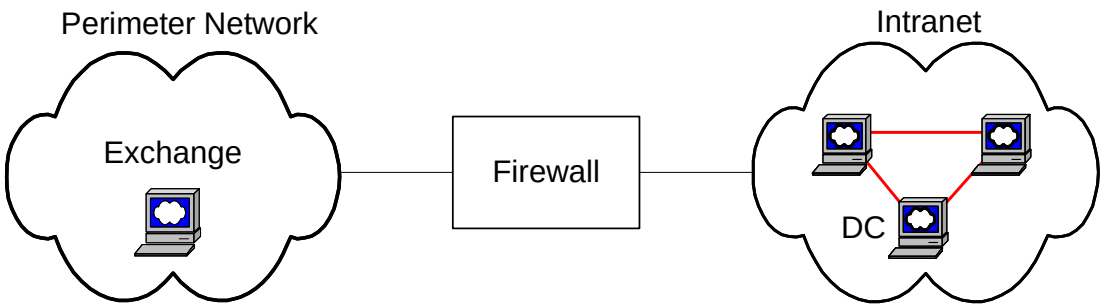

To enable a server located in the perimeter network to access data from an internal resource (such as a Global Catalog) open ports on the firewall and create point-to-point IP restrictions so that only specific computers are allowed to communicate across the firewall.

Open ports that are required by applications that you are using. For example, to enable Microsoft® Exchange 2000 Server to retrieve data from the Global Catalog, open port 3268. You need to obtain the exact list of ports required from your application vendors. For a list of other commonly used ports, see Appendix C.

Open the following ports for authentication traffic:

- Kerberos ports (88/tcp, 88/udp) used to perform mutual authentication between the member server and the domain controller. Kerberos traffic needs to be allowed in addition to the possible application specific traffic.
- DNS ports (53/tcp, 53/udp) used for name lookups.
- LDAP ports (389/udp, 389/tcp or 636/tcp for SSL) used for locator pings.
- Microsoft-DS traffic (445/tcp, 445/udp).

Deploying Domain Controllers in a Perimeter Network

Your forest determines the security boundary. Having a separate forest in the perimeter network is more secure than creating a forest that exists in both external and internal networks. This general rule applies to any scenario in which networks have different levels of security. If you need to enable users from the internal network to access resources in the perimeter network and the reverse, establish explicit trust between the domains. In this scenario, you can have one or more domains that have established trust with domains whose domain controllers are separated by a firewall.

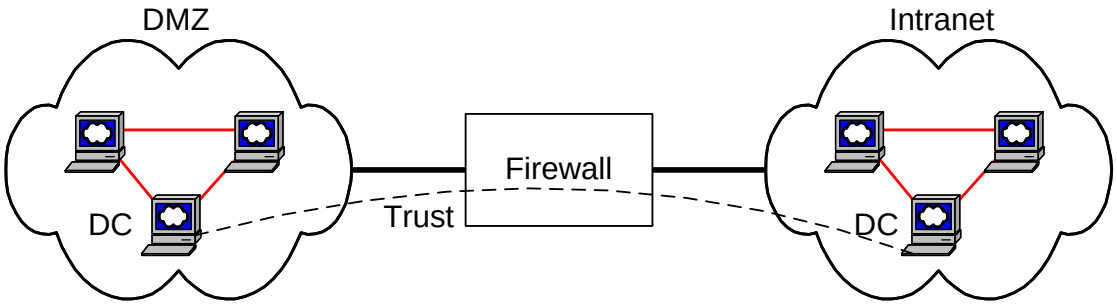

To enable internal users and resources residing in the perimeter network to have access to each other open the following ports on the firewall and create point-to-point IP restrictions so that only the specific domain controllers can communicate across the firewall. Use the User and Computer Authentication and Trust Validation operational building blocks described above to determine the type of traffic that needs to pass through the firewall and the ports that need to be opened.

- Port(s) used by specific applications if needed. You need to obtain the exact list of ports that need to be open from the application vendors. For a list of other commonly used ports, see Appendix C.
- Kerberos (88/tcp, 88udp)
- LDAP (389/udp, 389/tcp and/or 636/tcp if using LDAP over SSL)
- SMB over IP traffic (445/tcp, 445/udp)
- DNS ports (53/tcp, 53/udp) used for name lookups

If creating a separate forest for the perimeter network zone is not possible, and you are planning to deploy domain controllers of the same forest in the perimeter network and in the internal network, then replication between domain controllers needs to occur, and the appropriate configuration needs to be done. For details, see Domain Controller Replication Across a Firewall later in this document.

Deploying Active Directory in an Internal Network Containing Firewalls

Although it is not common, there are some internal corporate networks might have different portions of the network separated by firewalls.

The following diagram represents a scenario in which:

- A domain member computer (M1) is separated from the domain controller (M2) for its domain (Dom 1) by a firewall (Firewall 1 and Firewall 2).
- A user is separated from a domain controller (M3) for its domain (Dom2) by a firewall (Firewall 1 and Firewall 3).

- The user accesses a resource on a member computer (M4) of another domain (Dom3) that is located in a portion of network that is separated from the user by firewalls (Firewall 1, Firewall 4, and Firewall 5).

This scenario can take place across the Internet, an intranet, or both. In addition, the domains in this scenario can belong to a single forest or to one or mo re separate forests.

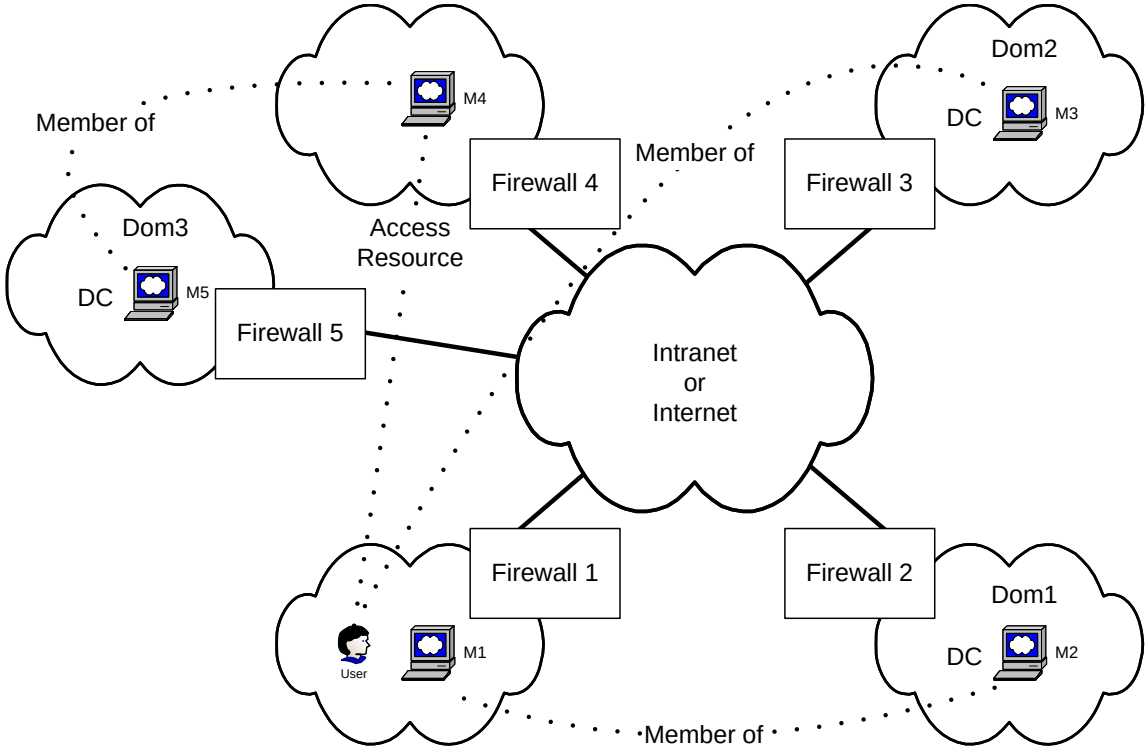

Deploying Active Directory in such networks requires additional configuration of the domain controllers and the firewalls to enable user and resource authentication and Active Directory replication across the firewalls. This scenario uses many of the operational building blocks described in Operational Building Blocks earlier in this paper.

For this scenario, open the required ports for Kerberos, LDAP, DNS, and SMB over IP. Depending upon which applications need to communicate across the firewalls, other ports might be required. You need to obtain the exact list of ports that need to be open from the application vendors. For a list of the corresponding ports, see Appendix C.

Additional configuration that is required to enable Active Directory replication is described in the following section. If your design requires that users and computers be authenticated only by local domain controllers (that is, domain controllers within the same site), and that domain controller administration is performed only from local computers (that is, computers within the same site), you need only to configure domain controller replication across the firewall, instead of opening the ports described in the previous section. For more information about configuring domain controller replication in this scenario, see Domain Controller Replication Across a Firewall later in this paper.

This scenario can also be broken down into smaller scenarios. The following illustrations show how you can use the operational building blocks described in Operational Building Blocks to construct complex scenarios.

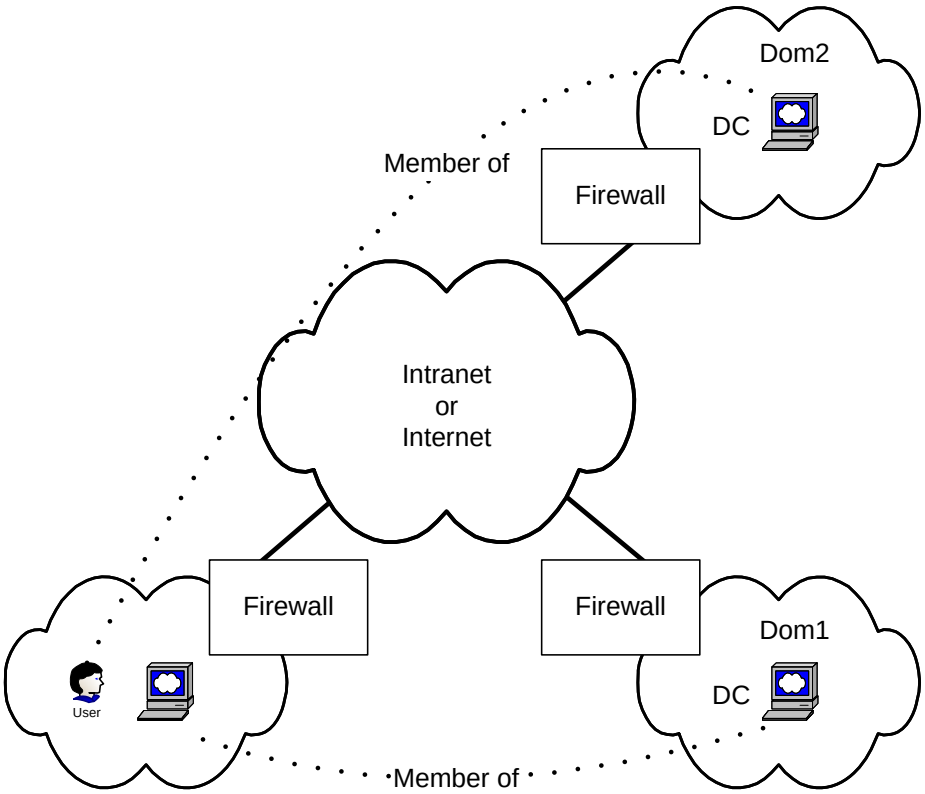

The computer and user login scenario uses the User and Computer Authentication operational building blocks and requires that the following ports be open:

- Kerberos (88/tcp, 88udp)
- LDAP (389/udp, 389/tcp and/or 636/tcp if using LDAP over SSL)
- SMB over IP traffic (445/tcp, 445/udp)

- DNS ports (53/tcp, 53/udp) used for name lookups

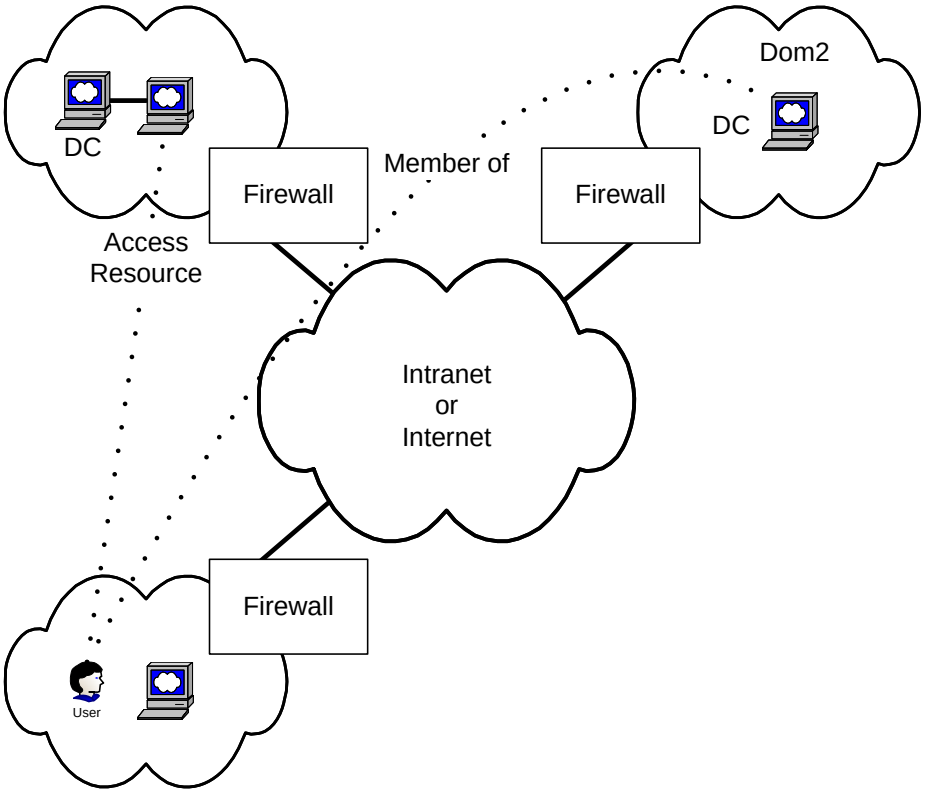

The file and resource access scenario uses the User and Computer Authentication operational building blocks and requires that the following ports be open:

- Kerberos (88/tcp, 88udp)
- LDAP (389/udp, 389/tcp and/or 636/tcp if using LDAP over SSL)
- SMB over IP traffic (445/tcp, 445/udp)
- DNS ports (53/tcp, 53/udp) used for name lookups

While opening these ports will allow these protocols through the firewall, user IDs, user password hashes and application data will be exposed in some of these protocols. Security concerns may require all traffic to be encrypted as it flows through the firewalls between the different subnets. Typically, a virtual private network (VPN) tunnel is implemented to secure traffic between firewalls. Windows 2000 Server Routing and Remote Access Service provides both VPN capabilities for site-to-site (also known as gateway-to-gateway) tunnels. However, if a VPN were not available, then an IPSec policy on each computer can enforce and automatically negotiate IPSec transport mode security for traffic between subnets using domain-based Kerberos trust. This could secure all server-to-server, client-to-client and client-to-server traffic, but should not be used to secure client-to-domain controller traffic, or domain controller to domain controller traffic. For more information, see Appendix A and article 254949, “Client-to-Domain Controller and Domain Controller-to-Domain Controller IPSec Support,” in the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462.

Domain Controller Replication Across a Firewall

Different configurations are required depending upon whether you need to enable inter- or intra-domain replication. Such configuration is described in the following two sections.

Inter-Domain Domain Controller Replication Across a Firewall

For replication between domain controllers that reside in separate domains and are separated by a firewall, encapsulate domain controller–to–domain controller traffic within IPSec and open the firewall for IPSec traffic. The IPSec filter rules should be set up to further limit the ports allowed within. For IPSec setup instructions, see Appendix A and Appendix B (for a sample script). For a list of commonly used ports, see Appendix C. For registry key modification information required to set a single port for directory service replication use, see Appendix D.

This solution provides a more secure communication channel between domain controllers while opening only a small defined set of ports on the firewall. It provides optimal firewall security as well as the ability to set up IPSec policies on each server. In addition, the Authentication Header (AH) variation of IPSec can be used to monitor the encapsulated traffic. This variation of IPSec is described in Appendix A. Although IPSec needs to be configured on each domain controller, by using Group Policy, you can apply the IPSec configuration to an organizational unit object that contains a set of domain controllers.

Intra-Domain Domain Controller Replication Across a Firewall

For Active Directory replication within a domain that has domain controllers separated by a firewall, encapsulate domain controller–to–domain controller traffic within IPSec and open the firewall for IPSec traffic. For IPSec setup instructions, see Appendix A.

This solution provides a more secure communication channel between domain controllers while opening a small defined set of ports on the firewall. It provides the highest level of firewall security, as well as the ability to set up IPSec policies on each server. Although IPSec needs to be configured on each domain controller, by using Group Policy, you can apply the IPSec configuration to an organizational unit object that contains a set of domain controllers. In this scenario dynamic RPC port allocation needs to occur to allow traffic across domain controllers; RPC port allocation also prevents specific port filters from being added in the IPSec policy.

Note

For more information about security considerations for using Active Directory in perimeter networks, see Chapter 5, “Security Design,” in the Reference Architecture Guide: Internet Data Center, at http://go.microsoft.com/fwlink/?LinkId=16468 and Chapter 8, “Directory Services,” in the Reference Architecture Guide: Enterprise Data Center, at http://go.microsoft.com/fwlink/?LinkId=16463.

Appendix A: Configuring and Managing IPSec Policies to Secure Traffic Between Domain Controllers

This appendix uses an example Internet Protocol security (IPSec) policy configuration to describe the recommended configuration for using IPSec in transport mode to secure all traffic (not just Active Directory replication traffic) between domain controllers that are on opposite sides of a firewall. Step-by-step procedures for configuring an appropriate IPSec policy, recommended practices, and considerations for managing IPSec policies that are assigned to domain controllers are provided.

You can use IPSec to secure all traffic between domain controllers in separate forests, between two domain controllers in the same domain (for example, to secure site-to-site replication traffic), between domain controllers in parent and child domains, and in other scenarios, as described earlier in this paper. For a detailed example of an IPSec policy configuration that you can use to secure traffic between two domain controllers in separate forests, see Deploying Domain Controllers in a Perimeter Network earlier in this white paper (this appendix uses this example to describe a recommended configuration for using IPSec. For more information, see Creating an IPSec Policy to Encrypt Traffic Between Two Domain Controllers later in this appendix). You can configure similar IPSec policies as needed for other scenarios. However, it is not currently recommended that you use IPSec to secure communication between domain members (either clients or servers) and their domain controllers. When domain members use IPSec-secured communication with domain controllers, increased latency might occur, and complex IPSec policy configuration and management is required. For more information, see article 254949, "Client-to-Domain Controller and Domain Controller-to-Domain Controller IPSec Support," in the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462.

Important

The example IPSec policy configuration described in this appendix secures traffic only between domain controllers that are on opposite sides of a firewall. Do not use these examples to enable IPSec-secured communication between domain members that are on opposite sides of a firewall.

Notes

The example IPSec policy configuration described in this appendix is based on a Windows 2000 Service Pack 3 or later environment. You can use the same IPSec policy configuration in a pure Windows Server 2003 environment or in a mixed-platform environment. However, note that there might be differences in the steps used to manage the Windows Server 2003 IPSec policy.

Portions of IPSec and related services for Windows 2000, Windows XP, and Windows Server 2003 were jointly developed by Microsoft and Cisco Systems, Inc.

Using the example IPSec policy configuration described in this appendix provides the following business benefits:

- IPSec allows you to configure security policies to meet the security requirements of a user, group, application, domain, site, or global organization. IPSec is integrated at the IP layer, so applications that use TCP/IP pass data to the IP layer, where it can be transparently secured by IPSec. In this way, IPSec provides security against vulnerabilities in upper-layer protocols and applications. For example, you can enhance security by using IPSec as a first layer of defense for the server message block (SMB) file-sharing protocol that is used extensively for replication and other file transfer functions. Two identified

SMB security issues were found in Windows 2000 and in Windows XP. Although supported fixes for the Windows 2000 and Windows XP issues are now available from Microsoft, you can enhance security by using IPSec as a first layer of defense for SMB or other protocols. For more information about the two identified SMB security vulnerabilities and supported fixes for Windows 2000 and Windows XP, see articles 329170, "MS02-070: Flaw in SMB Signing May Permit Group Policy to Be Modified," and 326830, "MS02-045: Unchecked Buffer in Network Share Provider May Lead to Denial-of-Service," in the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462.

- IPSec provides host-based authentication and encryption for all traffic between domain controllers to ensure that the administrative owner of the data retains full control of the data. The identity information in Active Directory constitutes the core of the data security inside the organization, so even if business and legal trust relationships that manage the trust of the network path are not enforced perfectly or are silently compromised, the IPSec-secured communications remain secured.
- IPSec allows for simple and secure firewall traversal. Firewalls interpret the many protocols that are used in communications between domain controllers as only IPSec Encapsulating Security Payload (ESP) (protocol 50) traffic or as Authentication Header (AH) (protocol 51) traffic. Firewalls permit traffic only for these protocols [and Internet Security Association and Key Management Protocol (ISAKMP) traffic], and these protocols are inherently secured against attacks.
- IPSec, with the Triple Data Encryption Standard (3DES) encryption algorithm and Secure Hash Algorithm 1 (SHA1) integrity algorithm, meets the requirement of many government, military, financial, and health care institutions that Common Criteria and FIPS 140-1-certified algorithms be used to secure their traffic. The algorithm used to encrypt traffic over most Windows protocols [for example, remote procedure call (RPC), Kerberos, and Lightweight Directory Access Protocol (LDAP)] is the RC4 stream cipher, which is not certified under Common Criteria or FIPS 140-1.
- As a software-based Windows solution, IPSec is more cost-effective for securing host-to-host communications than a hardware-based solution, such as purchasing and operating a virtual private network (VPN) or a private leased line.
- If you use IPSec offload adapters, IPSec provides lower CPU utilization than using protocol-specific security measures, such as SMB signing, provides because offload adapters accelerate the cryptographic operations that are used to secure IPSec packets, therefore minimizing the performance costs for encryption. As a result, IPSec-secured TCP/IP connections can achieve the same throughput as TCP/IP connections that are not IPSec-secured.

Note

If you cannot use IPSec offload network adapters, then IPSec encryption increases the CPU load on a domain controller. Accordingly, you might need to add more CPU capacity, depending on the available CPU and the amount of network traffic. Thorough testing is required to evaluate the performance impact of IPSec on domain controllers. For more information about the benefits of using IPSec hardware offload adapters, see IPSec Offload Performance and Comparison, at http://go.microsoft.com/fwlink/?LinkId=16469.

Although IPSec can greatly enhance security, be aware that deploying IPSec on your network requires additional training and administrative costs, and, if you need to purchase IPSec hardware offload network adapters or increase CPU capacity, it can increase hardware costs. Therefore, before deploying IPSec for any specific scenario, carefully consider and document the potential security threats that IPSec is intended to address, your security requirements, the costs of deploying IPSec, and the expected business benefits.

Creating an IPSec Policy to Encrypt Traffic Between Two Domain Controllers

When you design an IPSec policy to secure all traffic between domain controllers that are on opposite sides of a firewall, make sure that you have the following information available:

- A diagram of inter-domain and intra-domain communication paths that shows the location of each firewall and any network address translators, if network address translators are used.
- The static IP addresses of each domain controller on either side of the firewall(s).

In addition, before you create an IPSec policy, decide the following:

- Whether you want to use IPSec to encrypt traffic to provide confidentiality or to provide only data origin authentication and data integrity.
- Whether you can issue a computer certificate to each domain controller that uses IPSec (Recommended).

This appendix uses the example IPSec policy configuration that is introduced in Deploying Domain Controllers in a Perimeter Network to demonstrate how to create an IPSec policy to encrypt traffic between two domain controllers, SEA-PN-DC-01 and SEA-NA-DC-01. The figure shows that the domain controllers are in two forests that are deployed on opposite sides of a firewall. SEA-PN-DC-01 is in a perimeter network, and SEA-NA-DC-01 is in an internal network. An IP address is specified for each of these two domain controllers. To allow for communication across forest trusts between these two domain controllers, you can configure one external trust (if trust is needed in only one direction) or two external trusts, one for each direction.

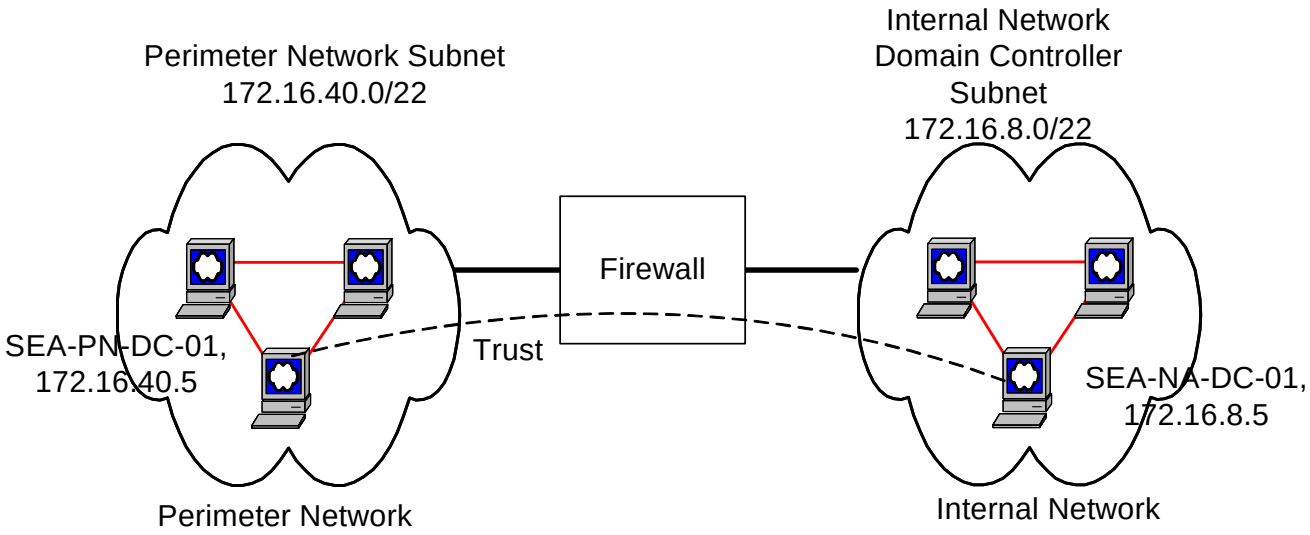

This figure is based in part on the network architecture that is documented in Windows 2000 Resource Kit Deployment Lab Scenarios. For more information, see Deployment Lab Scenarios, at http://go.microsoft.com/fwlink/?LinkId=504.

