



CHAPTER 9

VPNs

This chapter contains the following topics:

- [Updating VPNs That Include Routing Processes, page 9-1](#)
- [Loss of Communication with Spoke, page 9-2](#)
- [Configuring PKI with AAA on IOS Devices, page 9-2](#)
- [Defining Multiple CA Servers for Site-to-Site VPNs, page 9-2](#)
- [Unneeded Policy in Easy VPN Topology, page 9-4](#)
- [Discovering a VPN Already Configured in Security Manager, page 9-4](#)
- [Enabling and Disabling VRF on Catalyst Switches and 7600 Devices, page 9-4](#)
- [Commands That Cannot be Configured When Easy VPN is Enabled, page 9-5](#)
- [Defining VPNs with Multiple Spoke Definitions, page 9-6](#)
- [SSL VPN Limitations, page 9-7](#)
- [SSL VPN Limitations Due to Device OS Defects, page 9-8](#)
- [Removing Group Policy Attributes from SSL VPNs Defined on ASA 7.x Devices, page 9-9](#)



Note

For information about VPN features that are not discovered by Security Manager, see [Undiscovered VPN Features, page 6-6](#).



Note

For more detailed information on employing VPNs with Security Manager, see the [User Guide for Cisco Security Manager](#) for your release.

Updating VPNs That Include Routing Processes

Problem When you define and deploy changes to a routing process that is being used by a VPN topology (using either the Site-to-Site VPN Manager or the routing policies), the changes that you make are not reflected in the CLI commands configured on the device.

Solution When you discover a VPN topology that includes routing processes, such as GRE full mesh, Security Manager populates the GRE Modes policy in the Site-to-Site VPN Manager, as well as the relevant routing policies. However, changes made to one of these policies in Security Manager are not automatically reflected in the other policy, which can lead to unexpected results after deployment.

Therefore, if you make changes to the secured IGP in the Site-to-Site VPN Manager, be sure to go to Platform > Routing in Device view to make the necessary changes in the device's routing policies. Likewise, if you make changes directly to the routing policy, be sure to make the necessary changes in the Site-to-Site VPN Manager as well.

Loss of Communication with Spoke

Problem You lose communication with a spoke in the VPN.

Solution This problem can occur when the Security Manager server communicates with an external interface on the spoke from within the hub's protected network. We recommend that when you add the hub device to Security Manager that you define a management IP address that is located outside of the hub's protected network.

Configuring PKI with AAA on IOS Devices

Problem You cannot configure PKI with AAA authorization that uses the entire subject name on an IOS router.

Solution You can create this configuration using the predefined FlexConfig object named IOS_PKI_WITH_AAA. Please note that this FlexConfig is not supported on PIX/ASA devices.

Defining Multiple CA Servers for Site-to-Site VPNs

Problem You can select only one CA server when defining a Public Key Infrastructure (PKI) policy on a site-to-site VPN. This creates a problem when the devices in the VPN enroll with different CA servers. For example, the spoke devices might enroll with a different CA server than the hub, or the spokes in one part of the VPN might enroll with a different CA server than the spokes in another part of the VPN.

Solution To define a PKI policy, you select a PKI enrollment object that specifies the CA server to which the devices should enroll. Although by default the policy object refers globally to a single CA server, you can use device-level overrides to have the object refer to a different CA server on selected devices.

For example, if PKI enrollment object PKI_1 refers to a CA server named CA_1, you can create a device-level override for selected devices that has PKI_1 refer to a different CA server, for example, CA_2. Theoretically, you can use overrides to define a different CA server for each device in the VPN.

This procedure describes the basic steps for creating overrides for PKI enrollment objects.

**Note**

All topics that are referenced in the procedure can be found in the *User Guide for Security Manager*.

Procedure

-
- Step 1** To create the PKI enrollment object, open the PKI Enrollment dialog box. You can access this dialog box in two ways:
- Via the Public Key Infrastructure policy—Click the **Add** button beneath the Selected field. See *Configuring Public Key Infrastructure Policies*.
 - Via the Policy Object Manager—Select **PKI Enrollments** from the Object Type selector, then click the **New Object** button. See *Understanding the Policy Object Manager Window*.
- Step 2** Define the global definition of the PKI enrollment object, including the CA server to which the object refers. Be sure to select the **Allow Value Override per Device** check box. This option makes the object overridable on individual devices. For more information, see *Creating PKI Enrollment Objects*.