To create an IPSec policy for this example, deactivate the default response rule, and define only one rule to protect traffic between the specific IP addresses of the two domain controllers. The following table summarizes the IPSec rule that must be defined.

Source Address	Destination Address	Protocol	Source Port	Destination Port	Action	Certification Authority (CA)	Name of Rule/ Notes
172.16.40.5	172.16.8.5	Any	N/A	N/A	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	SEA-PN-DC-01<->SEA-NA-DC-01, all
172.16.8.5	172.16.40.5	Any	N/A	N/A	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirror (automatically generated)

Add a filter that specifies that all of the traffic that is sent between the IP addresses of the two domain controllers, including Kerberos, RPC, Domain Name System (DNS), LDAP, and Internet Control Message Protocol (ICMP), be IPSec-secured. If the **Mirrored** check box is selected in **Filter Properties** when you

specify the filter in a filter list, then you need to define a filter for only one direction (for clarity, the rule that is defined in the previous table includes two one-way filters).

When you define an IPSec policy with this filter, it is important to consider the following:

- If you use the Ping.exe command-line tool to verify connectivity between the two domain controllers, then Ping.exe triggers an Internet Key Exchange (IKE) negotiation. As a result, the output reports “Negotiating IP Security,” even if an IPSec security association (SA) does not exist, and the **ping** command fails. If an IPSec SA exists, or if an IPSec policy is not assigned, the **ping** command should succeed.
- If you use IPSec to protect ICMP traffic between domain controllers (as is recommended in this appendix), then you can use the **ping** command, which sends ICMP Echo Request messages to verify network connectivity, however, tools such as Tracert (which depend on the processing of ICMP traffic by routers) might not work. Tracert sends ICMP Echo Request messages with incrementally increasing Time to Live (TTL) values to determine the path taken to a destination. If IPSec protects ICMP traffic, then the Tracert output cannot show the correct path between the domain controllers because the ICMP packets are in IPSec format. In this case, the output shows only the destination computer, not the path of intermediate routers. However, if you use Tracert to determine the path taken from one domain controller to a destination other than the remote domain controller, then the output shows the intermediate destinations because the ICMP packets do match this filter and therefore are not secured by IPSec. To enable Tracert to show the routers in the path between two domain controllers, create a second rule in the IPSec policy with a filter to permit ICMP traffic between the two IP addresses of the two domain controllers, and make sure that any firewalls or other devices are configured to permit this traffic. Because IPSec does not protect ping ICMP Echo Request messages, you must use a TCP or UDP-based service to verify IPSec-secured connectivity.
- If the path between the IP addresses of the two domain controllers has a network address translator, then the network address translator cannot work because the Windows 2000 implementation of IPSec does not allow network address translators to modify IPSec packets.
- If you determine that IPSec is necessary to secure domain controller communications, then it is strongly recommended that you use the IPSec policy configuration described in this appendix. Specific IPSec policy configuration options, such as excluding only ICMP traffic, using subnet filters, using IPSec without encryption, and adding a filter to block traffic on specific ports, are described in this appendix, within the context of the overall recommended configuration. It is not recommended that you configure a customized IPSec policy with many filters to negotiate security for traffic on each protocol and port, due to the complexity of policy configuration, management, and troubleshooting required. Such a configuration can also increase costs and result in decreased performance. It is recommended that you keep your IPSec policy configuration as simple as possible.

Procedures for Defining an IPSec Policy to Encrypt Traffic Between Two Domain Controllers

To define an IPSec policy to encrypt traffic between two domain controllers, perform the following procedures:

1. Configure firewalls to permit ESP, ISAKMP, and AH traffic.
2. Remove the default filtering exemption for Kerberos and Resource Reservation Protocol (RSVP) traffic.
3. Create a filter list, and add a filter.

4. Create a filter action.

5. Create an IPSec policy, and add a rule to the policy.

6. Assign the IPSec policy.

Configuring Firewalls to Permit ESP, ISAKMP (IKE), and AH Traffic

When a firewall exists between IPSec peers, as it does in the example, you must configure the firewall to forward IPSec traffic on UDP source and destination port 500, IP protocol 50 (ESP), or IP protocol 51 (AH). First, to permit IPSec traffic on UDP source and destination port 500, use the following settings to create a firewall filter called **Permit ISAKMP traffic on UDP port 500**:

- Source address = *Specific_IP_address of domain controller*
- Destination address = *Specific_IP_address of domain controller*
- Protocol = **UDP**
- Source port = **500**
- Destination port = **500**

To permit IPSec traffic on IP protocol 50 (ESP) or IP protocol 51 (AH), use the following settings to create a firewall filter called **Permit IPSec traffic on ESP or AH protocol (50 or 51)**:

- Source address = *Specific_IP_address of domain controller*
- Destination address = *Specific_IP_address of domain controller*
- Protocol = **50** or **51**

In addition, when you configure the firewall, do the following:

- Configure the firewall to permit traffic between only the specific IP addresses of the two domain controllers (in the example, 172.16.40.5 and 172.16.8.5).
- Configure the firewall filter to permit or track fragments for ESP, ISAKMP, and AH traffic. In Windows 2000 releases through Service Pack 4 and in Windows XP and Windows XP Service Pack 1, IKE message fragmentation is required when certificate authentication is used. Also, many UDP applications do not attempt to avoid fragmentation, and therefore the UDP traffic is fragmented when IPSec protects it.
- Allow IKE and IPSec communications to flow statically in both directions between the IP addresses of the domain controllers. Do not configure the firewall to perform stateful filtering on UDP source and destination port 500 (ISAKMP), IP protocol 50 (ESP), or IP protocol 51 (AH).
- If communication requires TCP path maximum transmission unit (MTU) discovery, configure the firewall to permit ICMP Destination Unreachable messages.

Note

If a non-Microsoft firewall or other network device is performing network address translation on the traffic between the domain controllers, then you cannot use Windows 2000 IPSec to secure traffic end-to-end between the domain controllers. In addition, you cannot use Windows 2000 IPSec to secure traffic end-to-end between the domain controllers if a computer running Microsoft Internet Security and Acceleration (ISA) Server is used as a firewall between domain controllers, because ISA Server always performs network address translation on traffic that passes through it. For more information, see article 329807, "INFO: ISA Server Does Not Support Domain Members In Perimeter Network," article 254949, "Client-to-Domain Controller and Domain Controller-to-Domain Controller IPSec Support," in the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462.

Removing the Default Filtering Exemption for Kerberos and RSVP Traffic

After you configure the firewall to permit ESP, ISAKMP, and AH traffic, you must remove the default filtering exemption for Kerberos and RSVP traffic to ensure that IPSec can secure these traffic types. By default, in Windows 2000 and Windows XP, broadcast, multicast, Kerberos, RSVP, and ISAKMP traffic is exempt from IPSec filtering, even if you define a filter to match all IP traffic between the IP addresses of the two domain controllers. To secure Kerberos and RSVP traffic between the IP addresses of the two domain controllers, you must remove this default filtering exemption by modifying the registry.

Caution

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

To remove the default filtering exemption for Kerberos and RSVP traffic

1. Under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IPSEC**, add a new DWORD entry named NoDefaultExempt.
2. Assign this entry a value of 1. This specifies that Kerberos and RSVP traffic are not exempt from IPSec filtering (multicast, broadcast, and ISAKMP traffic are exempt).
3. Restart the IPSec service.

Important

In Windows Server 2003 IPSec, it is not necessary to set this registry key because the default filtering exemptions for Kerberos and RSVP traffic have been removed.

For more information, see article 254728, "IPSec Does Not Secure Kerberos Traffic Between Domain Controllers," and article 811832, "IPSec Default Exemptions Can Be Used to Bypass IPsec Protection in Some Scenarios," in the Microsoft Knowledge Base. To find these articles, see the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462.

Creating a Filter List and Adding a Filter

After you remove the default filtering exemption for Kerberos and RSVP traffic, create a filter list, and then add a filter to this list to define the traffic to secure between the two domain controllers. This procedure

creates a filter list and adds a filter to the list, using the IP addresses of the two domain controllers, SEA-NA-DC-01 and SEA-PN-DC-01, as an example.

To create a filter list and add a filter

1. On a domain controller, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Domain Controller Security Policy**.

   The Microsoft Management Console (MMC) window that opens displays the default domain controller Group Policy object (GPO) that is associated with the domain controller's organizational unit (OU) in Active Directory.

2. In the console tree, click **Windows Settings**, click **Security Settings**, and then click **IP Security Policies**.

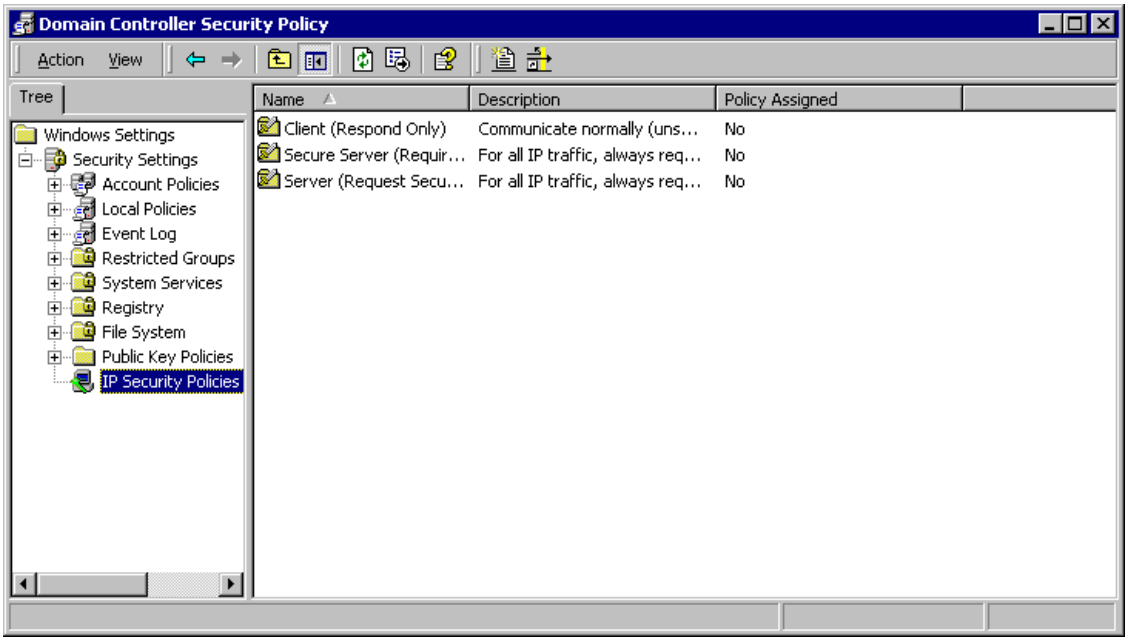

Note

If you change the filter settings or any other settings for an existing Active Directory-based IPSec policy that is assigned to a domain controller, or if you assign a new Active Directory-based IPSec policy to a domain controller, make sure to carefully coordinate these changes. Although the IPSec policy in this example is shown as part of the security settings for this GPO, the GPO contains only a reference to the IPSec policy. Group Policy detects changes only in IPSec policy assignments; it does not detect changes within an IPSec policy after it is assigned to a GPO. The IPSec service detects changes in the related IPSec policy. The differences between Group Policy and IPSec service polling intervals can result in incompatible policies if changes in policy settings or assignments are not carefully coordinated. For more information, see Considerations for Updating Active Directory-Based IPSec Policy, later in this appendix.

3. Right-click the details pane, click Manage IP filter lists and filter actions, and then, in Manage IP filter lists and filter actions, click Add.

4. In **IP Filter List**, type a name and description for the filter list, and then click **Add**.

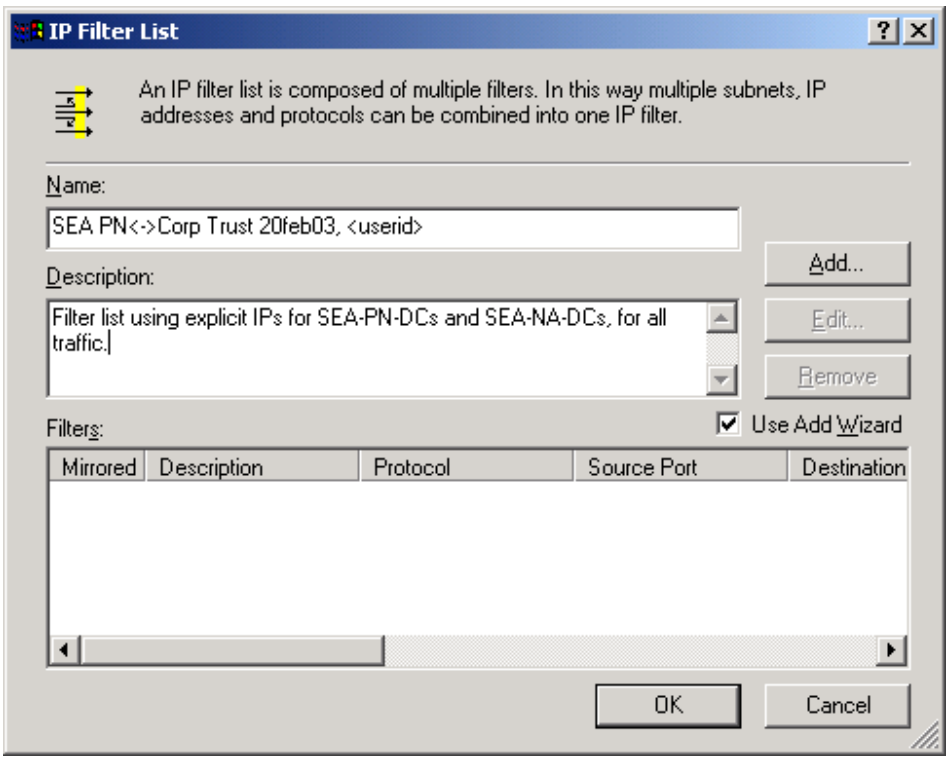

5. On the IP Filter Wizard welcome page, click **Next**.

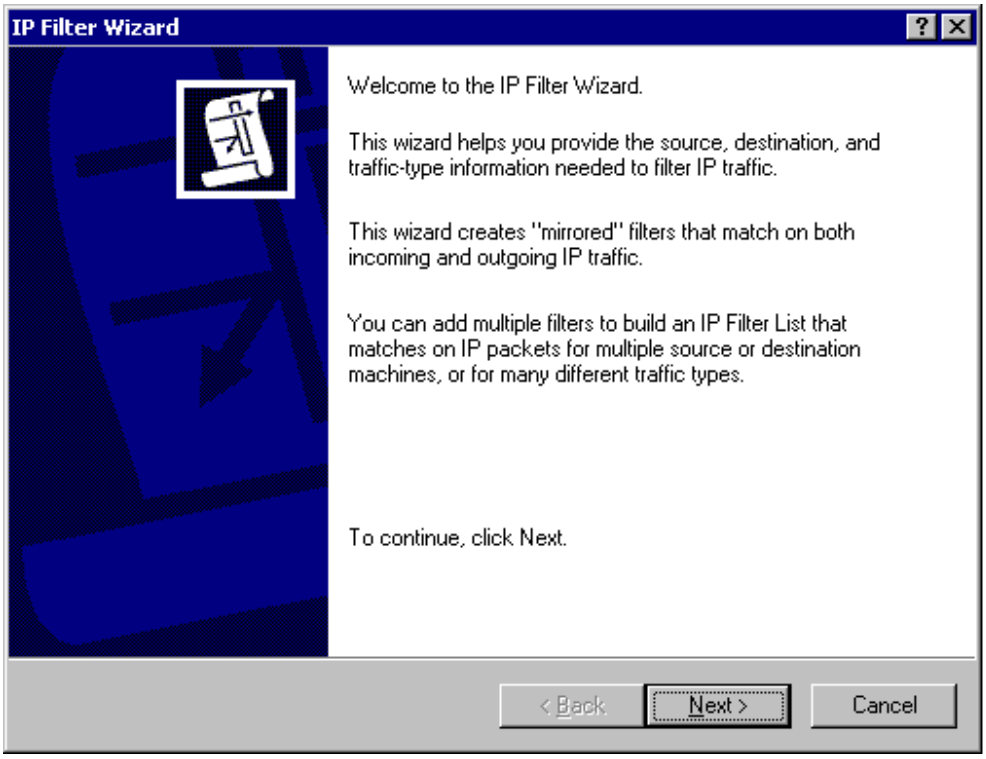

6. In **IP Traffic Source**, click **A specific IP Address**, type **172.16.40.5**, and then click **Next**.

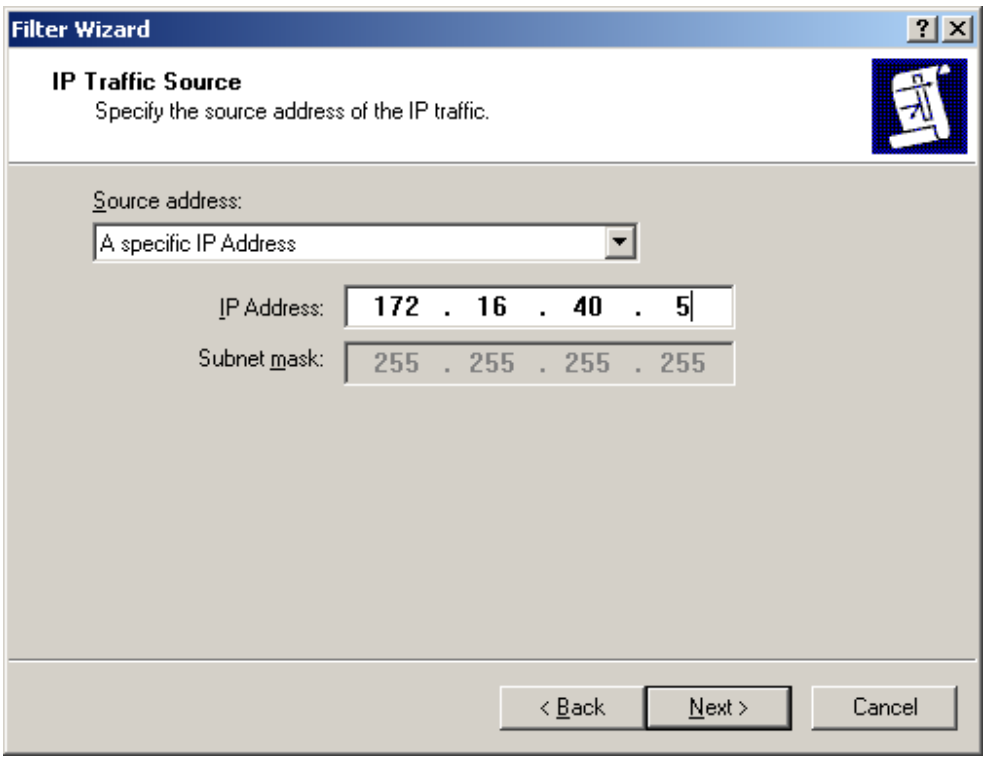

7. In **IP Traffic Destination**, click **A specific IP Address**, type **172.16.8.5**, and then click **Next**.

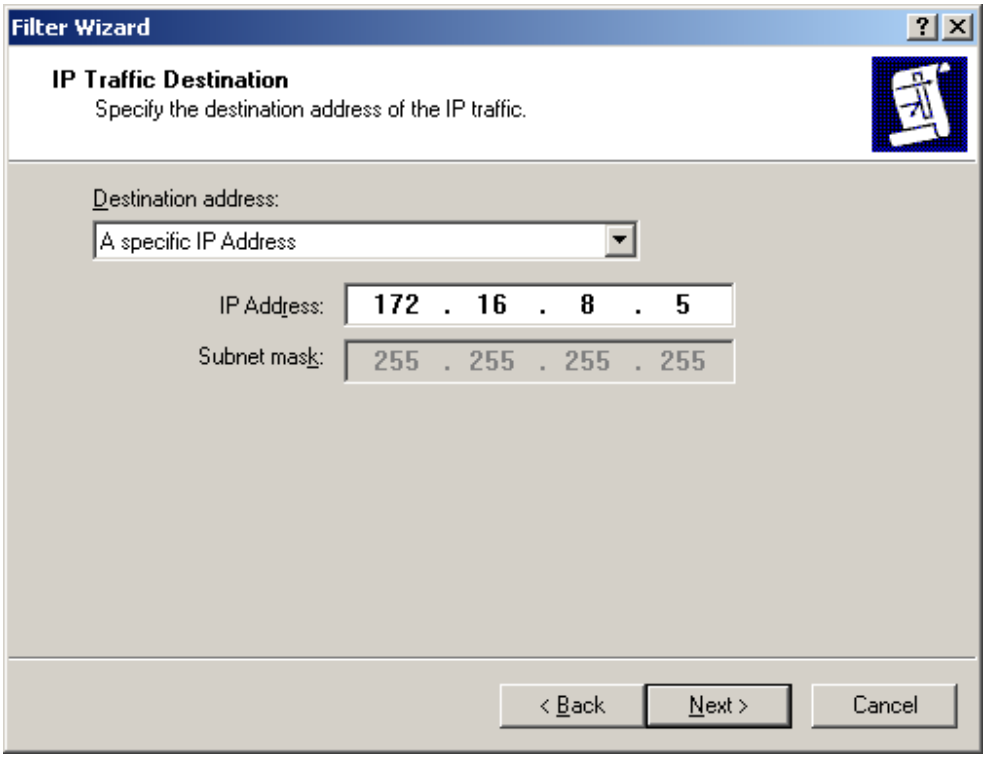

8. In **IP Protocol Type**, click **Any** (or click a specific protocol, and, if you click TCP or UDP, specify a port number), and then click **Next**.

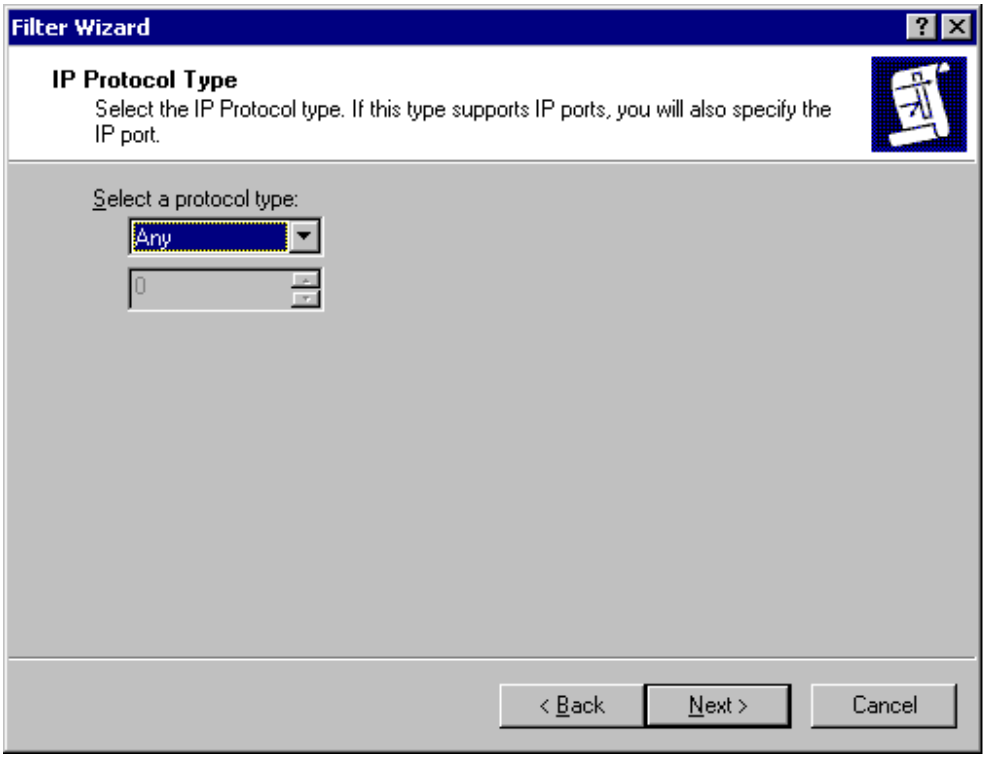

9. On the IP Filter Wizard completion page, click **Finish**, and then click **Close**.

After you create the filter, it appears as follows:

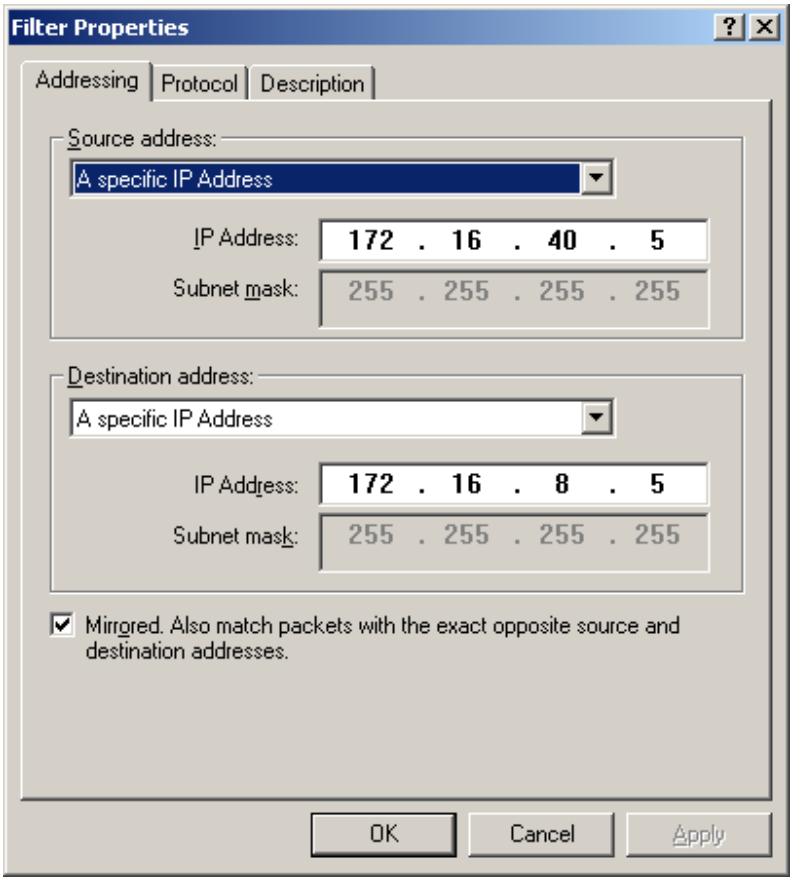

The filter list displays one mirrored filter, which appears as follows:

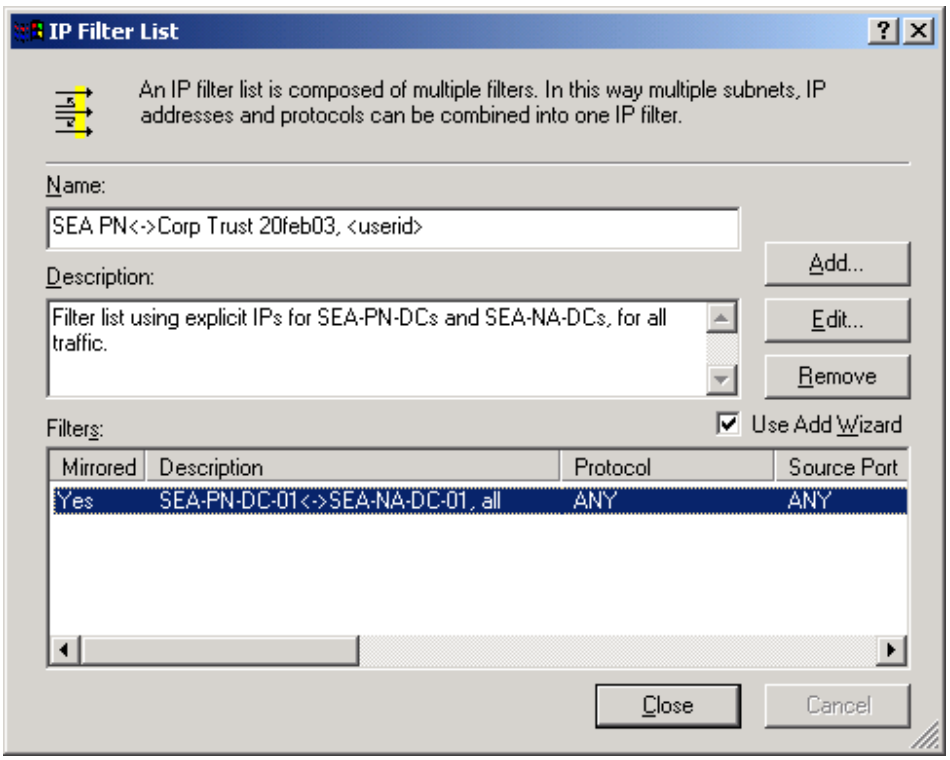

Note

The screenshot shows an optional description (SEA-PN-DC-01<->SEA-NA-DC-01, all) in the **Description** column in the properties for the filter. However, you cannot add this description by using the IP Filter Wizard. To add a description for a filter, in **Filter Properties**, click the **Description** tab, and, in **Description**, type an abbreviated description. Although such descriptions are not required, it is recommended that you specify them to aid in troubleshooting.

Creating filters to secure traffic between additional domain controller IP addresses

You can create additional filters in the same filter list to secure traffic between the specific IP addresses of other domain controller pairs. If you do so, all filters in the filter list use the action that is associated with the filter list in the rule.

If there are two domain controller IP addresses in each domain, add a filter list that contains four mirrored filters, as shown in the following table. When you do so, eight one-way filters are generated.

Mirrored	Source Address	Destination Address	Protocol	Source Port/ Destination Port	Name/ Comment
Yes	SEA-PN-DC-01	SEA-NA-DC-01	Any	N/A	SEA-PN-DC-01<->SEA-NA-DC-01, all
Yes	SEA-PN-DC-01	SEA-NA-DC-02	Any	N/A	Similar
Yes	SEA-PN-DC-02	SEA-NA-DC-01	Any	N/A	Similar
Yes	SEA-PN-DC-02	SEA-NA-DC-02	Any	N/A	Similar

This IPSec policy can be assigned to every domain controller in both forests (the perimeter network and the internal network). The policy takes effect only when a domain controller in either forest sends or receives packets that match the addresses in the filters.

Note that these filters are not defined to use **My IP Address** as a source or destination address. The four mirrored filters are defined to use explicit IP addresses for source and destination addresses, instead of **My IP Address**, for the following reasons:

- When you want IPSec to secure communication between only one domain controller in each forest, using explicit IP addresses in the IPSec policy allows the policy to be assigned to all of the computers in the domain controller's OU. Although the IPSec policy is enforced in this case, it does not affect computers that do not have one of the specified IP addresses.
- If you use IPSec to secure traffic between domain controllers in the same domain, then specifying **My IP Address** might cause IPSec to secure communication on internal paths between domain controllers that you do not want secured. For example, if you set up the perimeter network domain as a remote site for the internal domain, do not assign an IPSec policy to domain controllers in the perimeter network domain that includes only the **My IP Address** filter shown in the table below. If you specify this filter in an IPSec policy that is assigned to domain controllers in the internal network, then each domain controller in the internal network attempts to negotiate security with SEA-PN-DC-01. Although this behavior might be intended, if you use the same filter in an IPSec policy that is assigned to domain controllers in the perimeter network domain, then each domain controller requires IPSec when communicating with SEA-PN-DC-01 internally, within the perimeter network domain. To avoid using IPSec to secure communication between all domain controllers in a site, use static IP addresses for each domain controller in the site.
- If you do not explicitly specify the IP address of each domain controller in the IPSec filter list, then communication to one domain controller might fail. For example, if the **My IP Address** filter is the only filter in the IPSec policy that is assigned to all domain controllers in the perimeter network domain, then communication is blocked between SEA-PN-DC-01 and other domain controllers in the same domain because SEA-PN-DC-01 interprets the **My IP Address** filter as requiring IPSec communication to its own IP address and ignores the filter. Because SEA-PN-DC-01 does not have other filters that allow it to negotiate security with the specific IP addresses of other domain controllers, it attempts to send unsecured traffic to the other domain controllers. The other domain controllers, which require IPSec communication, drop the unsecured packets that SEA-PN-DC-01 sends, and communication fails. Likewise, SEA-PN-DC-01 ignores security negotiation requests from other domain controllers because it does not have an IPSec policy that allows it to respond to requests to negotiate security with the IP addresses of the other domain controllers.
- Using explicit IP addresses for the source or destination address for a domain controller that uses more than one network adapter (which is often the case) or a domain controller that has more than one IP address assigned on a single network adapter, prevents IPSec from generating filters for each IP address on the computer, which might cause more filters to be created than necessary, increasing CPU utilization.

However, it can be appropriate to define a filter to use **My IP Address**, as is shown in the following table:

Mirrored	Source Address	Destination Address	Protocol	Source Port/ Destination Port	Name/ Comment
Yes	My IP Address	SEA-PN-DC-01	Any	N/A	Any DC<->SEA-NA-DC-01, all

Use the **My IP Address** filter for the following conditions:

- When the IPSec policy is designed uniquely for each site (for example, when you want to secure all traffic from any internal domain controller to a specific domain controller in the perimeter network). You can add the **My IP Address** filter to the IPSec policy for the domain controllers in the internal network subnet, rather than creating a filter for each domain controller IP address pair.
- When you want to secure all internal traffic between many domain controllers, all domain controllers to use the same IPSec policy, and the number of IPSec filters that are required for the policy to be minimized. For example, if you have a domain with 100 domain controllers and you do not use the **My IP Address** filter, you need to specify a filter for each pair of domain controller IP addresses. In this case, 4,950 mirrored filters would be required [(100 x 100) – 100]/2. Because any domain controller could communicate with a maximum of only 99 other domain controllers in the same domain, most of these filters would not be needed. If you use **My IP Address** as a source address, in the same scenario, then you need only 100 filters to specify the destination IP address of each domain controller.

Creating filters to secure traffic between subnets

If there are many domain controllers in either domain, then the number of IP address combinations can be complex to manage. In such cases, consider creating filters to secure traffic between subnets. However, keep in mind that if you create filters to secure traffic between subnets, rather than between specific IP addresses, then communication between the domain controller in the source subnet and member computers in the destination subnet might be blocked, if the member computers in the destination subnet do not also have an appropriate IPSec policy and mutual authentication method for IKE to successfully negotiate security.

For the example described in this appendix, to secure traffic between the subnet for the perimeter network (which contains SEA-PN-DC-01) and the internal network (which contains SEA-NA-DC-01), define the following filter list:

Mirrored	Source Address	Destination Address	Protocol	Source Port/ Destination Port	Name/ Comment
Yes	172.16.40.0/255.255.252.0	172.16.8.0/255.255.252.0	Any	N/A	SEA-PN-DC subnet<->SEA-NA-DC subnet, all

When a filter like this applies to the domain controllers, the domain controllers negotiate security for any packet that is sent with a matching source and destination IP address. If the IP addresses of the domain controllers are closely grouped within a smaller subnet range, then you can use a smaller subnet in the source or destination address. For example, if the domain controller IP addresses in the perimeter network subnet range from 172.16.40.2 through 172.16.40.8, then you can use a subnet definition of 172.16.40.0/255.255.255.248 (172.16.40/29).

If you use subnet filters, keep in mind that IPSec-secured packets must be small enough to fit through the smallest link of the network path between the domain controllers. The maximum packet size supported for any link on a network path is known as the maximum transmission unit (MTU). The smallest MTU supported for all links on the network path is known as the path MTU (PMTU). Through the process of PTMU discovery, if a packet exceeds the MTU for any link along the network path, then the router sends an ICMP Destination Unreachable packet to the source address of the TCP packet, to indicate that the

packet cannot be forwarded unless it is fragmented. When the message is received, TCP adjusts its MTU for the connection so that any packets sent on the connection are no larger than the MTU.

The Windows implementation of IPSec is integrated with TCP so that TCP automatically reduces the TCP packet size when IPSec must add additional bytes to the packet for AH or ESP headers. As a result, IPSec-secured TCP packets are sized according to the MTU allowed by the outbound network adapter. A router or VPN gateway in the network path might still require a smaller packet, however, and therefore PTMU discovery must function correctly so that the MTU of the packet can be determined. For PTMU discovery to function correctly, you must add filters to your IPSec policy to permit inbound ICMP Destination Unreachable messages from routers or other gateways.

When IPSec filters are used to specify the IP addresses of individual domain controllers, inbound ICMP Destination Unreachable messages are not blocked. However, if you use the subnet filter that is specified in the preceding table on SEA-NA-DC-01 to require IPSec for any traffic to the perimeter network subnet, then an ICMP Destination Unreachable message from a router with an IP address of 172.16.40.1 is dropped. As a result, PTMU discovery does not function correctly and TCP communication to SEA-PN-DC-01 might be delayed or fail every time a packet that exceeds the PMTU is sent.

To ensure that PMTU discovery functions correctly when you use larger subnet filters, add the following one-way filter to the IPSec policy that is assigned to SEA-NA-DC-01. This filter permits inbound ICMP Destination Unreachable messages from the entire perimeter network subnet, 172.16.40.0/22.

Source Address	Destination Address	Protocol	Source Port	Destination Port	Action	Certification Authority (CA)	Name of Rule/ Notes
172.16.40.0/255.255.252.0	My IP Address	ICMP	N/A	N/A	Permit	N/A	Allow PMTU

If you can add a more specific filter to require IPSec communication between domain controllers (for example, if you use the smaller subnet range of 172.16.40.0/255.255.255.248), then you might not need to use the filter in the preceding table to permit inbound ICMP traffic. Alternatively, if you cannot specify a smaller subnet range for the filter, and only a few routers need to send ICMP Destination Unreachable messages, then add a filter that uses the specific IP address of the router as the source address for ICMP traffic.

Note that the ICMP filter in this example is not mirrored so that outbound ICMP packets used by the **ping** command will trigger an IKE negotiation and be protected by IPSec. The **ping** command allows different sizes of ICMP packets to be sent and forces routers to send ICMP Destination Unreachable messages when required. To determine the maximum size of an IPSec packet, run the following **ping** command on SEA-NA-DC-01, after you assign IPSec policy on both domain controllers and verify that IPSec SAs are established:

```
ping –f –l <size> 172.16.40.5
```

(where <size> starts at 1200 and increases until the **ping** command no longer elicits a reply).

If firewalls are configured to block inbound ICMP traffic to the domain controller that is sending IPSec-secured TCP traffic, or if the inbound ICMP traffic to the domain controller is blocked for other reasons, then consider configuring the network adapter on the domain controller to send smaller TCP packets, or

configure TCP/IP to detect when large packets are being dropped and to automatically attempt to reduce the size of the packet. However, because these two configuration options might cause delays in all other TCP/IP communications with the domain controller, you should implement these options only if you cannot add an IPSec filter to allow inbound ICMP messages on the domain controller that is sending IPSec-secured TCP traffic.

When you design the IPSec policy for SEA-PN-DC-01, if you plan to use subnet filters, consider whether inbound ICMP traffic must be permitted to SEA-PN-DC-01 because IPSec-secured TCP packets that are sent from the perimeter network to the internal network might also exceed the PMTU. For more information about PMTU discovery and how to configure TCP/IP and network adapters to use different methods for reducing MTU, see the “Path Maximum Transmission Unit (PMTU) Discovery” section and the **EnablePMTUBHDetect** and **MTU** registry key descriptions in Microsoft Windows 2000 TCP/IP Implementation Details, at http://go.microsoft.com/fwlink/?LinkId=16467.

Creating filters to secure traffic between IP addresses and subnets

If you create filters to secure traffic between a combination of IP addresses and subnets, make sure that both sides of the communication use exactly the same filter definition. For example, the following filter list corresponds to the example:

Mirrored	Source Address	Destination Address	Protocol	Source Port/ Destination Port	Name/ comment
Yes	172.16.40.0/255.255.255.0	172.16.8.5	Any	N/A	SEA-PN-DC subnet<->SEA-NA-DC 01, all

When you use this filter for an IPSec policy that is assigned to all domain controllers in the perimeter network (SEA-PN-DC) and in the internal network (SEA-NA-DC) subnet, only one domain controller in the internal network (SEA-NA-DC-01) will initiate a security negotiation with any of the domain controllers in the perimeter network subnet, because only SEA-NA-DC-01 has an IP address that matches this filter. If there is a second domain controller in the internal network (SEA-NA-DC-02) with an IP address of 172.16.8.6, traffic between that domain controller and any of the domain controllers in the perimeter network subnet is not matched against the filter, because SEA-NA-DC-02 has a different IP address. Therefore, SEA-NA-DC-02 does not initiate a security negotiation when attempting to communicate to any domain controller in the perimeter network subnet, and communication fails. Likewise, the domain controllers in the perimeter network subnet will always initiate a security negotiation with SEA-NA-DC-01 but not with any other domain controller IP address in the internal network subnet. Communication should fail in these cases because the firewall should be configured to block all unsecured communication between domain controllers.

If you use two IPSec policies with filters that do not match, the policies cannot work together. For example, if you use one IPSec policy for all domain controllers in the perimeter network subnet with a filter that matches all traffic from that subnet and the specific IP address of SEA-NA-DC-01, and you use a second IPSec policy with a different filter for SEA-NA-DC-01 that matches traffic between the specific IP addresses of SEA-NA-DC-01 and SEA-PN-DC-01, these two policies cannot work together, and in some cases, the security negotiation might fail. When any domain controller in the perimeter network subnet initiates a security negotiation first to SEA-NA-DC-01, the negotiation fails because SEA-NA-DC-01 is not configured to accept a subnet as a source address. However, if SEA-NA-DC-01 initiates a security negotiation first to any domain controller in the perimeter network subnet, the negotiation succeeds

because the filter that SEA-NA-DC-01 uses is more specific than the filter used by the domain controllers in the perimeter network subnet.

Enabling and disabling network adapters

A Windows 2000 IPSec policy with filters that match traffic between specific IP addresses is enforced when the IPSec service starts and detects that the specified IP addresses are already configured on the network adapter. If the network adapter is disabled during computer startup or when the IPSec service is started, and if the network adapter is configured to use an IP address that is specified in the policy, then you must restart the IPSec service after re-enabling the network adapter, or the IPSec policy is not enforced. If the network adapter is enabled when the IPSec service is started and then disabled and re-enabled while the IPSec service is still running, then the IPSec policy is still enforced. If you use a filter with a source address of **My IP Address** in the IPSec policy, then Windows 2000 automatically creates the appropriate IPSec policy when you add a network adapter or an IP address to the computer. Windows Server 2003, however, can automatically create the appropriate IPSec policy when you use an IPSec policy with filters that match traffic between specific IP addresses, and the network adapter is disabled during computer startup.

Creating a Filter Action

After you create a filter list and add filters to define the traffic that IPSec is to secure, create a filter action to specify how IPSec is to secure traffic between the domain controllers. This procedure defines how traffic between SEA-NA-DC-01 and SEA-PN-DC-01 is to be secured, as an example.

To create a filter action

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Right-click the details pane, click **Manage IP filter lists and filter actions**, click the **Manage Filter Actions** tab, and then, in **Manage IP filter lists and filter actions**, click **Add**.

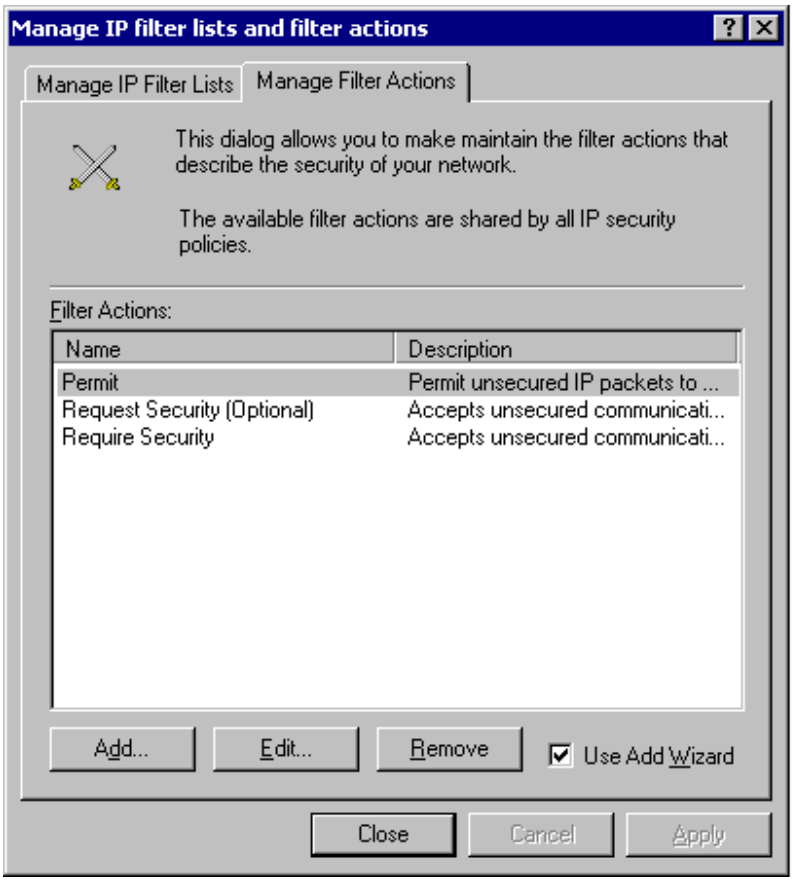

3. On the Filter Action Wizard welcome page, click **Next**.

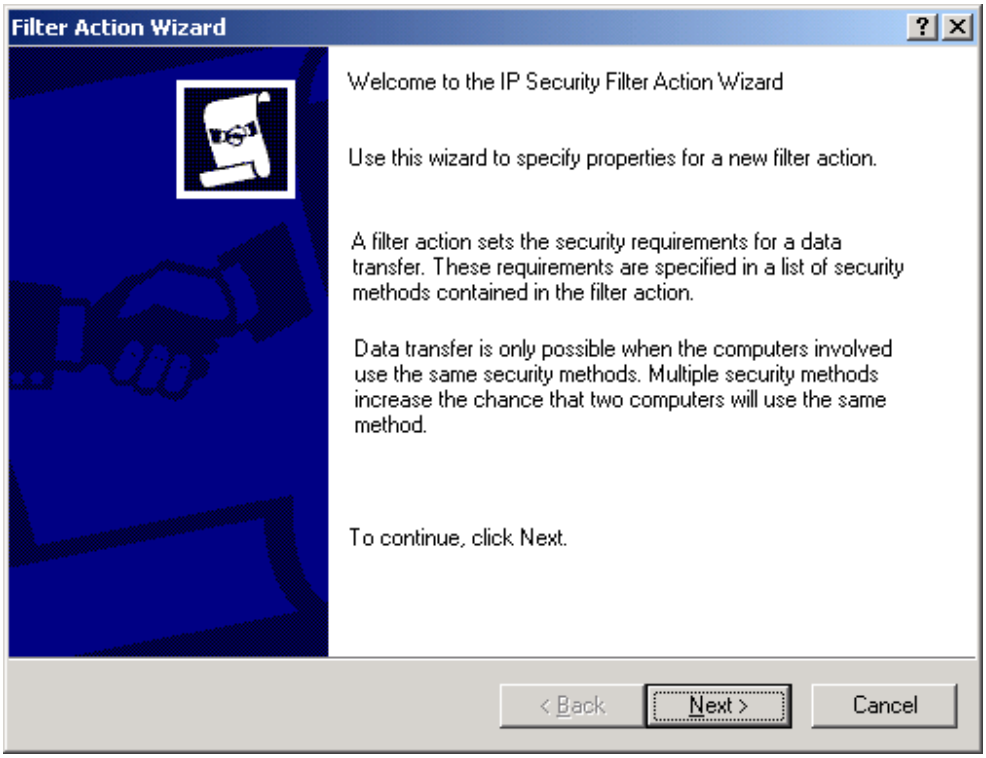

4. In **Filter Action Name**, type a name and description for the filter action, and then click **Next**.

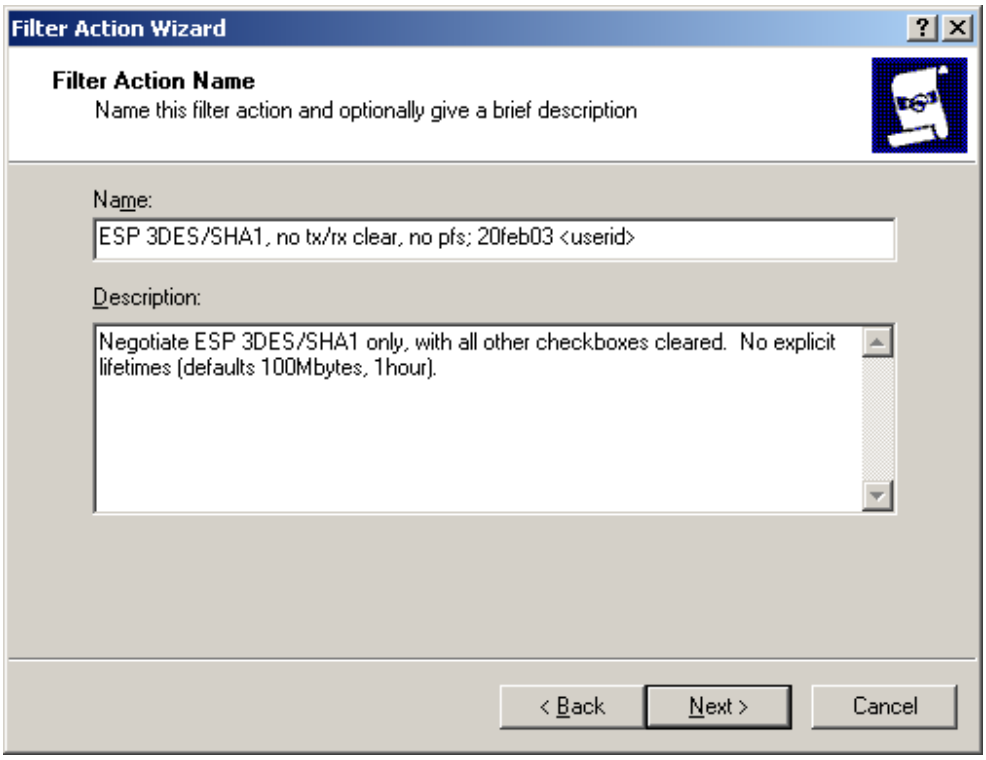

5. In **Filter Action General Options**, click **Negotiate security**, and then click **Next**.

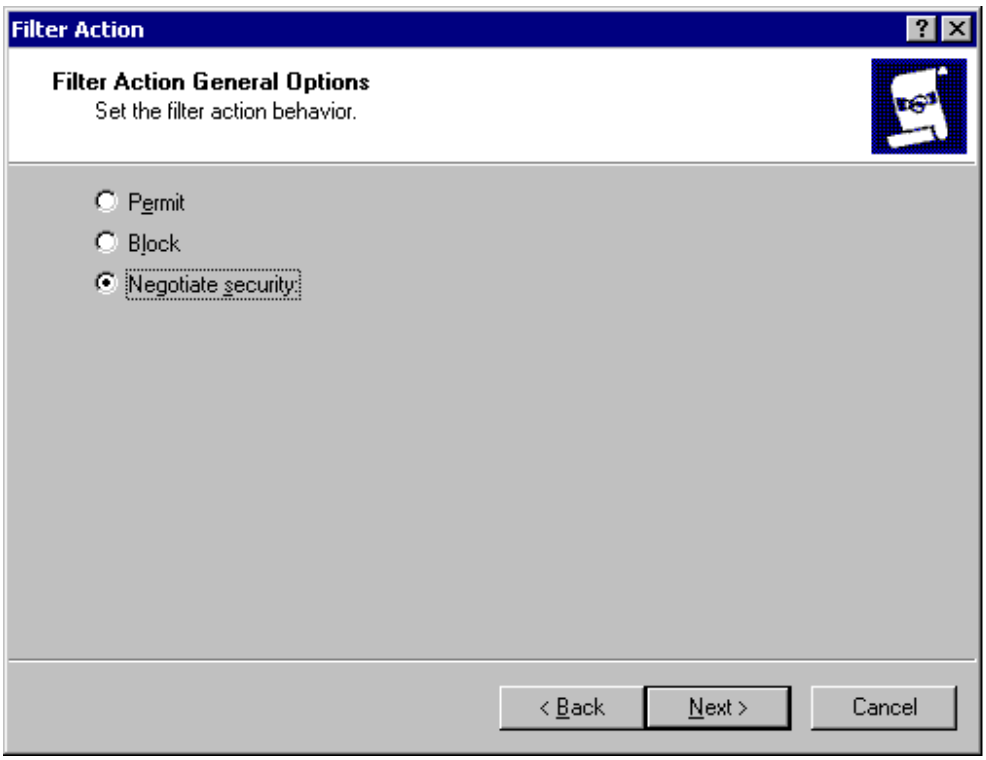

6. In **Communicating with computers that do not support IPSec**, click **Do not communicate with computers that do not support IPSec**, and then click **Next**.