Note We recommend that you base the global definition of the object on the CA server that is used by the most devices in the VPN. Doing this reduces the number of device-level overrides that are required.

- Step 3** When you finish defining the PKI enrollment object, click **OK**. As a result:
- If you accessed the dialog box via the PKI policy, the new object appears in the Selected field of the policy page.
 - If you accessed the dialog box via the Policy Object Manager, the new object appears in the work area of the Policy Object Manager window. A green check mark in the Overridable column indicates that device-level overrides *may* be created for this object. (The check mark does *not* indicate whether any overrides actually exist.)

- Step 4** Create the device-level overrides for the PKI enrollment object. You can do this in one of two ways:
- Via Device Properties—This option is recommended when you want to create a device-level override for a single device. Select **Policy Object Overrides > PKI Enrollments**, select the PKI enrollment object that you want to override, then click the **Create Override** button. You can then define the content of the override, including the CA server defined by the object.
For more information, see *Creating Object Overrides for a Single Device*.
 - Via the Policy Object Manager—This option is recommended when you want to create a device-level override for multiple devices at the same time. Double-click the green check mark in the Overridable column, select the devices to which the override should apply, then define the content of the override, including the CA server defined by the object.
For more information, see *Creating Object Overrides for Multiple Devices*.
-



Note You can also use device-level overrides when the CA servers are arranged in a PKI hierarchy beneath a common, trusted CA server. To do this, you must ensure that both the global definition of the object and the device-level override specify the trusted CA server in the Trusted CA Hierarchy tab of the PKI Enrollment dialog box. See *Defining the Trusted CA Hierarchy*.

Unneeded Policy in Easy VPN Topology

Problem According to the Site-to-Site VPN Manager, your Easy VPN topology contains a policy that is not relevant to the types of devices contained in the topology.

Solution When you configure an Easy VPN topology, IOS routers, Catalyst 6500/7600 devices, and PIX 6.3 devices require you to define a user group policy. PIX 7.0 and ASA devices, however, require a tunnel group policy instead. To streamline the process, the Create VPN wizard automatically configures both policies with default values, including matching keys and group names.

If your topology contains both devices that require the user group policy and devices that require the tunnel group policy, each policy receives the relevant policy during deployment. If your topology consists entirely of devices that require the same policy (either the user group policy or the tunnel group policy), the unneeded policy is simply ignored during deployment.



Note

If you make any changes to the user group or tunnel group policies, you must make sure that the group name and the key match in both policies. Otherwise, deployment will fail.

Discovering a VPN Already Configured in Security Manager

Problem After you perform discovery, you see duplicate VPN topologies configured in the Site-to-Site VPN Manager. This situation can occur if you discover a VPN that you have already configured manually in Security Manager. If the VPN topology you discover matches the one you configured, the discovered VPN is imported into Security Manager without overwriting the VPN that you configured manually.

Solution When you add existing site-to-site VPNs to Security Manager, you should either:

- Use discovery to import the VPN into Security Manager *instead* of configuring the topology manually.
- Perform rediscov^{er} *after* configuring the VPN manually. Performing rediscov^{er} after configuring the VPN does not result in duplicate topologies. To perform rediscov^{er}, right-click the VPN in the Site-to-Site VPN Manager, then select **Re-Discover Site-To-Site VPN**.



Note

Rediscov^{er} discovers the VPN endpoints only; it does not discover the policies configured for the VPN.

Enabling and Disabling VRF on Catalyst Switches and 7600 Devices

Problem Deployment fails when you change the virtual routing and forwarding (VRF) mode on the Catalyst switches and 7600 hub of an existing site-to-site VPN. For example, if you initially configured VRF in the Create VPN wizard and deployed, but later return to the Peers policy and deselect the Enable VRF Settings check box, deployment fails. (This setting is found in the VRF Aware IPsec tab of the Edit Endpoints dialog box.) Deployment likewise fails if you try to enable VRF on a VPN that was not initially configured with it.