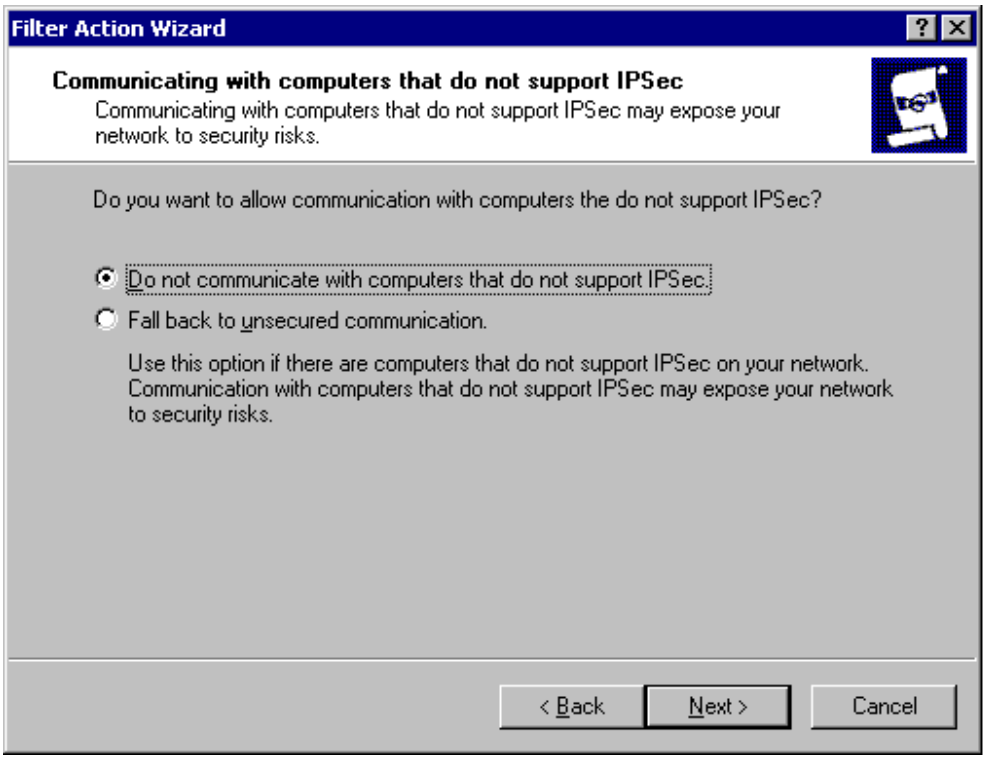

7. In **IP Traffic Security**, click the security method that you want to use.

The IPSec policy described in the example in this appendix uses ESP with 3DES encryption. To configure IPSec to use ESP with 3DES encryption, click **Custom**, and then click **Settings**.

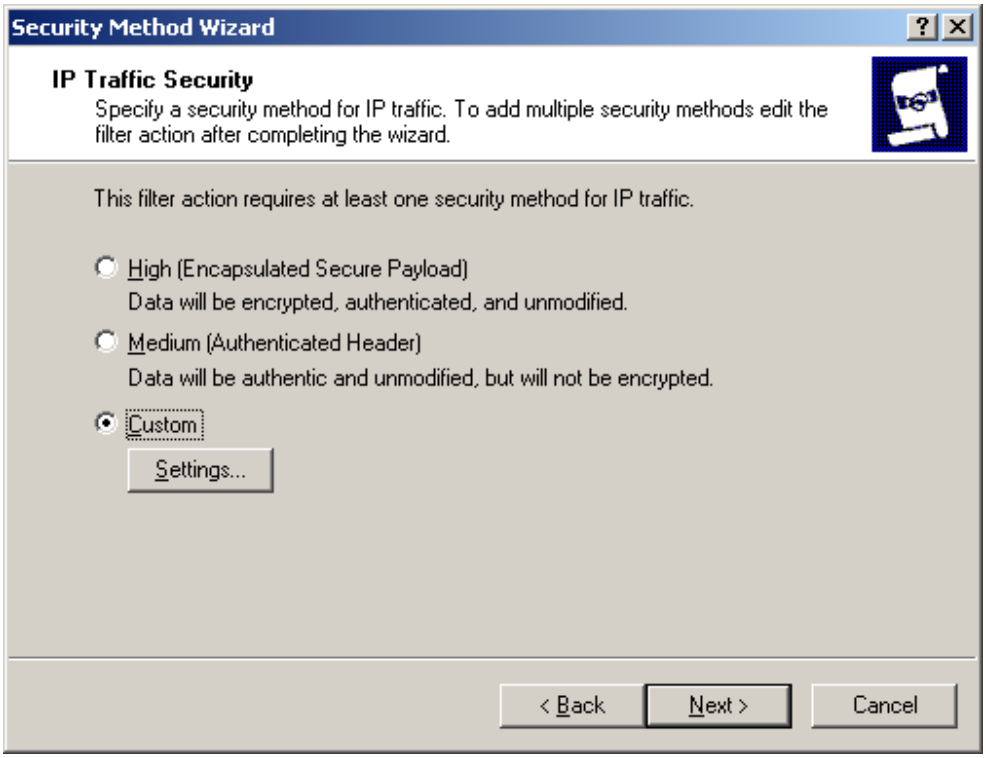

Notes

If you select **High** rather than **Custom**, IPSec also uses ESP. However, 56-bit DES encryption is used by default, rather than 3DES encryption. If 3DES encryption is available in your locale, it is recommended that you use this stronger level of encryption, rather than DES encryption. Computers running Windows 2000 must have the High Encryption Pack or Service Pack 2 (or later) installed in order to use 3DES. If a computer running Windows 2000 is assigned a policy that uses 3DES encryption, but does not have the High Encryption Pack or Service Pack 2 (or later) installed, the security method defaults to the weaker DES algorithm. To ensure at least some level of privacy for communication, make sure to allow DES as a fallback option whenever a 3DES setting is applied to a group of computers, in case some of them cannot support 3DES. For more information, see Windows 2000 High Encryption Pack, at http://go.microsoft.com/fwlink/?LinkId=7272.

You can use IPSec AH encapsulation, rather than ESP encryption, if you do not need to encrypt all traffic, or if you are troubleshooting IPSec. To use AH, click **Medium**.

8. To continue to configure IPSec to use ESP with 3DES encryption, in **Custom Security Method Settings**, select the **Data integrity and encryption (ESP)** check box. Under **Integrity algorithm**, click **SHA1**, under **Encryption algorithm**, click **3DES**, and then click **OK**.

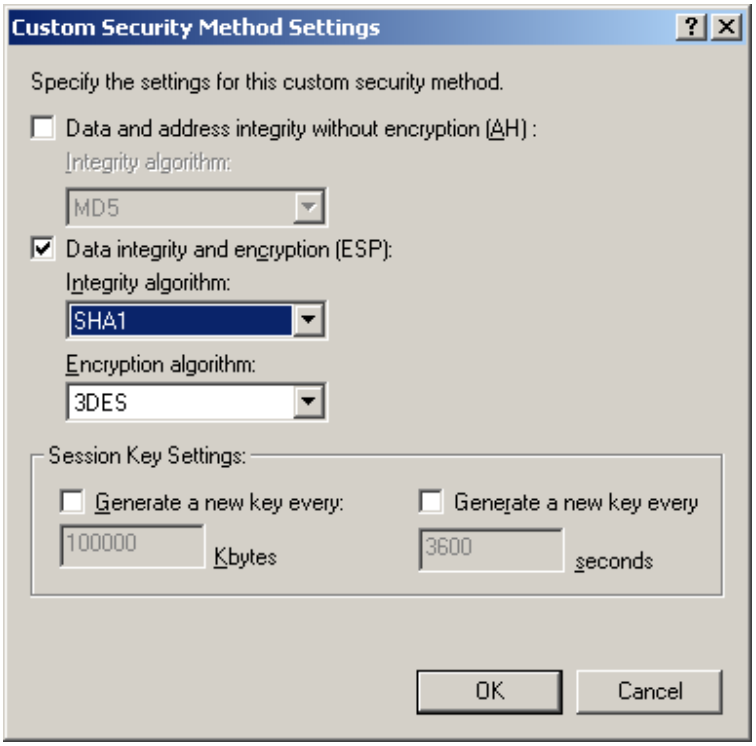

Note

You should not need to modify the default session key settings. These defaults are used if you do not configure other values.

9. In **IP Traffic Security**, click **Next**.

10. On the Filter Action Wizard completion page, select the **Edit Properties** check box, and then click **Finish**.

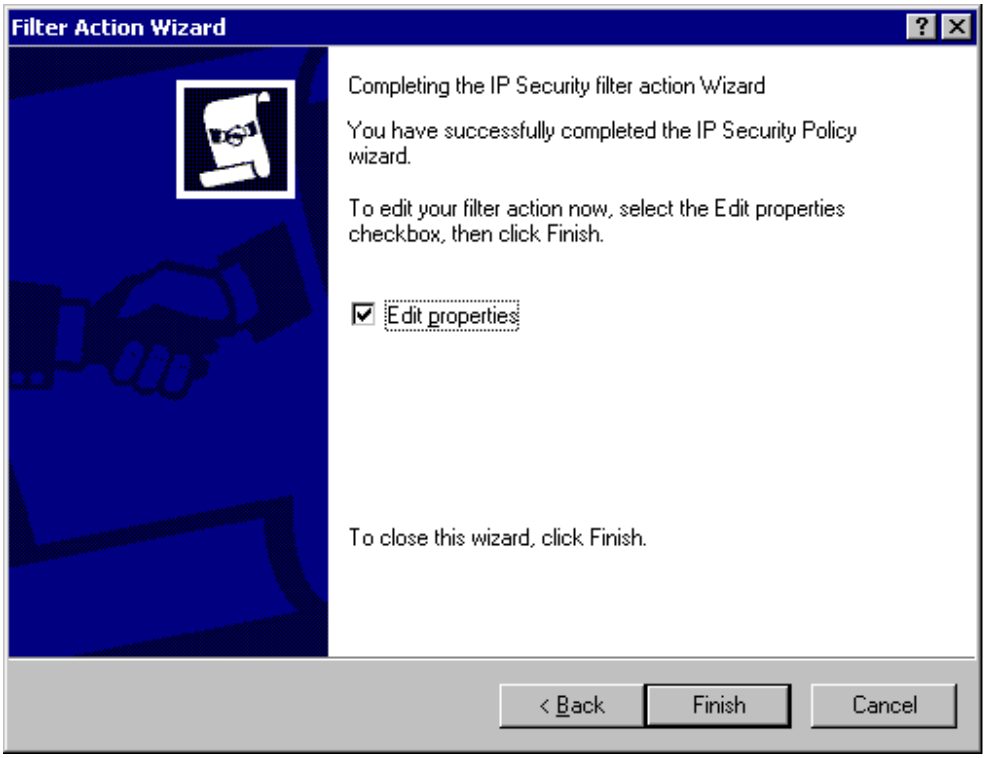

11. In the filter action properties dialog box, clear the **Accept unsecured communication, but always respond using IPSec** check box, and then click **OK**.

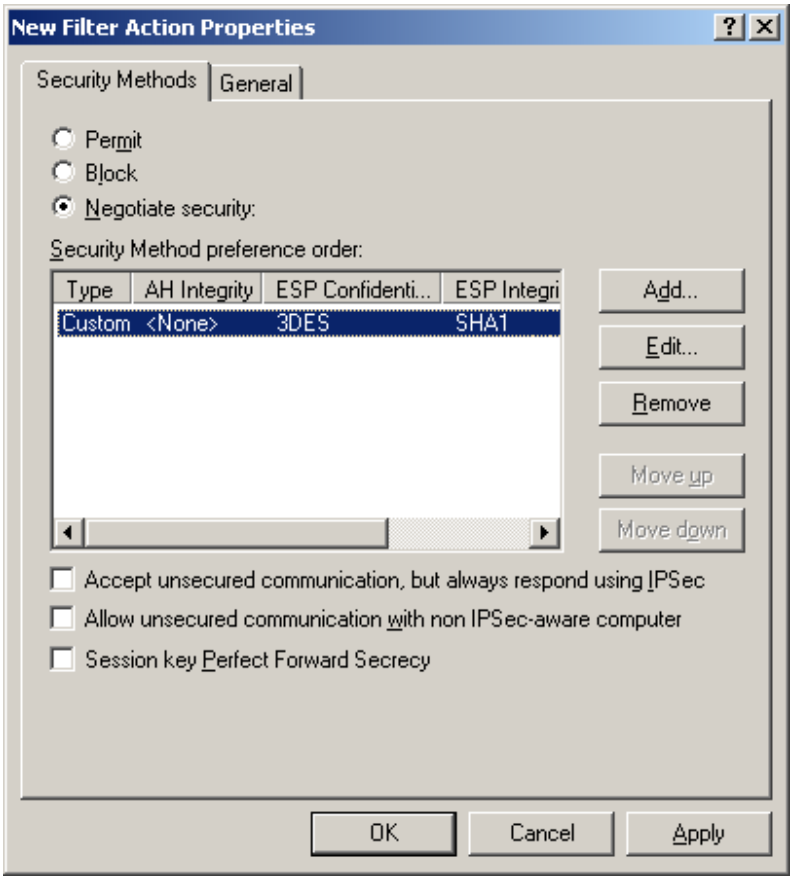

Notes

When you define the **Negotiate security** filter action so that unsecured communication is not accepted (by verifying that the **Accept unsecured communication, but always respond using IPSec** and **Allow unsecured communication with non-IPSec aware computer** check boxes are cleared), the security negotiation fails if IKE does not receive a response from the destination computer. However, if you have two domain controllers in the same domain that must replicate IPSec policy settings, or if you are not assigning an active IPSec policy to all domain controllers simultaneously, select the **Allow unsecured communication with non-IPSec-aware computers** check box during the rollout phase, so that communication is not blocked after some of the domain controllers start using IPSec policy. After all of the domain controllers receive the appropriate IPSec policy and are successfully using IPSec, you can clear this check box.

For maximum-security environments, when you must protect highly sensitive data for long periods of time, select the **Session key Perfect Forward Secrecy** check box. If you select this check box, make sure that it is selected in the IPSec policy that is assigned to each domain controller. If session key PFS is enabled on one IPSec peer and disabled on the other peer, then negotiation fails.

12. In **Manage IP filter lists and filter actions**, click **Close**.

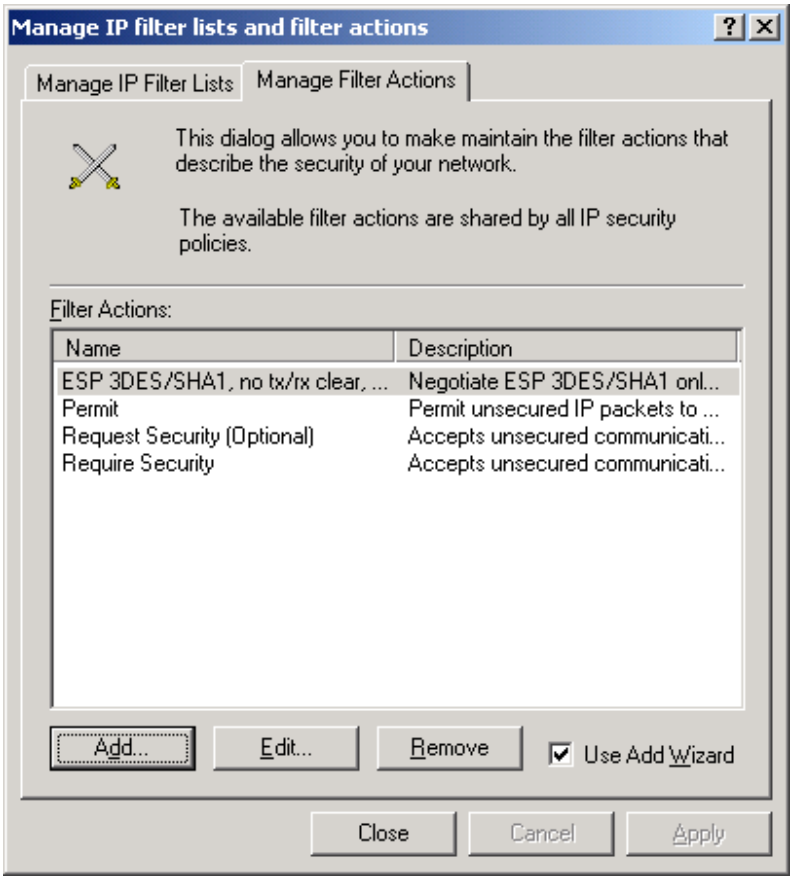

Creating an IPSec Policy and Adding a Rule to the Policy

After creating the filter list and filter actions that you want to use for your IPSec policy, you must create the policy and add a rule to combine the filter list with the filter actions. This procedure creates a policy to secure traffic between SEA-NA-DC-01 and SEA-PN-DC-01 and adds the rule as specified in the previous section, as an example.

To create an IPSec policy and add a rule to the policy

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. Right-click the details pane, and then click **Create IP Security Policy**.

3. On the IP Security Policy Wizard welcome page, click **Next**.

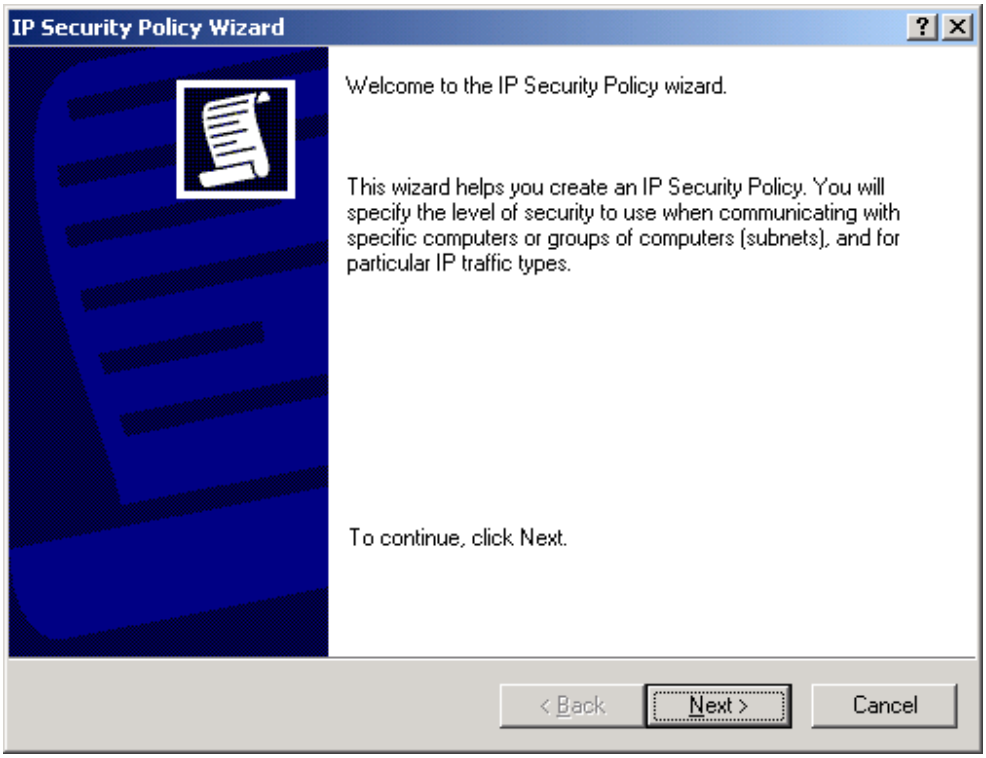

4. In **IP Security Policy Name**, type a name and a description for the policy, and then click **Next**.

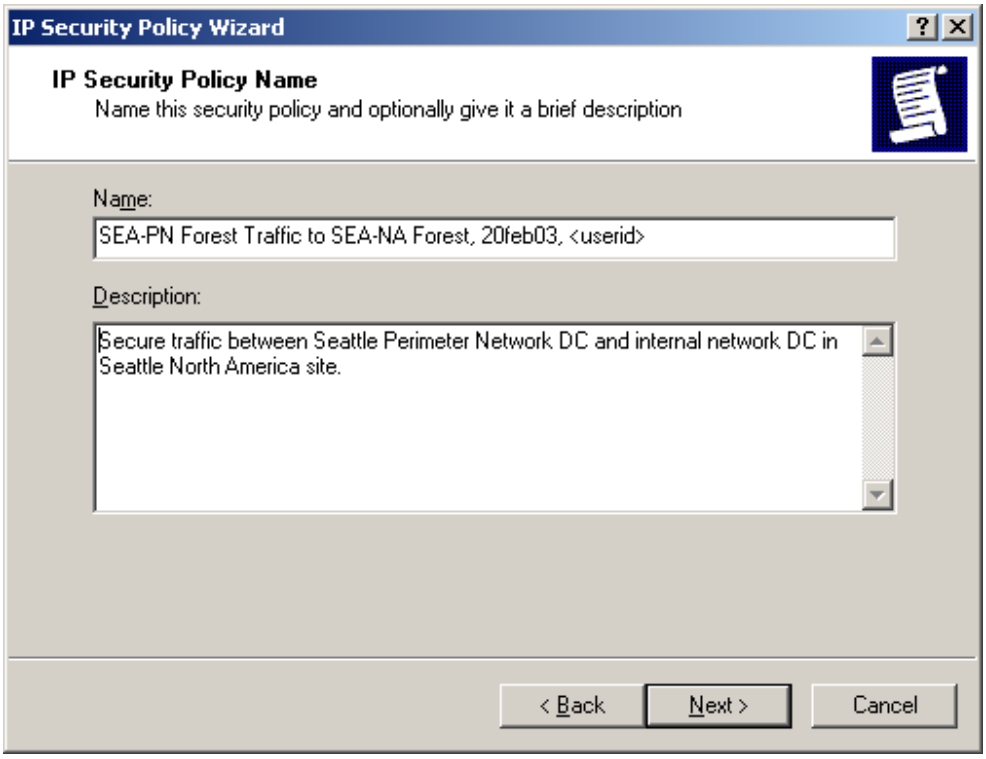

5. In **Requests for Secure Communication**, clear the **Activate the default response rule** check box, and then click **Next**.

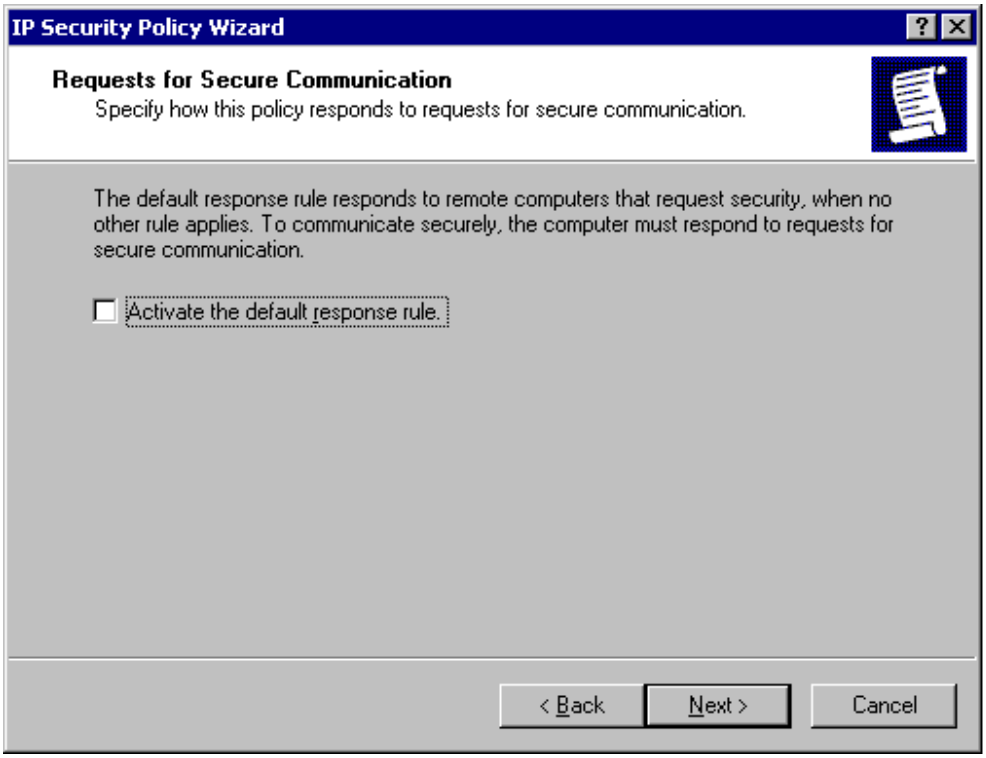

6. On the IP Security Policy Wizard completion page, verify that the **Edit Properties** check box is selected, and then click **Finish**.

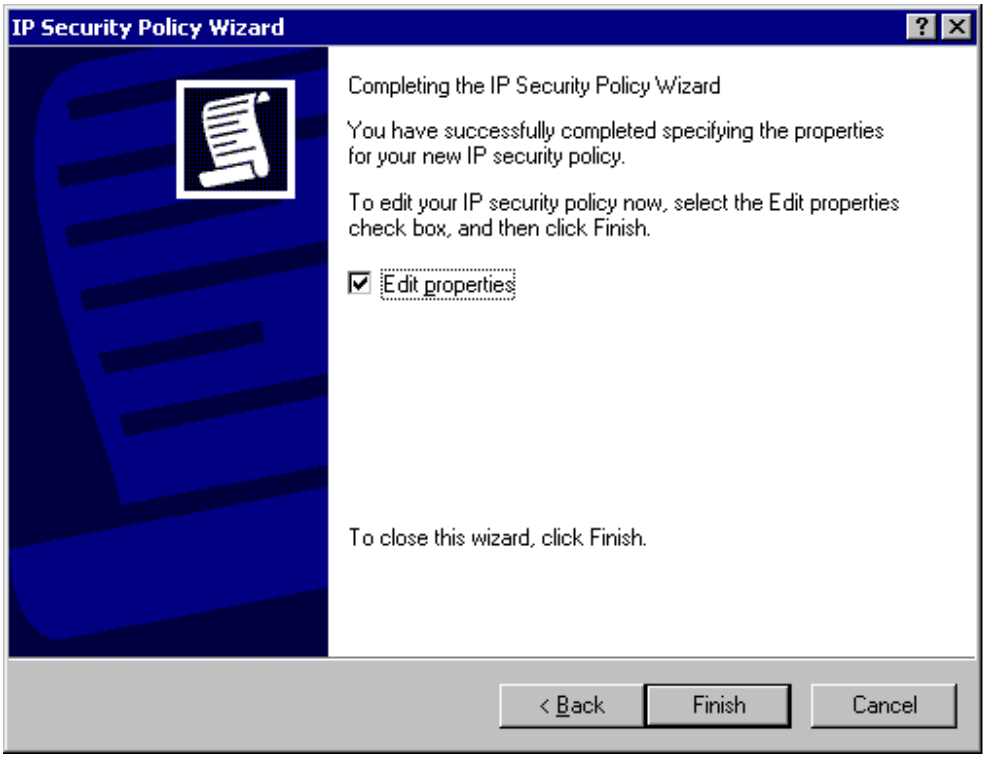

7. On the **Rules** tab, to add a new rule to this policy, click **Add**.

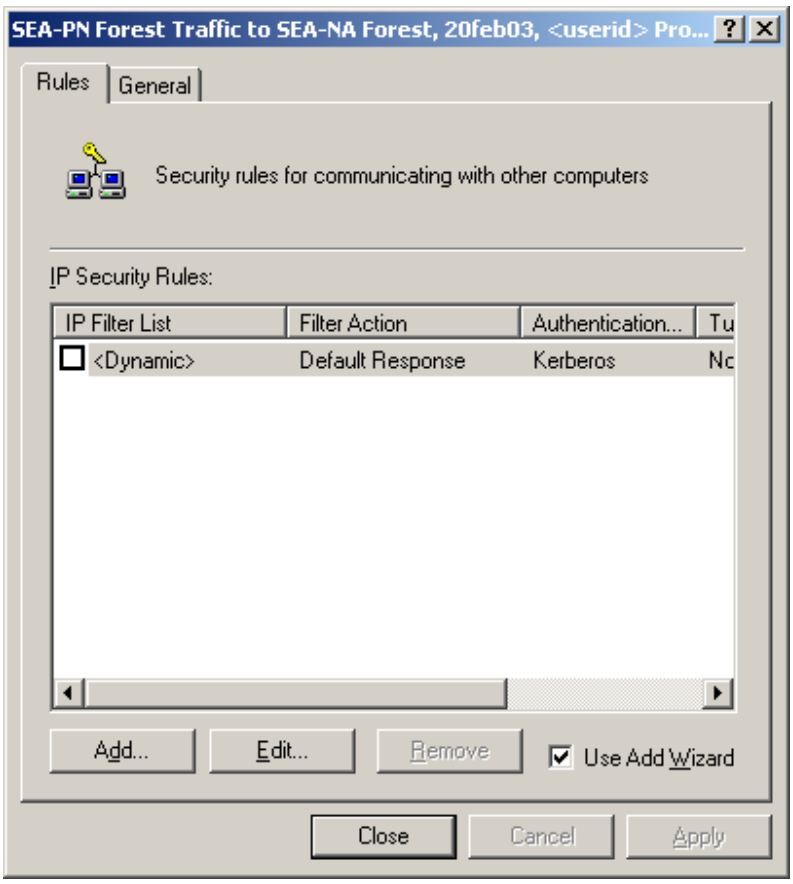

8. On the Security Rule Wizard welcome page, click **Next**.

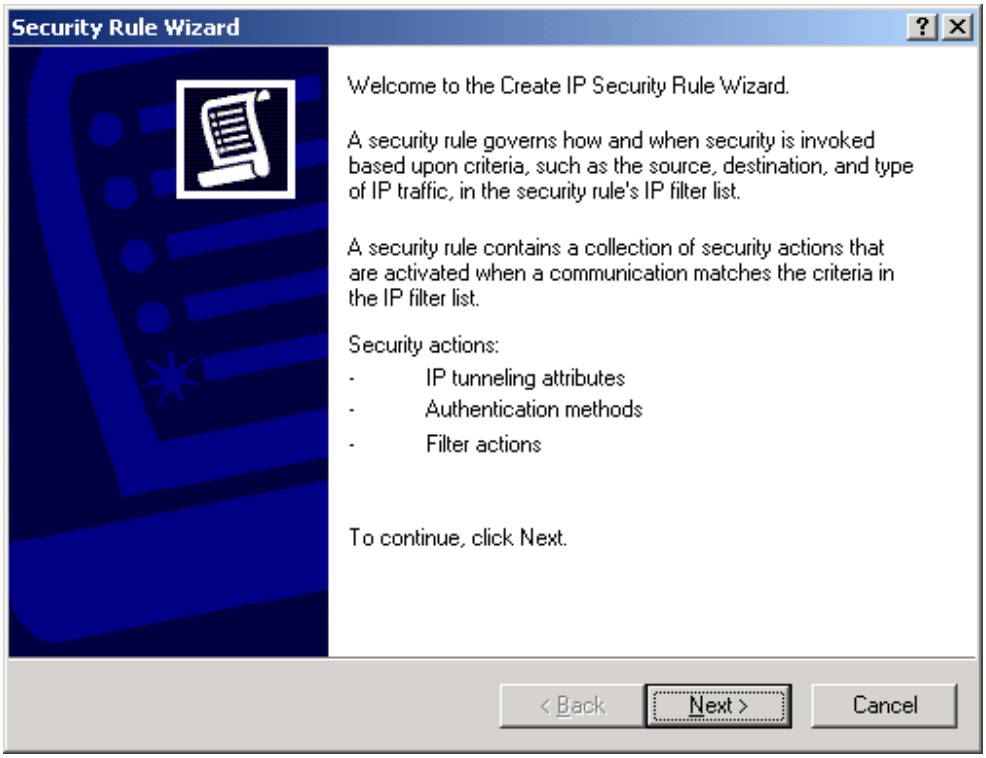

9. In **Tunnel Endpoint**, click **This rule does not specify a tunnel**, and then click **Next**.

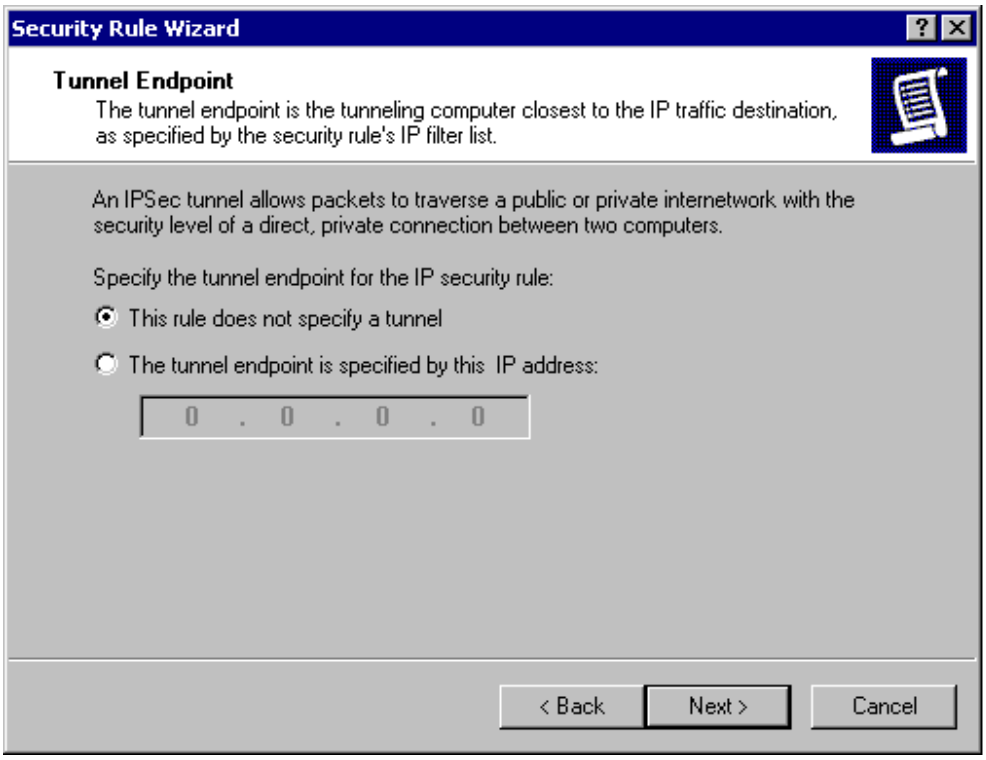

10. In **Network Type**, click **All network connections**, and then click **Next**.

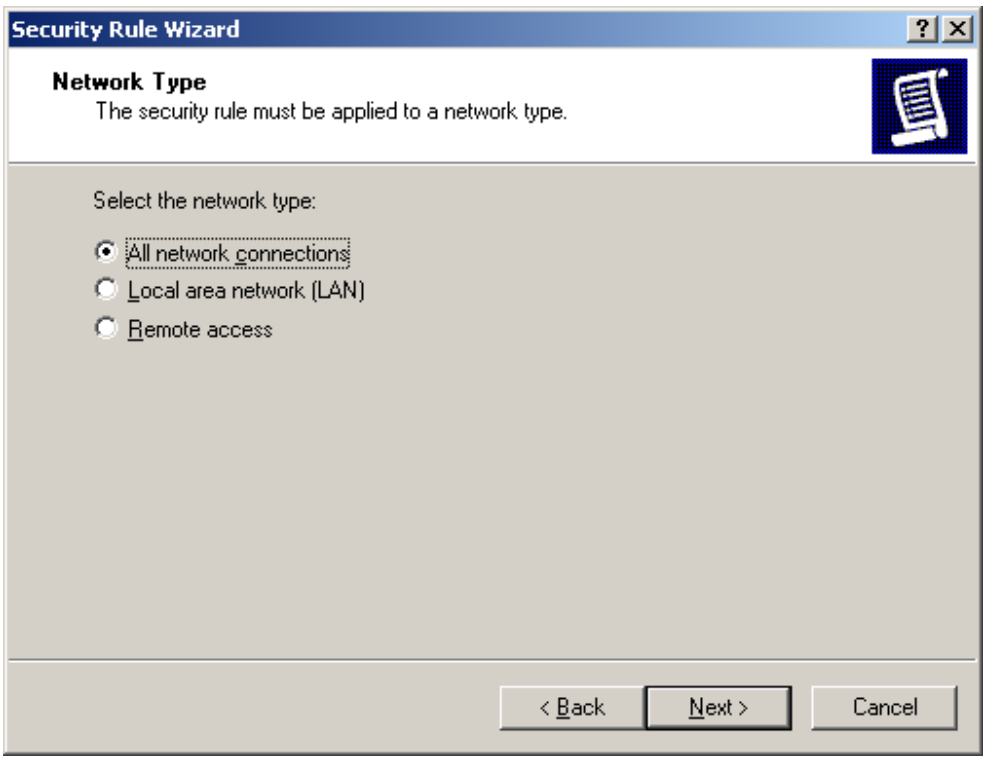

11. In **Authentication Method**, choose an authentication method to use.

    To use certificate authentication (Recommended. For more information about certificate authentication and other authentication considerations, see “Considerations for Selecting an Authentication Method,” after this procedure), click **Use a certificate from this Certificate Authority (CA)**, click **Browse** to choose the CA, and then complete steps 11a and 11b.

Note

For certificate authentication to be successful, the computer must have a certificate in the computer store that chains to this root CA. The Microsoft Root Authority CA is shown as an example.

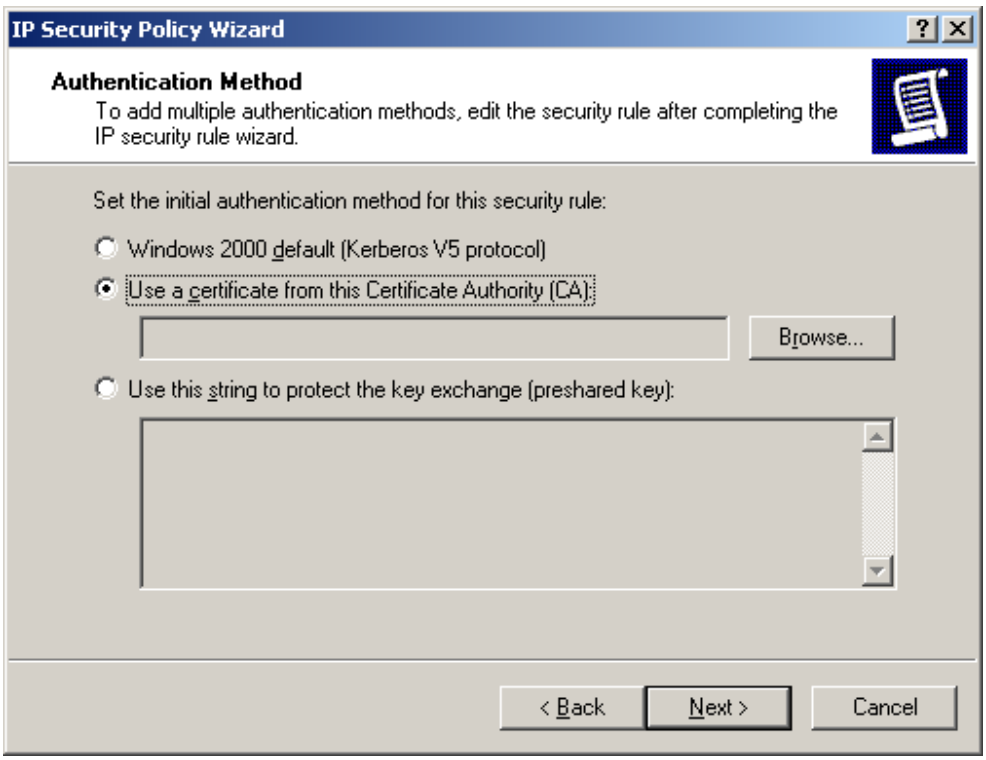

a. In **Select Certificate**, select the trusted root CA that you want to use, and then click **OK**.

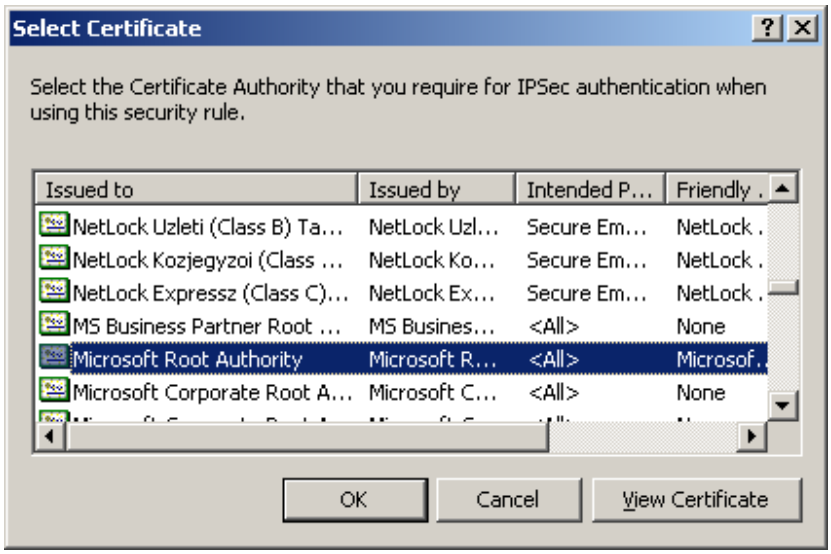

b. In **Authentication Method**, the name of the selected root CA is displayed. Click **Next**.

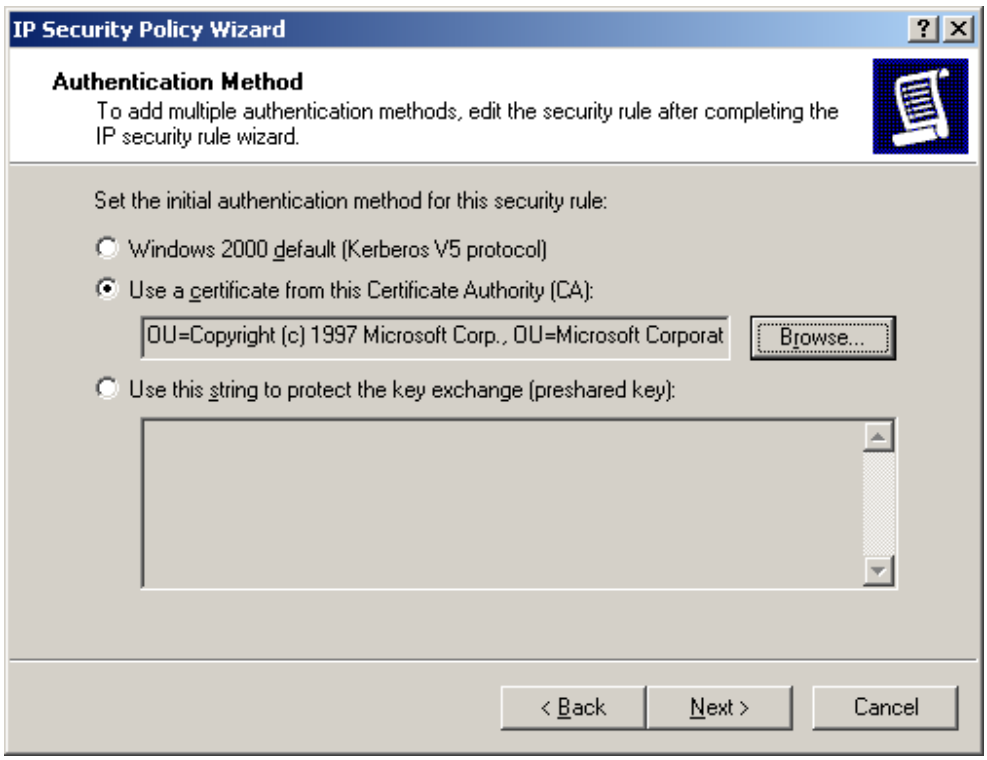

12. In **IP Filter List**, click the filter list that you created earlier (**SEA PN-DC-01<->SEA-NA-DC-01 all traffic**), and then click **Next**.

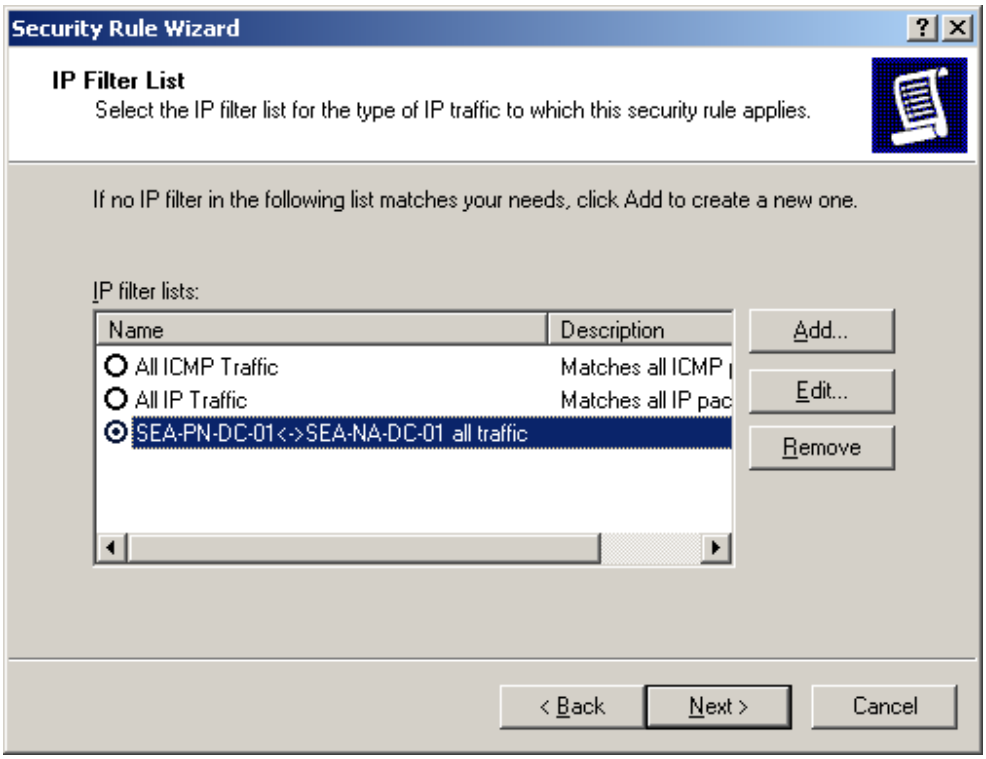

13. In **Filter Action**, click the filter action that you created earlier (**ESP 3DES/SHA1, no tx/rx clear, no pfs 20feb02 <userid>**), and then click **Next**.

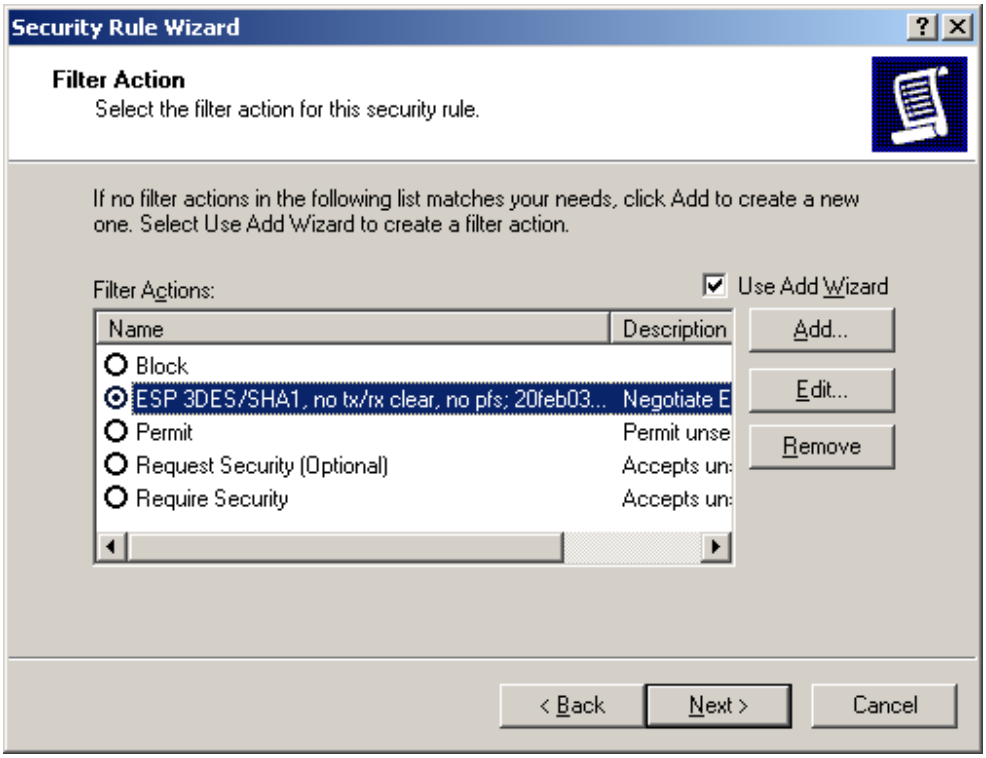

14. On the Security Rule Wizard completion page, click **Finish**.

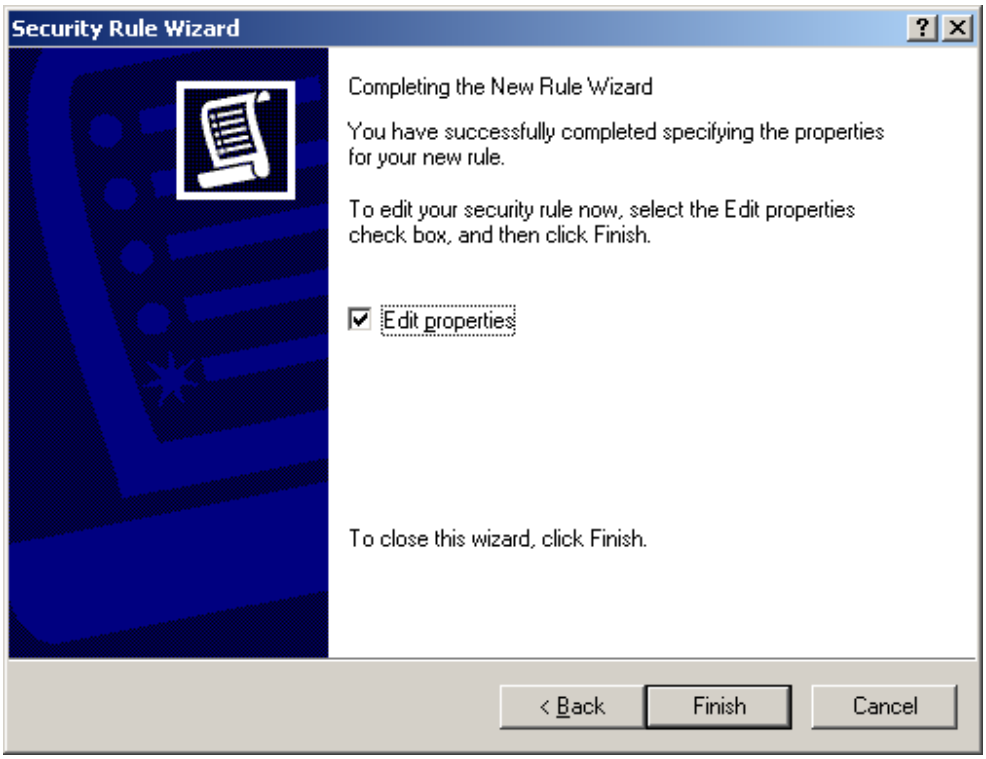

15. On the **Rules** tab, the IPSec policy is displayed as follows, with the rule that you just added. Click **Close** to close the dialog box and return to the IP Security Policy Management snap-in.

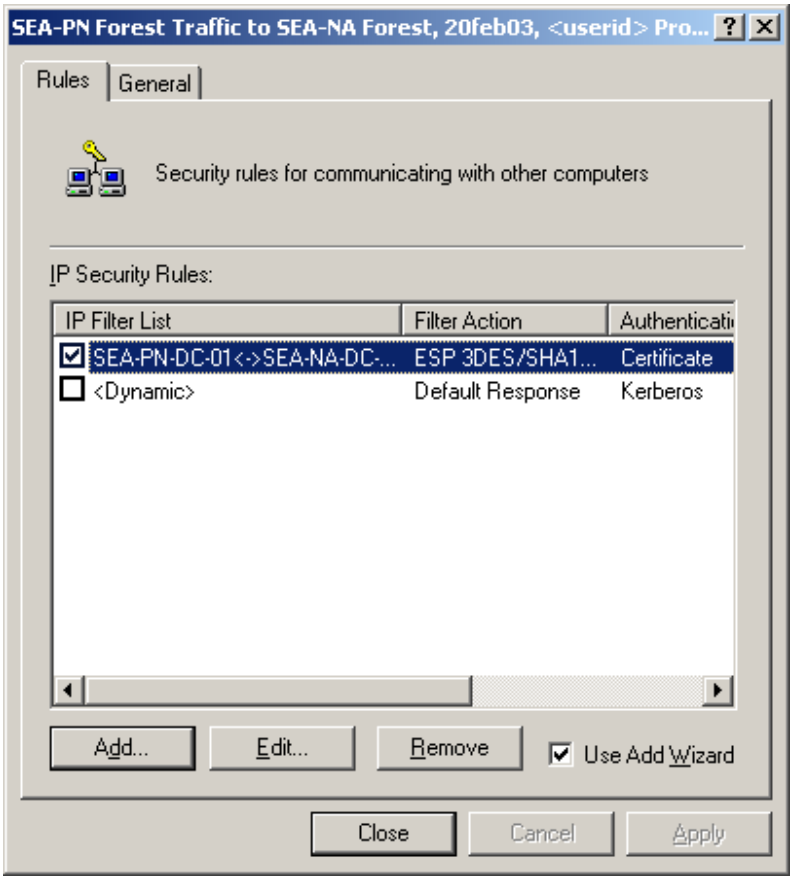

Considerations for selecting an authentication method

When security is negotiated for the traffic between domain controllers in two separate forests, the forest trust relationship is not used to establish the trust required for IKE authentication. Instead, you must configure an IPSec policy to use either certificate authentication or preshared key authentication. Although IPSec also provides Kerberos as an option for IKE authentication, using Kerberos for IKE authentication between domain controllers is not supported. It is recommended that you use certificate authentication when you configure an IPSec policy to secure traffic between domain controllers. The following list includes additional considerations for each authentication method:

- Certificate authentication

  In Windows 2000 Server, you can use Certificate Services to automatically manage computer certificates for IPSec throughout the certificate lifecycle. Certificate Services is integrated with Active Directory and Group Policy, and it simplifies certificate deployment by enabling certificate auto-enrollment and renewal and by providing several default certificate templates that are compatible with IPSec.

  The choice of a root CA is a very important security decision because IKE must use a certificate or certificate chain from this root CA to authenticate the remote domain controller. IKE authentication does not verify the name or IP address in the certificate, so it can protect only against untrusted

attacks against IPSec communication. If a computer with a certificate from this root CA were compromised, an attacker could use that computer to conduct trusted man-in-the-middle attacks on IPSec communication or to establish IPSec-secured connectivity from one of the remote domain controller's IP addresses. Likewise, if too many computers have certificates from this root CA, a trusted attack might be possible. In this IPSec policy design, any computer that uses the same IP address as the remote domain controller (or that is in the same subnet as the remote domain controller, if subnet filters are used) and that has obtained a certificate from this root CAis trusted for IPSec communication.

Upper-layer authentication protocols, such as Kerberos and RPC authentication, can provide defense-in-depth against a trusted man-in-the middle attack. Accordingly, many organizations deploy their own CA root, to ensure that only their computers have a certificate that is trusted for IPSec communication. It is recommended that you use a Windows 2000 Server issuing CA to provide certificates for IPSec because the issuing CA is available on all computers running Windows 2000 Server and because it can automatically enroll and update certificates. However, IPSec also supports the use of a variety of non-Microsoft X.509 public key infrastructure (PKI) systems. Windows 2000 IKE is compatible with many certificate systems, including those offered by Microsoft, Entrust, VeriSign, and Netscape. If you are using a non-Microsoft PKI system, the PKI system must be able to issue certificates to computers and store their certificates in the CryptoAPI computer certificate store. If you have already deployed a non-Microsoft PKI system, you can create a Windows 2000 Server issuing CA as a child CA of the non-Microsoft root CA.

For information about the use of certificates for IKE negotiation, see Step-by-Step Guide to Internet Protocol Security (IPSec), at http://go.microsoft.com/fwlink/?LinkId=269. For information about Windows 2000 Certificate Services, see the Step-by-Step Guide to Administering Certificate Services, at http://go.microsoft.com/fwlink/?LinkId=326.

- Preshared key authentication

  If certificates are not available, you can use preshared key authentication instead. To use preshared key authentication, in **Authentication Method**, click **Use this string to protect the key exchange (preshared key)**.

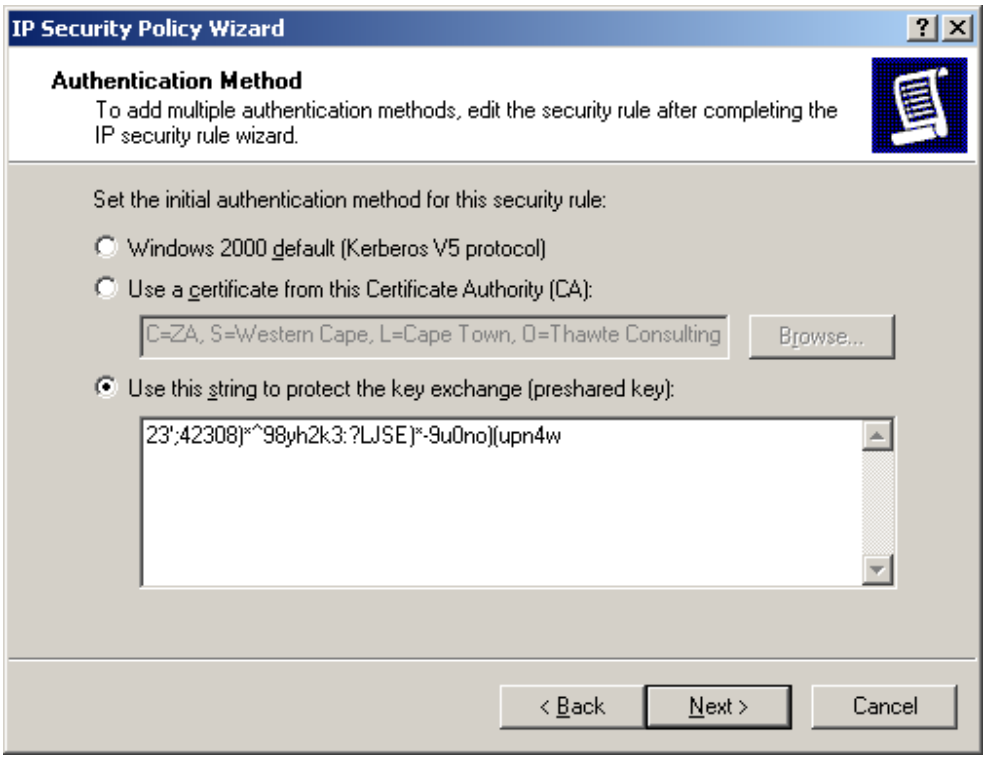

Preshared key authentication is provided for interoperability purposes and for compliance with RFC standards. However, Microsoft does not recommend the use of preshared key authentication because the key value is not securely stored, and it is therefore difficult to keep secret. The preshared key value is stored in clear text in an IPSec policy. Any member of the local Administrators group can view a local IPSec policy, and such a policy can be read by any system service with Local System user rights. Additionally, any authenticated user in the domain can view Active Directory-based IPSec policy. Even if you try to limit access to the IP Security Policies container in Active Directory by denying Read access to this container to members of the Domain Users group and by granting Read access to members of the Domain Computers group, any member of the local Administrators group on a domain computer can query Active Directory for policy information by using the Local System context. Because IPSec policies comprise many Active Directory objects in the IP Security Policies container, setting Read permissions on individual objects in this container is not recommended.

For these reasons, use preshared key authentication only for testing and when it is not possible to use certificate authentication in a production environment.

If you must use preshared key authentication, use only local IPSec policy (not Active Directory-based policy), a 25-character or longer random key value, and a different preshared key value for every IP address pair. These practices result in different security rules for each destination, and ensure that a compromised preshared key compromises only the IKE authentication between the pair of computers that share the key, rather than many IPSec communication paths.

- Kerberos authentication

  Although you can use Kerberos for IKE authentication between member computers, using Kerberos for IKE authentication between domain controllers is not supported. Because Kerberos requires that traffic be permitted over several dependent network protocols (DNS, ICMP, UDP, and LDAP) in addition to UDP and TCP ports 88, the IPSec policy configuration that is required to enable Kerberos authentication to a domain controller is very complex. Additionally, Kerberos is not available in Windows 2000 Server for authentication across forest trusts.

Accessing and Assigning an IPSec Policy

After you create an IPSec policy, you must assign it for the policy to take effect. IPSec policies can be assigned to the GPO of a site, domain, or an OU. In addition, each computer has one local GPO, which is also known as the *local computer policy*.

You can use IP Security Policy Management to create, modify, and store IPSec policies in Active Directory. To activate an IPSec policy, you must use Group Policy Object Editor to assign the IPSec policy to the OU in Active Directory that contains the server. Keep in mind that in an Active Directory domain environment, the IPSec policies that are assigned to a site, domain, or OU override the IPSec policies that are assigned to the local computer. To prevent this from occurring, deny the computer account Read access to the GPO in the Active Directory domain that is delivering the IPSec policy assignment.

The following sections describe several different methods for accessing an IPSec policy, depending on whether the IPSec policy is Active-Directory based or local, and then how to assign an IPSec policy.

You can access Active Directory-based IPSec policy by doing any of the following on the domain controller from which you want to manage policy:

- Start IP Security Policy Management in Domain Controller Security Policy.
- Start IP Security Policy Management from the Domain Controllers OU in Active Directory (Group Policy).
- Add IP Security Policy Management for Active Directory-based IPSec policy to MMC.

You can access local IPSec policy for a domain controller by doing either of the following on the domain controller for which you want to manage policy:

- Add IP Security Policy Management for local IPSec policy to MMC.
- Add Group Policy Object Editor for local IPSec policy to MMC.

To start IP Security Policy Management in Domain Controller Security Policy

1. On the domain controller from which you want to manage IPSec policy, click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Domain Controller Security Policy**.

   The MMC window that opens displays the default domain controller GPO that is associated with the domain controller's OU in Active Directory.

2. In the console tree, click **Windows Settings**, click **Security Settings**, and then click **IP Security Policies**.

When you open the security policy settings in Domain Controller Security Policy, the graphical user interface (GUI) for IP Security Policies appears as follows:

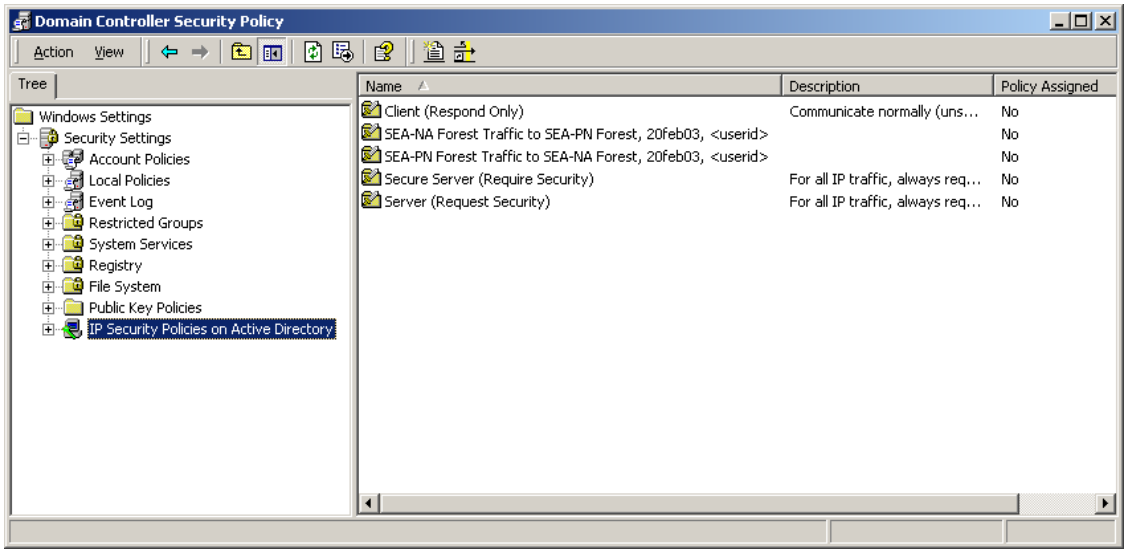

To start IP Security Policy Management from the Domain Controllers OU in Active Directory (Group Policy)

The default GPO that is associated with a domain controller is assigned to the Domain Controllers OU in Active Directory.

1. On the domain controller from which you want to manage policy, click **Start**, click **Control Panel**, double-click **Administrative Tools**, double-click **Active Directory Users and Computers**, and then navigate to **Domain Controllers**.

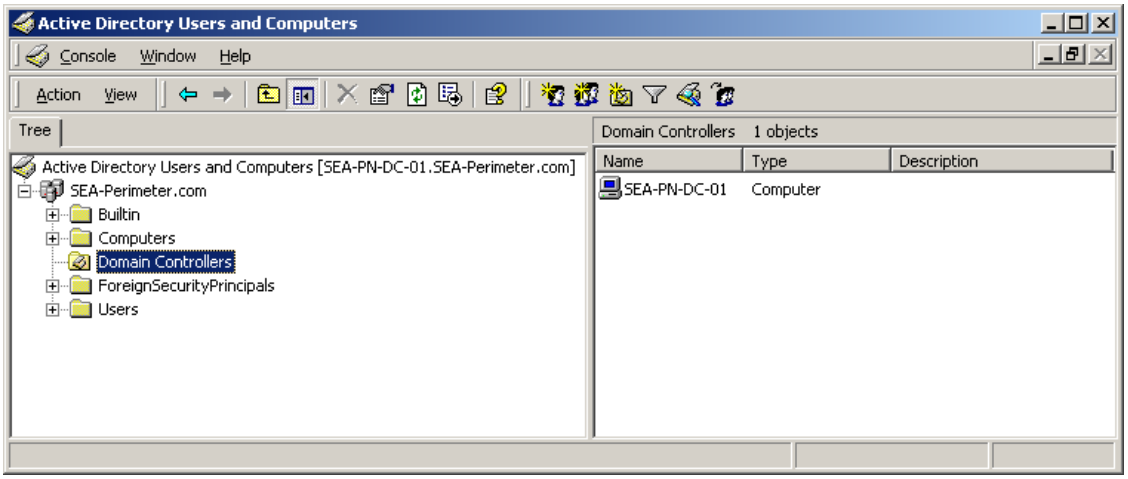

2. In the console tree, right-click **Domain Controllers**, and then click **Properties**.

3. In **Domain Controllers Properties**, click the **Group Policy** tab, click the GPO that you want to edit, and then click **Edit**.

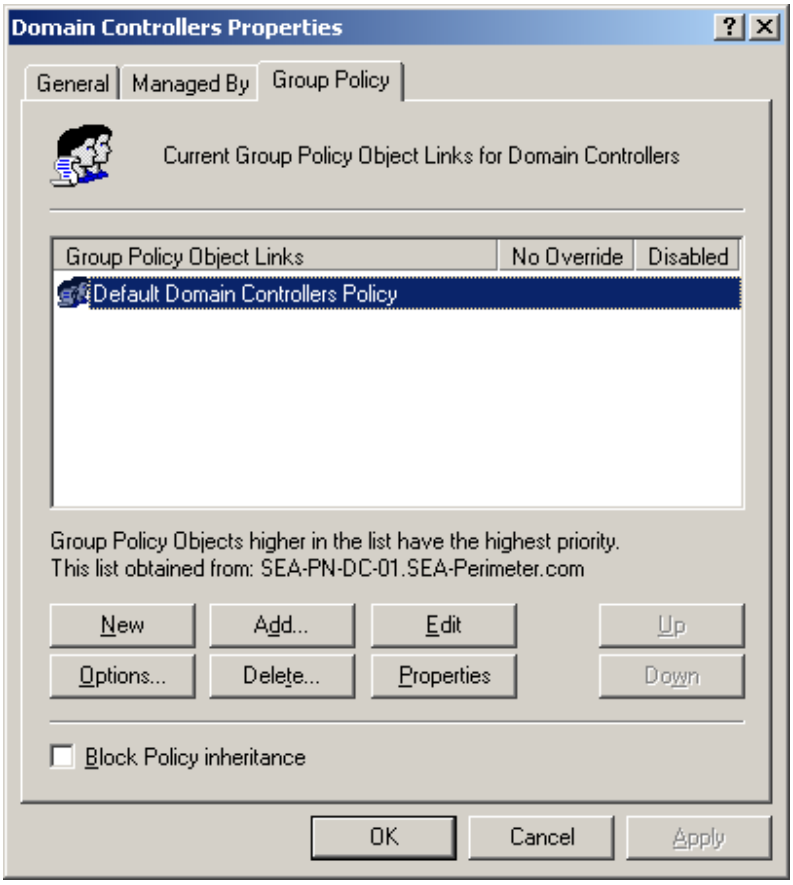

The GPO that you want to edit (in this example, the domain controller for the internal network domain, SEA-Perimeter.com), appears as follows:

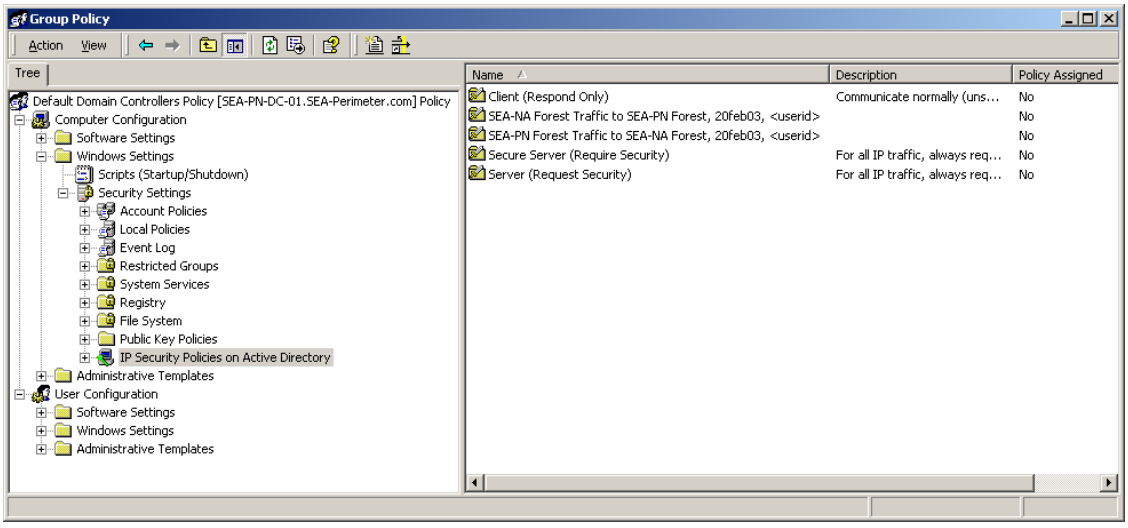

To add IP Security Policy Management for Active Directory-based IPSec policy to MMC

1. Click **Start**, click **Run**, type **MMC**, and then click **OK**.
2. In **MMC**, click **Console**, click **Add/Remove Snap-in**, and then click **Add**.
3. Click **IP Security Policy Management**, and then click **Add**.
4. Click **Manage domain policy for this computer's domain**.

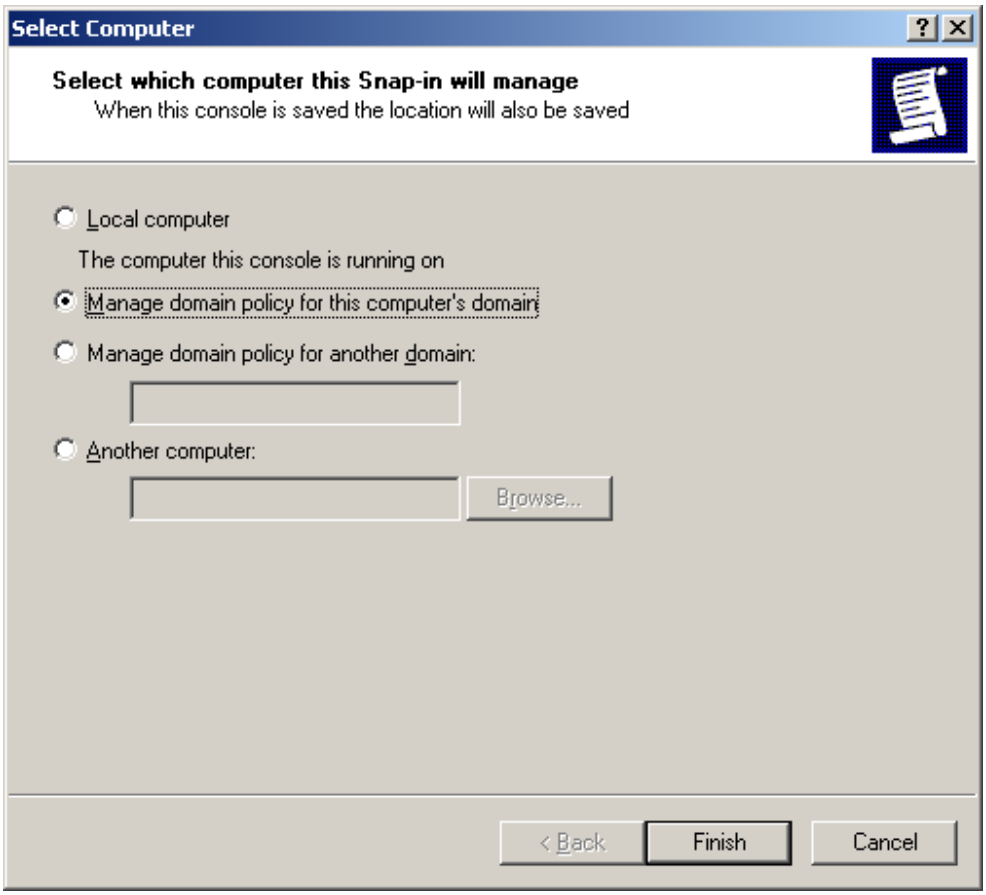

If you access Active Directory-based IPSec policy both by opening Active Directory Users and Computers and adding IP Security Management to this console and by managing domain policies for the local domain or a remote domain, then the GUI appears as follows:

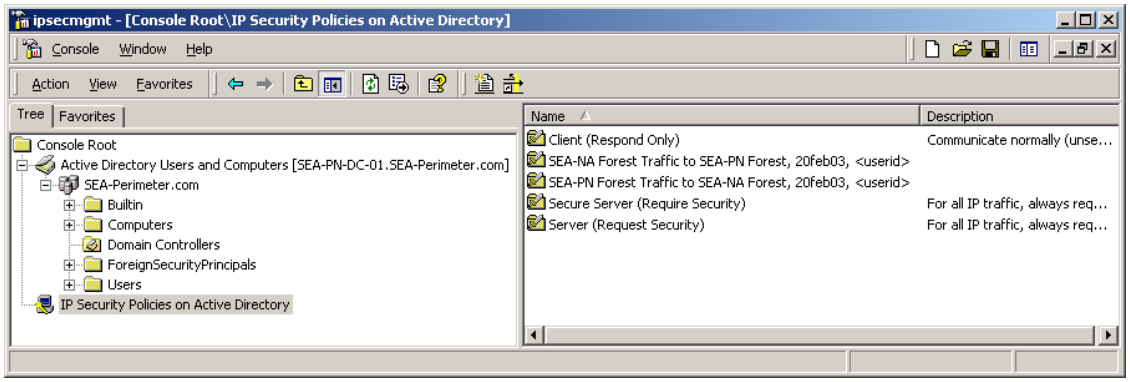

Note

To view or modify Active Directory-based policy settings, you must be logged on to the computer with the credentials required to access Active Directory-based IPSec policy on that computer or on remote computers. You cannot specify different credentials to accomplish this task.

To add IP Security Policy Management for a local IPSec policy to MMC

Do one of the following on the domain controller for which you want to manage policy:

- Add IP Security Policy Management to MMC as a stand-alone snap-in. The GUI appears as follows:

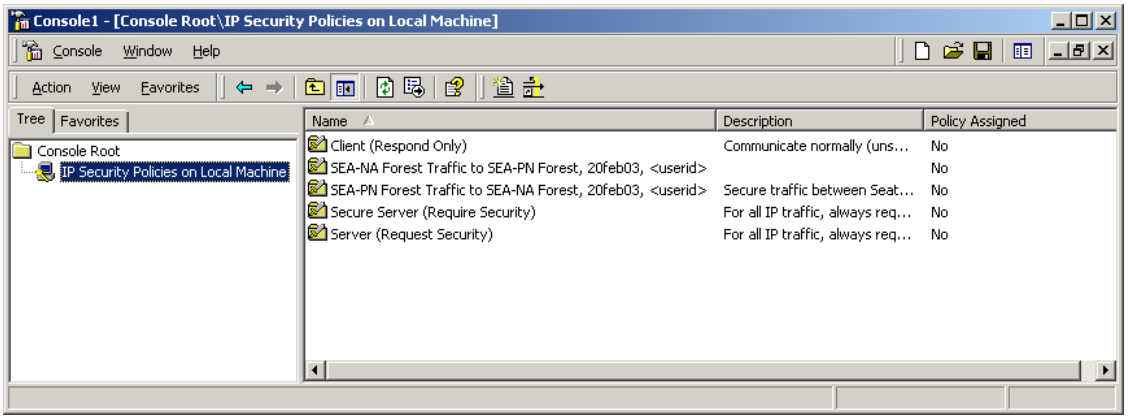

- Add Group Policy Object Editor to MMC, choose Local Computer Policy as the GPO, and then navigate to **IP Security Policies on Local Machine**. The GUI appears as follows:

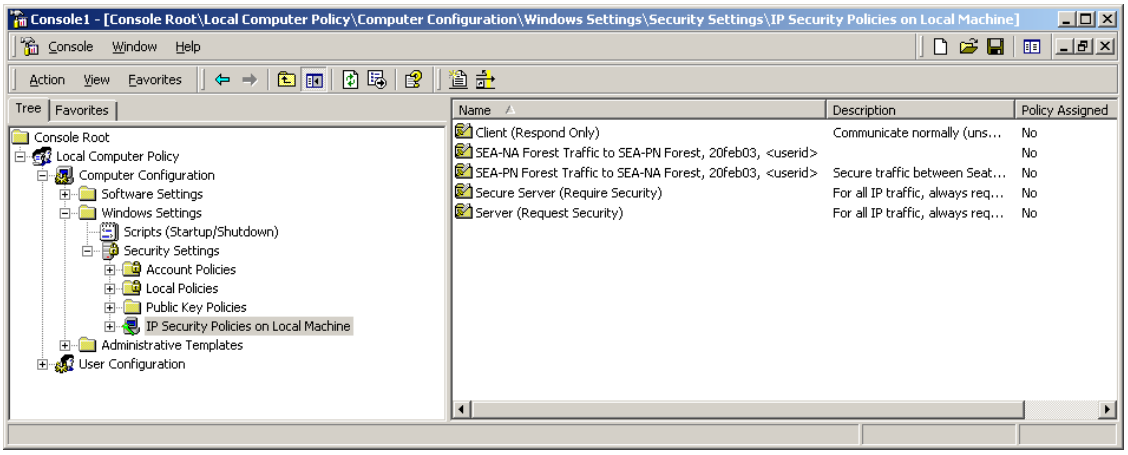

To assign an IPSec policy

1. Access the IPSec policy by using one of the methods described in this section.
2. In the details pane, right-click the policy that you just created (for this example, **SEA-PN Forest Traffic to SEA-NA Forest, 20feb03 <userid>**), and then click **Assign**.
3. In the console tree, right-click **IP Security Policies**, and then click **Refresh**.

Note

You cannot assign an Active Directory-based policy from the IP Security Policies on Active Directory console; you can only configure the policy. To assign Active Directory-based policy, you must use Group Policy Object Editor.

Exporting and Importing IPSec Policies

It is often helpful to dedicate one computer to use local IPSec policy. You can export the IPSec policy from this computer to a file, and then import the file into the local IPSec policy store on another computer.

To export and import Active Directory-based IPSec policies, start IP Security Policy Management from MMC or access IP Security Policy Management from Group Policy (Active Directory), select the appropriate Active Directory-based IPSec policy, and then follow the procedures in this section.

The **Export Policies** menu command exports all IPSec policy objects from the policy store into one .ipsec file. Each policy object is internally identified by a globally unique identifier (GUID). You can also use the **Export Policies** menu command to back up an IPSec policy in the event that you need to recover the policy later.

Importing an IPSec policy file overwrites existing IPSec policy objects that have the same GUID or creates new IPSec policy objects if they do not exist. If you change only the name of an IPSec policy, filter list, or filter action, and then re-import the policy file, then these IPSec policy objects are overwritten with the new names. However, if you select the **Delete all existing policy information** check box when you select the policy file to import, then all policies, filter actions and filter lists in the target store are deleted before the objects in the export file are written into the target policy store. Note that you cannot export or import a single IPSec policy, filter list, or filter action.

Caution

If you reimport changes to an IPSec policy that is already assigned to a GPO, the IPSec policy is unassigned. You must edit the GPO and reassign the IPSec policy after the import is complete.

Additionally, when importing or editing IPSec policy in Active Directory, do not close IP Security Policy Management before all of the IPSec policy data is written to Active Directory. If IP Security Policy Management cannot finish writing all of the policy data into Active Directory, then IPSec policy corruption might result. If you detect IPSec policy corruption, you can try to reimport the IPSec policy file. In some cases, the IPSec policy objects must be deleted so that a new IPSec policy import operation can be successfully completed. You can use either LDP.exe or ADSIedit.exe to delete the IPSec policy objects. LDP.exe is a Windows Support Tool that is included in the Support Tools folder of the Windows 2000 operating system CD. ADSIedit.exe is a Windows Support Tool that uses the Active Directory Service Interfaces (ADSI) and that is also provided on the Windows 2000 operating system CD. If you are managing IPSec policy remotely over slow links, then use a file copy technique to transfer the IPSec policies in .ipsec export files, before you delete the IPSec policy objects. After you transfer the IPSec policies, use the Terminal Services client to connect to the remote server and perform the import operation quickly.

To export local IPSec policies

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.

2. In the console tree, right-click **IP Security Policies on Local Machine**, point to **All Tasks**, and then click **Export Policies**.
3. In **Save As**, specify where to save the .ipsec policy file, and then click **Save**.

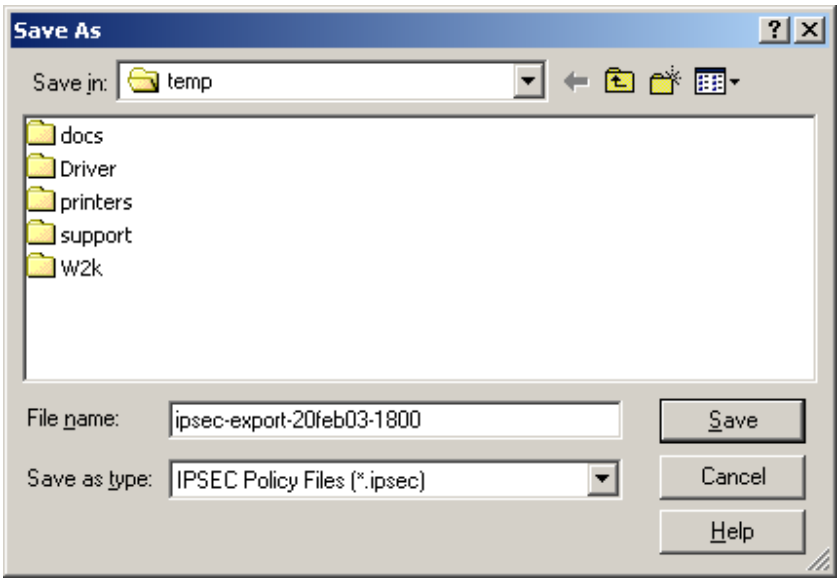

Note

Make sure to properly name and secure access to the file.

To import local IPSec policies from a file

1. Create a console containing IP Security Policies. Or, open a saved console file containing IP Security Policies.
2. In the console tree, right-click **IP Security Policies on Local Machine**, point to **All Tasks**, and then click **Import Policies**.

3. In **Open**, specify the .ipsec policy file, and then click **Open**.

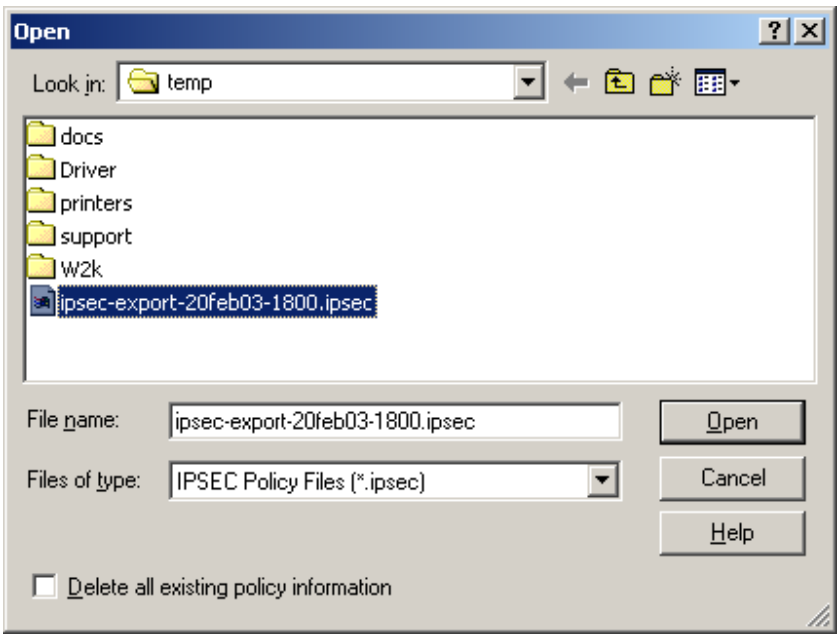

Considerations for Updating Active Directory-Based IPSec Policy

When you plan updates to Active Directory-based IPSec policy, make sure that the IP addresses that are used in the IPSec policy are consistent with the IP addresses that are used in the domain. Coordinate IPSec policy updates when you add new domain controllers to the domain so that the communications that the domain controller requires are not blocked. For example, you might need to import a local IPSec policy and certificate on a new domain controller to enable successful IPSec-secured communications when the domain controller first joins the domain, when the domain controller is promoted (using DCPromo), and when Active Directory synchronization operations occur.

Verify that successful IPSec security associations are established where expected. If Active Directory replication failure events occur, confirm that there are no communication problems first. In the IPSec policy example in this appendix, ICMP traffic is protected by IPSec so that you can use the **ping** command to verify that IKE and IPSec traffic is not blocked on any communication paths.

Keep in mind that changes to Group Policy assignments and IPSec policy settings might take effect at different times. Group Policy is used only to deliver the IPSec policy assignment to the domain controller's IPSec service. Assigning an IPSec policy to a GPO records the LDAP distinguished name (also known as DN) of the IPSec policy that is inside the GPO attribute **ipsecOwnersReference**. The IPSec service retrieves the assigned IPSec policy from Active Directory, maintains a current cache of the policy in the local registry of the domain controller, and updates the assigned policy settings by using a polling interval that is specified in the IPSec policy. Therefore, if you change IPSec policy settings on the domain controller, the changes are applied only when the IPSec service on the domain controller polls for changes.

While the IPSec service polls for changes in IPSec policy settings, the Winlogon service polls for changes in policy assignment within GPOs. Because a domain controller might detect changes in the settings of its assigned IPSec policy and changes in IPSec policy assignment within a GPO at different times, make sure that you carefully coordinate IPSec policy changes. If IPSec is used to protect traffic between

domain controllers in the same domain, incompatible IPSec policies might cause communication (for example, the sending and receipt of replication traffic) between domain controllers to fail. Therefore, if you plan to either change settings for an existing IPSec policy that is assigned to a domain controller or to assign a new IPSec policy to a domain controller, plan and test these changes to ensure compatibility with the IPSec policy used by other domain controllers. Consider creating an IPSec policy first and not assigning it to a GPO, so that the policy can be replicated throughout the domain. After you verify that the policy is replicated throughout the domain, assign the policy to a GPO. This strategy ensures that the IPSec policy exists in the domain before domain controllers receive the GPO assignment.

Scripting IPSec Policy

The primary method of creating and managing IPSec policy is by using the IP Security Policy Management console. To script the creation of IPSec policy, you can use Ipsecpol.exe, a command-line tool that is provided with the Windows 2000 Server Resource Kit. However, you can use Ipsecpol.exe only to create an IPSec policy. Ipsecpol.exe is not a full-featured command-line or scripting tool (for example, you cannot use Ipsecpol.exe to delete or rename filter lists or filter actions), nor is it supported under any Microsoft standard support program or service. You can download the latest version of Ipsecpol.exe from Ipsecpol.exe: Internet Protocol Security Policies Tool, at http://go.microsoft.com/fwlink/?LinkId=16466.

An Ipsecpol.exe script is provided in Appendix B as an example that you can use to create an IPSec policy for the perimeter network domain and for the internal network domain controller used as an example in this appendix. For another example of how to use Ipsecpol.exe, see article 813878, “How to Block Specific Network Protocols and Ports by Using IPSec,” in the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462.

Note

In Windows Server 2003, the Netsh IPSec command-line tool provides a full-featured and Microsoft-supported interface. Ipsecpol.exe does not work on computers running Windows Server 2003.

IPSec Policy Compatibility Considerations

Windows IPSec policies are compatible across the following operating systems: Windows 2000, Windows XP, and Windows Server 2003. For example, you can export IPSec policies that you create by using Ipsecpol.exe from computers running Windows 2000 and then import these policies into Windows Server 2003 Active Directory or local IPSec policy stores on computers running Windows Server XP or Windows 2003. Additionally, you can assign Windows 2000 Active Directory-based policy to computers running Windows 2000, Windows XP, or Windows Server 2003. Any IPSec policies that you create on computers running Windows XP and Windows Server 2003 can be stored in Windows 2000 Active Directory or exported and then imported for use on computers running Windows 2000, if the policies use the basic IPSec features supported by Windows 2000.

However, if you plan to apply IPSec policies that use any of the new features that are available only in the Windows Server 2003 implementation of IPSec, do not assign these policies to computers running Windows XP or Windows 2000. Also, do not use the IP Security Policy Management snap-in in Windows 2000 or Windows XP to manage policies that use features that are new for Windows Server 2003. If you use Windows 2000 or Windows XP to manage Windows Server 2003 policies, any new Windows Server 2003 features are lost.

To ensure that the IPSec policy functions consistently and as expected on computers running the Windows Server 2003 family and on computers running Windows XP or Windows 2000, test the policy thoroughly on all relevant operating systems before deployment.

Important

It is recommended that you use a computer running Windows Server 2003 as an IPSec policy management station, due to the enhanced scripting support provided by the Netsh IPSec command-line tool and the availability of IP Security Monitor, which allows you to view filter details and other information about the assigned IPSec policy.

Performance and Troubleshooting Considerations

The following performance and troubleshooting considerations can help simplify the administration of IPSec policies.

Using IPSec Hardware Offload Network Adapters

Hardware offload network adapters can accelerate IPSec processing by performing hardware offload of IPSec cryptographic functions. For information about the 10/100 MB hardware offload network adapters that are available for Windows 2000 IPSec, see the following:

- The Intel Web site, at http://go.microsoft.com/fwlink/?LinkId=16474. The Pro 100 S series of hardware offload adapters can be used for Windows 2000 IPSec hardware offload.
- The 3Com Web site, at http://go.microsoft.com/fwlink/?LinkId=16475. 3Com network adapters with 3XP processors can be used for Windows 2000 IPSec hardware offload.

For information about additional network cards that are compatible with Windows 2000, see Search for Compatible Hardware Devices, at http://go.microsoft.com/fwlink/?LinkId=3787.

Viewing IPSec and Other Network Communication with Network Monitor

The Network Monitor component that is provided with Windows 2000 does not include a parser for the IPSec ESP protocol. However, Network Monitor in Windows Server 2003 contains a parser for ESP traffic. The Windows Server 2003 Network Monitor parser can interpret ESP traffic if this traffic is not encrypted or decrypted in software (that is, if encryption or decryption is performed by an IPSec hardware acceleration adapter, or if ESP without encryption is used). Therefore, if you need to use Windows Server 2003 Network Monitor to troubleshoot network traffic flows, change the filter actions in your IPSec policy so that the AH protocol is used for security negotiation, rather than ESP. For ease of configuration, when you create an IPSec policy to secure traffic between two domain controllers, consider configuring AH as the second security method in the filter action that is used by both domain controllers. That way, you need to change only one domain controller policy filter action to AH. If one domain controller is configured to use only AH, and the other domain controller is configured to use only ESP, security negotiation fails.

If you plan to switch from using ESP to AH for troubleshooting purposes, remember to ensure that your firewall is configured to forward traffic on protocol 51.

Evaluating Bad SPI Events

Bad SPI events indicate the total number of packets for which the Security Parameters Index (SPI) was incorrect. The IPSec driver records Bad SPI events in the Event Viewer System log when it receives an

IPSec-formatted packet that it cannot interpret. These events, which appear in the System log as Event 4283: “<number> packets discarded due to bad SPI,” are usually benign.

You can expect a small number of Bad SPI events per each IPSec peer IP address, due to the way in which IKE processes the transition of IPSec SAs during rekeys. This event is usually logged when an SA has been deleted on one IPSec peer, and the other IPSec peer still sends secure traffic on the deleted SA. This event is also logged when one IPSec peer begins sending IPSec-secured traffic using a new SA that the other peer is not yet ready to receive. These brief periods between rekeys typically do not cause problems for traffic over upper-layer protocols, which can accommodate the loss of a few packets. This number is likely to increase if rekey intervals are short and there are a large number of SAs. Windows Server 2003 IPSec has been enhanced to eliminate packet loss during the periods between rekeys.

You can use IP Security Monitor to examine the number of rekeys. If the number of rekeys is very large compared to the amount of time that the connections have been active, consider setting the key lifetimes in the IPSec policy to be longer.

Note that you cannot disable logging for Bad SPI events.

Evaluating Events Generated by Automatically Starting Services during Computer Startup

Some services, such as DNS, start automatically on a domain controller and are designed to communicate immediately with other domain controllers during computer startup. When an IPSec policy is assigned to a domain controller, these services might record events in Event Viewer that indicate that the remote domain controller cannot be reached. These events are expected during computer startup because the IPSec service starts at the same time as many other automatically starting services. If other services are configured to depend on the IPSec service during computer startup, the IPSec service does not delay these services. IPSec-secured connectivity is established as soon as the IPSec service starts, IPSec policy is applied, and outbound traffic triggers a security negotiation. However, if such events are recorded during normal operation, you might need to troubleshoot IPSec connectivity.

Security Considerations

When you configure an IPSec policy to secure traffic between domain controllers, keep the following security considerations in mind.

Configuring Domain Controller Baseline Security Option Settings

IPSec does not replace the need to implement additional measures to enhance the security of domain controllers. Such measures play a critical role in defending your network against attacks that use upper-layer protocols and services and attacks mounted from remote domain controllers that are compromised. Keep in mind that IPSec provides a limited defense against attacks from domain controllers that can successfully authenticate and negotiate security. For information about how to enhance the security of domain controllers and other computers, see Chapter 7, “Hardening Specific Server Roles,” in Securing Windows 2000 Server, at http://go.microsoft.com/fwlink/?LinkId=15394.

Combining IPSec Policy Configurations

In addition to configuring an IPSec policy to negotiate security for traffic between domain controllers, as described in this appendix, you can configure an IPSec policy that uses IPSec blocking filters to enhance the security of domain controllers. For information about defining specific IPSec filters to permit or block traffic to domain members, see Table 7.5, “Domain Controller IPSec Network Traffic Map,” in Chapter 7,

"Hardening Specific Server Roles," in Securing Windows 2000 Server, at http://go.microsoft.com/fwlink/?LinkId=15394. You can combine IPSec blocking filters with the IPSec policy described in this appendix. To implement both IPSec policies, combine the rules for the two policies into one policy. The IPSec policy in this appendix uses more specific filters than IPSec blocking filters to require that IPSec negotiate security for all traffic between the source and destination IP addresses of the two domain controllers. Note that the filters that negotiate security between the specific IP addresses of the two domain controllers, as described in this appendix, will override the more general filters used for domain member traffic described in Table 7.5 of "Securing Windows 2000 Server."

Security During Computer Startup

In Windows 2000, the IPSec service does not provide security during computer startup. As a result, IPSec does not protect traffic that is sent or received during startup. However, if you configure a firewall between two separate forests, as described in this appendix, then the firewall should block unprotected traffic between domain controllers in the forests. After the IPSec service starts, the traffic will match IPSec filters and trigger security negotiations.

In the Windows implementation of IPSec, the IPSec service starts automatically, by default. A delay in any service that starts automatically can delay all other dependent services that also start automatically. Therefore, to minimize the time required for a computer to start and to allow for IPSec policy to be applied as quickly as possible during this time, the IPSec service is configured by default to not have other dependent services. If the IPSec service is configured to have a dependent service, however, the IPSec service might be unable to apply the assigned IPSec policy before the dependent service is notified that the IPSec service is ready.

With Windows Server 2003 IPSec, you can ensure that IPSec provides security during computer startup through a combination of two methods. First, configure a computer startup mode (also known as the *bootmode*) for the IPSec driver so that traffic can be filtered before the IPSec service starts. To configure the computer startup mode for the IPSec driver, you must use the **netsh ipsec dynamic set config bootmode** command. Second, create and assign a *persistent IPSec policy*. Persistent policy secures a computer even if a local or Active Directory-based IPSec policy cannot be applied. After the IPSec service starts and applies the persistent policy, the IPSec service notifies the IPSec driver to move out of startup mode and to apply the persistent policy. To configure persistent policies, you must use the **netsh ipsec static set store location=persistent** command. Because persistent policy adds to or overrides the local or Active Directory-based policy, and it remains in effect regardless of whether other policies are applied, you should use persistent policies to secure traffic between domain controllers only when:

- Computer startup security is required before local or Active Directory-based IPSec policy is applied.
- Faster IPSec-secured connectivity with a remote domain controller is required during computer startup. You can eliminate minor delays in the application of local or Active Directory-based IPSec policy by using persistent policy.
- You can use Netsh IPSec command-line scripts to manage policy configuration on domain controllers, rather than local or Active Directory-based IPSec policy.

For more information about Netsh IPSec, see "Netsh commands for Internet Protocol security" in Help and Support Center for Windows Server 2003.

Security during Safe Mode with Networking and Directory Services Restore Mode

In Windows 2000, the IPSec service does not start when the computer is started in Safe Mode with Networking or when Directory Services Restore Mode is used. Accordingly, IPSec does not protect traffic that is sent or received during these times. However, if you configure a firewall between the domain controllers, as described in this appendix, and if the firewall allows only IPSec traffic, then connectivity with the remote domain controller is blocked by the firewall and by the IPSec policy configuration on the remote domain controller.

In Windows Server 2003, the IPSec service also does not start when either Safe Mode with Networking or Directory Services Restore Mode is used. However, if you have already assigned a persistent IPSec policy or a local or Active Directory-based IPSec policy to a computer, by default, the IPSec driver provides stateful filtering of inbound traffic to that computer. For more information about using IPSec during Safe Mode with Networking, see Chapter 6, “Deploying IPSec” in *Deploying Network Services* in the Windows Server 2003 Deployment Kit, on the Web at http://go.microsoft.com/fwlink/?LinkId=8195.

Using IPSec Filters to Secure Traffic between Domain Controllers over Specific Protocols and Ports

For the domain trust scenario described in this appendix, you can configure IPSec filters to secure traffic over only the protocols and ports that are required for communication between domain controllers, rather than using IPSec to secure all traffic between domain controllers. You cannot do this, however, when RPC is required for communication between domain controllers because the port numbers are not statically defined for all RPC programs.

When a domain trust is configured, the following protocols and ports are required for communication between domain controllers:

- Kerberos (88/tcp, 88/udp, but not used between domain controllers if an explicit trust between two separate forests is configured)
- LDAP (389/udp, 389/tcp, and/or 636/tcp if using LDAP over SSL)
- SMB over IP (445/tcp, 445/udp)
- DNS (53/tcp, 53/udp used for name lookups)

To secure only the communication used between domain controllers in domain trusts, add the following rules to your IPSec policy (SEA-PN-DC-01 and SEA-NA-DC-01 are used as an example):

Rules for outbound connections initiated from SEA-PN-DC-01 to SEA-NA-DC-01

Source Address	Destination Address	Protocol	Source Port	Destination Port	Action	Certification Authority (CA)	Name of Rule/ Notes
172.16.40.5	172.16.8.5	UDP	Any	88	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, Kerberos UDP, outbound
172.16.40.5	172.16.8.5	TCP	Any	88	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, Kerberos TCP, outbound

172.16.40.5	172.16.8.5	UDP	Any	389	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, LDAP UDP, outbound
172.16.40.5	172.16.8.5	TCP	Any	389	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, LDAP TCP, outbound
172.16.40.5	172.16.8.5	TCP	Any	636	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, LDAP over SSL (LDAPS) TCP outbound
172.16.40.5	172.16.8.5	UDP	Any	445	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, SMB UDP, outbound
172.16.40.5	172.16.8.5	TCP	Any	445	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, SMB TCP, outbound
172.16.40.5	172.16.8.5	UDP	Any	53	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, DNS UDP, outbound
172.16.40.5	172.16.8.5	TCP	Any	53	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, DNS TCP, outbound

Rules for inbound connections initiated from SEA-NA-DC-01 to SEA-PN-DC-01

172.16.8.5	172.16.40.5	UDP	Any	88	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, Kerberos UDP, inbound
172.16.8.5	172.16.40.5	TCP	Any	88	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, Kerberos TCP, inbound
172.16.8.5	172.16.40.5	UDP	Any	389	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, LDAP UDP, inbound
172.16.8.5	172.16.40.5	TCP	Any	389	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, LDAP TCP, inbound

172.16.8.5	172.16.40.5	TCP	Any	636	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, LDAPS TCP, inbound
172.16.8.5	172.16.40.5	UDP	Any	445	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, SMB UDP, inbound
172.16.8.5	172.16.40.5	TCP	Any	445	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, SMB TCP, inbound
172.16.8.5	172.16.40.5	UDP	Any	53	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, DNS UDP, inbound
172.16.8.5	172.16.40.5	TCP	Any	53	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, DNS TCP, inbound
172.16.8.5	172.16.40.5	ICMP	N/A	N/A	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirrored, ICMP, both directions
172.16.8.5	172.16.40.5	Any	N/A	N/A	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	N/A	Mirrored, block all other between IPs

The second to last rule in the table, “Mirrored, ICMP, both directions,” allows ICMP traffic between domain controllers so that you can use the **ping** command to test network connectivity. However, you cannot use the **ping** command to test IPSec policy configurations because IPSec filters are defined for each protocol and port. The last rule in the table, “Mirrored, block all else between IPs,” specifies that all other traffic between the two IP addresses of the two domain controllers is blocked and that, therefore, any attempts to communicate over RPC fail. With this IPSec policy, security associations use ESP 3DES/SHA1 IPSec encapsulation for only the protocols and ports that are specified in the table. When you use Network Monitor to view the traffic that is sent over these protocols and ports, that traffic appears as ESP IPSec packets.

This IPSec policy requires that you create many rules because you must specify the direction of the traffic. It is recommended that you test an IPSec policy that uses a more basic “All Traffic” filter configuration first, before you test a policy that uses this more detailed filtering configuration (in general, you should use the simplest IPSec policy possible to meet your security requirements and management capabilities). It is also recommended that you test an IPSec policy that uses IPSec AH encapsulation first, rather than ESP, so that you can use Network Monitor captures to interpret the IPSec traffic for

troubleshooting. To further aid in troubleshooting, you can also modify the last rule in the table to use an AH action, rather than a Block action.

Keep in mind that if you use the **My IP Address** filter, you must customize an IPSec policy for each domain, but if you use specific IP addresses as source and destination addresses, you can apply this IPSec policy to both SEA-NA-DC-01 and SEA-PN-DC-01.

Using IPSec Filters to Block a Subset of IPSec-Secured Traffic

Many TCP and UDP ports on a domain controller are used to support domain members and to host other services and applications. However, some ports are not required for communication between domain controllers. A successful network attack on any port or protocol can compromise a domain controller. You can configure firewalls to permit traffic only on certain ports that are required for communication for domain trusts between a domain controller in a perimeter network and a domain controller in an internal network, as described in Deploying Domain Controllers in a Perimeter Network earlier in this paper. However, if you use IPSec to bypass these firewalls, then the firewalls might be unable to filter IPSec-secured traffic over certain ports and protocols. If you are using IPSec AH or ESP without encryption for digital signing of traffic only, verify whether the firewalls or routers that you plan to use can inspect IPSec-secured traffic. If you use IPSec ESP encryption to secure traffic, however, then network devices cannot inspect IPSec-secured traffic.

As an alternative, you can configure IPSec filters to block a subset of IPSec-secured traffic by blocking access to specific ports. Consider using this approach if you are using IPSec encryption or if your network devices cannot inspect traffic that is digitally signed by IPSec. As an example, this section describes how to configure IPSec filters to block access to NetBIOS over TCP/IP (NetBT) ports on a specific domain controller IP address, when IPSec is used to secure all other traffic with that IP address. Before implementing an IPSec policy that uses detailed filtering such as this, consider whether it is feasible to devote the time and administrative effort required to do so. If not, you might depend instead on trust configuration and measures to enhance security on domain controllers to defend against attacks on IPSec-secured traffic.

The recommended simple rule described earlier in this appendix secures all traffic between the IP addresses of the two domain controllers. The following table summarizes this rule, using the two domain controllers in the example:

Source Address	Destination Address	Protocol	Source Port	Destination Port	Action	Certification Authority (CA)	Name of Rule/ Notes
172.16.40.5	172.16.8.5	Any	N/A	N/A	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	SEA-PN-DC-01<->SEA-NA-DC-01, all traffic
172.16.8.5	172.16.40.5	Any	N/A	N/A	Require ESP 3DES/SHA1, no inbound clear, no fallback to clear	CA Root	Mirror (automatically generated)

In Windows 2000, when this filter (known as the “All Traffic” filter) is used, IPSec-secured traffic is received, verified, and then accepted without further inbound filter comparisons. On the domain controller in the internal network (SEA-NA-DC-01, in the example), IPSec cannot secure all traffic between 172.16.40.5 and 172.16.8.5 and also block inbound NetBT traffic from 172.16.40.5. However, outbound traffic is matched against the most specific outbound IPSec filter. You can use this approach as a limited defense against certain port-based attacks within an “All Traffic” filter.

For example, if an attacker compromises the domain controller in the perimeter network (SEA-PN-DC-01, in the example), the IPSec filters on this domain controller cannot provide a defense against further attacks. However, you can block access to the NetBT ports on the internal domain controller by creating a filter to block outbound NetBT responses from that domain controller to the perimeter network domain controller.

To do this, add the following additional rule to the IPSec policy for SEA-NA-DC-01. This rule secures all traffic between 172.16.40.5 and 172.16.8.5 and blocks outbound NetBT traffic from 172.16.8.5, on UDP and TCP ports 137, 138, and 139.

Source Address	Destination Address	Protocol	Source Port	Destination Port	Action	Certification Authority (CA)	Name of Rule/ Notes
172.16.8.5	172.16.40.5	UDP	137, 138, 139	Any	Block	N/A	One-way filter to block outbound NetBT UDP response traffic to SEA-PN-DC-01
172.16.8.5	172.16.40.5	TCP	137, 138, 139	Any	Block	N/A	One-way filter to block outbound NetBT TCP response traffic to SEA-PN-DC-01

Because the filters for this rule block only outbound responses from the NetBT ports on SEA-NA-DC-01, SEA-NA-DC-01 can still receive incoming traffic on these ports from SEA-PN-DC-01. Under normal operations, however, IPSec filters on SEA-PN-DC-01 can block outbound traffic on these ports.

These filters are more specific than the “All Traffic” filter because they specify certain protocols and ports. If the blocking filter specified a more general address combination, such as from 172.16.8.5 to 172.16.40.0/255.255.255.0 (the perimeter network subnet), then the blocking filter would be more general than the “All Traffic” filter, and, therefore, this filter would block IPSec-secured traffic over UDP and TCP ports 137, 138, and 139.

In Windows Server 2003 IPSec, after inbound IPSec-secured traffic is decapsulated, an additional filter lookup is performed against the entire IPSec policy. This enhancement allows further inbound filter comparisons to be made for IPSec-secured traffic, so that you can configure an inbound filter to block a specific subset of IPSec-secured traffic. For example, if SEA-NA-DC-01 was running Windows Server 2003 rather than Windows 2000, you could configure a set of inbound IPSec filters to block inbound traffic on UDP and TCP ports 137, 138, and 139. However, these inbound blocking filters would have no effect if they were applied with the “All Traffic” filter to a domain controller running Windows 2000.

Note

The use of inbound filters to block IPSec-secured traffic on NetBT ports is provided only as an example. Some domain controller deployment scenarios might require communication using NetBT or other ports. Likewise, in some domain controller environments, additional ports might need to be used. To ensure that IPSec filters do not block traffic that is required in your environment, make sure to test your IPSec policy configuration thoroughly. For example, trust validation between two domain controllers in different domains requires the use of the Net Logon service. Because the Net Logon service cannot use a single RPC port, it uses the RPC endpoint mapper and, therefore, requires that traffic be allowed over TCP port 135, UDP port 135, and a range of dynamic RPC ports that are used by the mapper. Accordingly, if you configure a one-way trust across forests so that the domain in the perimeter network trusts the domain in the internal network, and you also configure an IPSec policy on the domain controllers to block IPSec-secured traffic over TCP or UDP port 135, or any other protocols and ports used by the Net Logon service (or dependent traffic), the Net Logon service will not validate the trust.

Resources

This section provides a summary of the links that are referenced in this appendix and additional links to relevant resources for Windows 2000 IPSec.

Windows 2000 IPSec

- For information about Windows 2000 IPSec, see Internet Protocol Security, in the Windows 2000 Resource Kit, at http://go.microsoft.com/fwlink/?LinkId=16465.
- For information about Windows 2000 IPSec and other security features of Windows 2000, see IPSec, at http://go.microsoft.com/fwlink/?LinkId=13283.
- For information about diagnostics and troubleshooting procedures for Windows 2000 IPSec, see Step-by-Step Guide to Internet Protocol Security (IPSec), at http://go.microsoft.com/fwlink/?LinkId=269.
- For procedures published by the National Security Agency (NSA) for configuring IPSec policy to secure communication among all domain controllers in a Windows 2000 domain, see “Chapter 5: Configuring IPsec Policy for Secure Domain Controller Communications,” in the Microsoft Windows 2000 IPsec Guide, at http://go.microsoft.com/fwlink/?LinkId=16470.

Windows Server 2003 IPSec

- For more information about Windows Server 2003 IPSec, see Help and Support Center for Windows Server 2003.
- For deployment information about Windows Server 2003 IPSec, see Chapter 6, “Deploying IPSec” in *Deploying Network Services* in the Windows Server 2003 Deployment Kit, on the Web at http://go.microsoft.com/fwlink/?LinkId=8195.

Windows 2000 General

- For example configurations that show the deployment of Windows 2000 technologies on an actual network that simulates a large organization and the Internet, see Deployment Lab Scenarios, at http://go.microsoft.com/fwlink/?LinkId=504.
- For information about TCP/IP, see Microsoft Windows 2000 TCP/IP Implementation Details, at http://go.microsoft.com/fwlink/?LinkId=16467.

Security for Windows 2000 Active Directory

- For information about how to configure IPSec blocking filters to harden domain controllers, see Chapter 7, "Hardening Specific Server Roles," in Securing Windows 2000 Server, at http://go.microsoft.com/fwlink/?LinkId=15394.
- For information about how to configure a firewall to enable replication traffic in an Active Directory environment, see Active Directory Replication over Firewalls, at http://go.microsoft.com/fwlink/?LinkId=16461.
- For information about security considerations for using Active Directory in perimeter networks, see:
  - Chapter 5, "Security Design," in the Reference Architecture Guide: Internet Data Center, at http://go.microsoft.com/fwlink/?LinkId=16468.
  - Chapter 8, "Directory Services," in the Reference Architecture Guide: Enterprise Data Center, at http://go.microsoft.com/fwlink/?LinkId=16463.

Windows 2000 Certificate Services

- For information about Windows 2000 Certificate Services, see the "Step-by-Step Guide to Administering Certificate Services, at http://go.microsoft.com/fwlink/?LinkId=326.

Microsoft Knowledge Base Articles

- To find Knowledge Base articles, see the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462. The following Knowledge Base articles are referenced and/or relevant to the subject of this appendix:
  - Knowledge Base article 254728, "IPSec Does Not Secure Kerberos Traffic Between Domain Controllers"
  - Knowledge Base article 254949, "Client-to-Domain Controller and Domain Controller-to-Domain Controller IPSec Support"
  - Knowledge Base article 321169, "Slow SMB Performance When You Copy Files from Windows XP to a Windows 2000 Domain Controller"
  - Knowledge Base article 811832, "IPSec Default Exemptions Can Be Used to Bypass IPsec Protection in Some Scenarios"
  - Knowledge Base article 813878, "How to Block Specific Network Protocols and Ports by Using IPSec"

Microsoft Downloads

- To download the Windows 2000 High Encryption Pack, see Windows 2000 High Encryption Pack, at http://go.microsoft.com/fwlink/?LinkId=7272.

  If you are running Windows 2000 or Windows 2000 Service Pack 1, you must install the Windows 2000 High Encryption pack to use 3DES. If you are using Windows Service Pack 2 or later, you do not need to install the Windows 2000 High Encryption Pack to use 3DES.

- To download the latest version of Ipsecpol.exe, see Ipsecpol.exe: Internet Protocol Security Policies Tool, at http://go.microsoft.com/fwlink/?LinkId=16466.

IPSec Hardware Offload Adapters and Hardware Compatibility

- For information about the benefits of using IPSec hardware offload adapters, see Intel PRO/100S Network Adapter, IPSec Offload Performance and Comparison, at http://go.microsoft.com/fwlink/?LinkId=16469.
- For information about hardware that is certified for Windows 2000, see Search for Compatible Hardware Devices, at http://go.microsoft.com/fwlink/?LinkId=3787.

Appendix B: Ipsecpol Sample Script

```
@Echo off
echo AD Deployment in Segmented Networks Deployment Whitepaper
echo Windows 2000 IPsec Policy Command Line Example Script
rem #########################
rem AD Deployment in Segmented Networks deployment whitepaper: Windows 2000 IPSec Policy Command Line Example
rem 9apr03
rem
rem Download latest version of ipsecpol_setup.exe from Windows 2000 Resource Kit, Free Tools http://www.microsoft.com/windows/reskits/default.asp
rem See also: KB Q813878 for additional examples of how to use permit/block filters
rem
rem
set SEA-NA-DC-01 = 172.16.8.5
set SEA-PN-DC-01 = 172.16.40.5
set ROOTCA=OU=Copyright (c) 1997 Microsoft Corp, OU=Microsoft Corporation, CN=Microsoft Root Authority
echo SEA-NA-DC-01 is %SEA-NA-DC-01% (Corp)
echo SEA-PN-DC-01 is %SEA-PN-DC-01% (Perimeter)
echo ROOT CA is "%ROOTCA%"
rem
rem could set environment variable with policy name
rem could set environment variable with location
rem make sure ipsecpol & .dlls are in PATH
rem
echo.
echo.


echo Deleting prior versions of IPsec policy
ipsecpol -w REG -p "SEA-PN Forest Traffic to SEA-NA Forest, 20feb03, <userid>" -o
ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -o


rem #########################
rem Perimeter Network DC's IPsec policy
rem created as an empty policy that specifies appropriate General tab security settings
rem  Keeps rule settings focused on driver permit/block filters & IKE quick mode settings
rem
rem  -1s 3DES-SHA-2 = IKE Main Mode setting to use only 3DES, SHA1 and DH Group 2 (1024bits)
rem  IKE MM lifetime must be set if -1s MM security setting is used
rem  -1k 3600S = IKE MM lifetime set as 1hr/3600 seconds, unlimited quick modes under 1 main mode, no MM PFS
```

```
rem  :5 = Polling Interval 5 min - this can be set to whatever time you think best
rem     Polling has no effect on local store IPsec policy after assigned when IPSec snapin changes are made
rem     However, ipsecpol scripted changes to an already assigned local policy do not take effect until polled (without re-assignment)
rem    You want to avoid unassigning IPSec policy and reassigning it because it will interrupt IPsec secured communications
rem   possibly breaking TCP connections if too much traffic is lost
rem ##########################
echo.
echo.
echo Create: SEA-PN Perimeter Forest Traffic to SEA-NA Corp Forest, 20feb03, userid
ipsecpol -w REG -p "SEA-PN Forest Traffic to SEA-NA Forest, 20feb03, <userid>":5 -1s 3DES-SHA-2 -1k 3600S


rem ##########################
rem add rule to secure all traffic between DC IP addresses using single IP pair all traffic filter
rem   -soft = allows soft security associations for initial rollout
rem   you MUST uncheck “Allow unsecured communication with non-IPsec aware computer” once all DCs have IPsec policy
rem   IKE Quick Mode uses only ESP 3DES, SHA1
rem   no IKE quick mode PFS
rem   ipsecpol cannot specify individual names for filter lists, filters and filter actions.
rem   The name of the rule is used as the name of the filter list and the filter action.
rem
rem ##########################
echo   Add rule: SEA-PN-DC-01 to SEA-NA-DC-01 all traffic, ESP 3DES/SHA1, Cert Auth
ipsecpol -w REG -p "SEA-PN Forest Traffic to SEA-NA Forest, 20feb03, <userid>" -r "SEA-PN-DC-01<->SEA-NA-DC-01 all traffic, ESP 3DES/SHA1" -f %SEA-PN-DC-01%:*+%SEA-NA-DC-01%:*:* -n ESP[3DES,SHA] -soft -a CERT:"%ROOTCA%"


rem identical ways to specify the same filter are:
rem   172.16.40.5+172.16.8.5
rem   172.16.40.5:0+172.16.8.5:0:0


rem ##########################
rem Assign the SEA-PN-DC-01 policy active
rem ##########################
rem echo Assign: SEA-PN Perimeter Forest Traffic to SEA-NA Corp Forest, 20feb03, userid
rem ipsecpol -w REG -p "SEA-PN Forest Traffic to SEA-NA Forest, 20feb03, <userid>" -x


rem ###############################################################################


rem ##########################
rem SEA-NA-DC-01's IPsec policy created as local policy
```

```
rem  General Tab settings - some IKE Main Mode options are specified here

rem  -1s 3DES-SHA-2 = IKE Main Mode setting to use only 3DES, SHA1 and DH Group 2
(1024bits) to secure the IKE negotiation itself

rem  IKE MM lifetime must be set if -1s MM security setting is used

rem  -1k 3600S = IKE MM lifetime set as 1hr/3600 seconds

rem  use defaults for unlimited quick modes under 1 main mode, no MM PFS

rem  :5 = Polling Interval 5 min - this can be set to whatever time you think best

rem     Polling has no effect on local store IPsec policy after assigned when IPSec
snapin changes are made

rem     However, ipsecpol scripted changes to an already assigned local policy do not
take effect until polled (without re-assignment)

rem    You want to avoid unassigning IPSec policy because it will interrupt IPsec secured
communications

rem

rem ##########################

echo.

echo.

echo Create: SEA-NA Corp Forest Traffic to SEA-PN Perimeter Forest, 20feb03, userid

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>":5 -1s
3DES-SHA-2 -1k 3600S


rem ##########################

rem add rule to secure all traffic between DC IP addresses using single IP pair all
traffic filter

rem  -f %SEA-PN-DC-01%:*+%SEA-NA-DC-01%:*:* = mirrored filter for all traffic between
172.16.40.5 and 172.16.8.5

rem  -n ESP[3DES,SHA] = filter action to negotiate security, IKE QM negotiates one
security method:

rem     ESP transport mode with 3DES & SHA1

rem  -soft = filter action to allow fall back to clear (create a soft SA) if peer does
not reply.

rem     this actually applies to IKE Main Mode, but it is a filter action setting

rem  -a CERT:"%ROOTCA%"=

rem     use this root CA both to select the computer cert to offer and to validate the
peer's cert

rem     Authentication method is part of the IKE Main Mode, but specified here as part of
a rule,

rem     since rules are used to specify filters

rem  use default setting for no inbound unsecured traffic

rem  use default setting for IPsec SA & IKE QM lifetimes: 100Mbytes and 1 hour, no QM PFS

rem

rem ##########################


echo   Add rule: SEA-NA-DC-01 to SEA-PN-DC-01 all traffic, ESP 3DES/SHA1, Cert Auth

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -r "SEA-
NA-DC-01<->SEA-PN-DC-01 all traffic, ESP 3DES/SHA1" -f %SEA-NA-DC-01%:*+%SEA-PN-DC-
01%:*:* -n ESP[3DES,SHA] -soft -a CERT:"%ROOTCA%"
```

```
rem ##########################

rem more specific filters to block responses from SEA-NA-DC-01 back to SEA-PN-DC-01 IP
inside IPsec SA

rem

rem this filter must use explicit IPs so it is more specific than the “all traffic”
filter above that secures

rem all other traffic between the two IPs with IPsec ESP transport mode

rem

rem outbound responses are blocked because more specific Windows 2000 inbound filters
will not be matched

rem if traffic is received already protected by IPsec.  Windows Server 2003 does support
more specific inbound filters.

rem

rem ##########################

echo   Add rule: Block NetBT UDP 137 responses from SEA-NA-DC-01 IP

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -r "Block
NetBT UDP 137 reply from SEA-NA-DC-01 IP" -n BLOCK -f %SEA-NA-DC-01%:137=%SEA-PN-DC-
01%:0:UDP


echo   Add rule: Block NetBT UDP 138 responses from SEA-NA-DC-01 IP

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -r "Block
NetBT UDP 138 reply from SEA-NA-DC-01 IP" -n BLOCK -f %SEA-NA-DC-01%:138=%SEA-PN-DC-
01%:0:UDP


echo   Add rule: Block NetBT UDP 139 responses from SEA-NA-DC-01 IP

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -r "Block
NetBT UDP 139 reply from SEA-NA-DC-01 IP" -n BLOCK -f %SEA-NA-DC-01%:139=%SEA-PN-DC-
01%:0:UDP


echo   Add rule: Block NetBT TCP 137 responses from SEA-NA-DC-01 IP

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -r "Block
NetBT TCP 137 reply from SEA-NA-DC-01 IP" -n BLOCK -f %SEA-NA-DC-01%:137=%SEA-PN-DC-
01%:0:TCP


echo   Add rule: Block NetBT TCP 138 responses from SEA-NA-DC-01 IP

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -r "Block
NetBT TCP 138 reply from SEA-NA-DC-01 IP" -n BLOCK -f %SEA-NA-DC-01%:138=%SEA-PN-DC-
01%:0:TCP


echo   Add rule: Block NetBT TCP 139 responses from SEA-NA-DC-01 IP

ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -r "Block
NetBT TCP 139 reply from SEA-NA-DC-01 IP" -n BLOCK -f %SEA-NA-DC-01%:139=%SEA-PN-DC-
01%:0:TCP
```

```
rem ########################
rem Assign the SEA-PN-DC-01 policy active
rem ########################
rem echo Assign: SEA-NA Corp Forest Traffic to SEA-PN Perimeter Forest, 20feb03, userid
rem ipsecpol -w REG -p "SEA-NA Forest Traffic to SEA-PN Forest, 20feb03, <userid>" -x
echo.
echo.
echo Script done.
```

Appendix C: Port Punching

The following tables list the ports used for Active Directory replication, mutual authentication and the domain controller location mechanism.

Service	Port/protocol
RPC endpoint mapper	135/tcp, 135/udp
RPC static port for Active Directory replication	See Appendix D
Kerberos	88/tcp, 88/udp
LDAP	389/tcp
LDAP over SSL	636/tcp
Global Catalog LDAP	3268/tcp
Global Catalog LDAP over SSL	3269/tcp
SMB over IP (Microsoft-DS)	445/tcp, 445/udp
DNS	53/tcp, 53/udp
Network Time Protocol (NTP)	123/udp
Other non-AD network ports used:	
NetBIOS name service	137/tcp, 137/udp
NetBIOS datagram service	138/udp
NetBIOS session service	139/tcp

The following table describes other non-AD network ports that are used.

Service	Port/protocol
NetBIOS name service	137/tcp, 137/udp
NetBIOS datagram service	138/udp
NetBIOS session service	139/tcp

Appendix D: Using a Static Port for Active Directory Replication

For each service that needs to communicate across a firewall there is a fixed port and protocol. Normally, the directory service and FRS use dynamically allocated ports that require a firewall to have a wide range of ports open. Although FRS cannot be restricted to a fixed port, the directory service can be restricted to communicate on a static port which can be set using the following registry entry:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters]

"TCP/IP Port"=dword:0000c000
```

Changing this registry key on a domain controller and rebooting it causes the directory service to use the TCP port named in the registry entry. In the case above, it is port 49152(hexadecimal 0000c000).

Appendix E: Limiting the Range of Dynamic RPC Ports

You can use the registry key to limit the range of the dynamic RPC ports assigned by a particular computer. This procedure can be used to limit services that normally do not have a fixed RPC port by allowing only their dynamic port to be assigned from a smaller well-known range.

It is recommended that the dynamic ports range start at or above 5000 and consist of at least 20 ports. If additional applications that use dynamic RPC are installed on a computer, increase this range. Rebooting is necessary for the registry change to take effect.

To limit the range of dynamic RPC ports, set the following registry key:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Rpc\Internet]

"Ports"=REG_MULTI_SZ:5000-5020
```

The above registry key example shows adding a “Ports” value of type Multi-String and setting it to “5000-5020”. The value is added under the Internet key, which is added to the default RPC key. For more information about this procedure, see article 154596, "Configure RPC Dynamic Port Allocation to Work with Firewall,” in the Microsoft Knowledge Base, at http://go.microsoft.com/fwlink/?LinkId=16462.