Solution You cannot change the VRF mode during operation. Therefore, you must do the following:

Procedure

-
- Step 1** Delete the VPN topology from Security Manager.
- Step 2** Deploy your changes.
- Step 3** Reload (restart) the Catalyst 6500/7600 device.
- Step 4** Right-click the device and select **Discover Policies on Device**.
- Step 5** Open the Create VPN wizard and redefine the VPN topology. At this point, you can select a different VRF mode.



Note

- This restriction applies only to Catalyst 6500/7600 hubs, not other device types.
 - This restriction does not apply to changes made to the VRF settings themselves. For example, if VRF is configured on the VPN topology, you can return to the Peers policy and change the VRF name or route distinguisher.
-

Commands That Cannot be Configured When Easy VPN is Enabled

Problem You cannot modify the configuration of a VPN client, including interface settings, on an ASA device when Easy VPN is enabled.

Solution The following commands (including their 'no' form) cannot be modified when Easy VPN is enabled:

- `aaa authentication listener`
- `aaa mac-exempt`
- `clear configure aaa`
- `clear configure crypto`
- `clear configure crypto isakmp`
- `clear configure crypto map`
- `clear configure nat`
- `clear configure sysopt`
- `clear configure tunnel-group`
- `crypto isakmp`
- `crypto map`
- `interface name-if`
- `interface security-level`
- `isakmp keepalive`
- `nat...access list`
- `sysopt connection permit-vpn`
- `tunnel-group`
- `webvpn enable`



Note

The `clear configure interface` command disables Easy VPN Remote.

Defining VPNs with Multiple Spoke Definitions

Problem If you discover a VPN whose spokes contain different definitions (for example, different client modes for Easy VPN spokes), Security Manager changes the definitions during discovery to create a uniform definition for all spokes. This behavior occurs because VPN topologies in Security Manager can contain only one set of spoke definitions.

Solution You can choose one of two approaches:

- Define multiple VPN topologies in Security Manager, where each topology includes spokes containing matching spoke definitions.
- Define a FlexConfig policy that contains the specialized definition, then assign the policy to the spokes that require this definition, as described in the procedure below.

Procedure

-
- Step 1** Create a shared FlexConfig policy in Policy view:
- Select **View > Policy View**.
 - Right-click **FlexConfigs** in the Policy Type selector, then select **New FlexConfigs Policy**.
 - Enter a name for the policy, then click **OK**. The new shared policy is displayed in the Shared Policy selector in the lower-left pane of Policy view.
- Step 2** Define the FlexConfig policy by creating and selecting a FlexConfig object:
- In the work area of Policy view, click the **Add** button on the Details tab.
 - In the FlexConfigs Selector, click the **Create** button in the lower-left corner of the window. The FlexConfig dialog box is displayed.
 - Define an appended FlexConfig object that contains the required client definition. For example, to define the client mode on an Easy VPN spoke, enter the following commands:


```
crypto ipsec client ezvpn CSM_EASY_VPN_CLIENT_1
mode client
exit
```
 - After you create the FlexConfig object, add it to the FlexConfig policy using the selector.
- Step 3** In the work area of Policy view, use the Assignments tab to select the spokes to which this policy should be assigned, then click **Save**.
- Step 4** Deploy the policy.
-



Note

For more information about the steps described in this procedure, see the following topics in [User Guide for Cisco Security Manager](#) for your release:

- Creating a New Shared Policy
 - Creating FlexConfig Policy Objects
 - Modifying Policy Assignments in Policy View
-

SSL VPN Limitations

The current implementation of SSL VPN in Security Manager is subject to the following limitations:

- SSL VPN license information cannot be imported into Security Manager. As a result, certain command parameters, such as **vpn sessiondb** and **max-webvpn-session-limit**, cannot be validated.
- You must configure DNS on each device in the topology in order to use clientless SSL VPN. Without DNS, the device cannot retrieve named URLs, but only URLs with IP addresses.
- If you share your Connection Profiles policy among multiple ASA devices, bear in mind that all devices share the same address pool unless you use device-level object overrides to replace the global definition with a unique address pool for each device. Unique address pools are required to avoid overlapping addresses in cases where the devices are not using NAT.
- Bear in mind that you must use interface roles, not physical interfaces, when defining SSL VPN gateways on IOS devices. On ASA devices, however, you can select physical interfaces when defining an Access policy. For more information about interface roles, see “Working with Interface Role Objects” in the *User Guide for Cisco Security Manager*.
- Security Manager (and ASA devices in general) do not check whether proxy-bypass interfaces are also configured as SSL VPN-enabled. If proxy-bypass is enabled on an interface that is not SSL VPN-enabled, certain 7.2 releases prevent you from reusing the proxy-bypass port after the rule is removed. The only solution to this problem is to reboot the device.
- If the device configuration contains an address pool for SSL VPN with a name that begins CSM_ (the naming convention used by Security Manager), Security Manager cannot detect whether the addresses in that pool overlap with the pool configured in your SSL VPN policy. (This can occur, for example, when the pool was configured by a user on a different installation of Security Manager.) This can lead to errors during deployment. Therefore, we recommend that you configure the same IP address pool as a network/host object in Security Manager and define it as part of the SSL VPN policy. This enables the proper validation to take place.
- The same IP address and port number cannot be shared by multiple SSL VPN gateways on the same IOS device. As a result, deployment errors can occur if a duplicate gateway exists in the device configuration but was not redefined using the Security Manager interface. If such an error occurs, you must choose a different IP address and port number and redeploy.
- If you define AAA authentication or accounting as part of an SSL VPN policy, the **aaa new-model** command is deployed to enable AAA services. Bear in mind that this command is not removed if you later delete the SSL VPN policy, as there might be other parts of the device configuration that require the **aaa new-model** command for AAA services.

**Note**

In addition, we recommend that you define at least one local user on the device with a privilege level of 15. This ensures that you will not be locked out of the device if the **aaa new-model** command is configured without an associated AAA server.

SSL VPN Limitations Due to Device OS Defects

The current implementation of SSL VPN in Security Manager is subject to the following limitations caused by existing IOS and ASA defects:

- Deployment fails if you remove a port forwarding list used by an SSL VPN user group. This problem occurs in IOS 12.4(9)T and was corrected in IOS 12.4(12.15)T. The workaround is to delete all the attributes of the port forwarding list (other than the name, which is mandatory) instead of removing it from the user group. For more information, see [CSCsh50799](#).



Note If the port forwarding list is used by other user groups, you can ignore the deployment error.

- Deployment fails if you modify the attributes of a WINS master server (for example, the timeout) used by an SSL VPN user group. This problem occurs in IOS 12.4(9)T and was corrected in 12.4(13.11)T. The workaround is to remove the WINS server from the user group and deploy. After deployment, you can make the necessary changes to the WINS server and add it back to the user group. For more information, see [CSCsg16935](#).
- Deployment fails if the addresses in the address pool used by an SSL VPN user group do not belong to the same subnet as one of the interfaces on the device. This problem occurs in IOS 12.4(11)T. The workaround is to create a loopback interface that is on the same subnet as the addresses in the pool.
- If you define a AAA accounting server in the SSL VPN policy, you must have a default accounting server defined on the device. Otherwise, accounting functions (such as keeping track of how many times an SSL VPN connection is used, by whom, and for how long) are not performed. This problem occurs in IOS 12.4(9)T and was corrected in 12.4(10.04)T. For more information, see [CSCse90029](#). To assign an accounting server to SSL VPN, enter the following CLI command:

```
aaa accounting network default start-stop group radius
```



Note If you use a FlexConfig to enter this command, be sure to remove the FlexConfig after deployment. Otherwise, the command will be reissued each time that you redeploy.

- When CNS is configured, the port forwarding list and the URL list defined in Security Manager are assigned to the wrong SSL VPN context. For example, if these lists are defined to context 1, they are deployed to context 2. This problem occurs in IOS 12.4(11)T. The workaround is remove the CNS configuration before defining these lists and restoring the configuration afterwards. For more information, see [CSCsh72072](#).

Removing Group Policy Attributes from SSL VPNs Defined on ASA 7.x Devices

Problem After you deploy changes to a Tunnel Group policy for a remote access VPN on a PIX/ASA 7.x device, you find that the **group-policy** commands defined on the device for SSL VPN have been removed.

Solution Security Manager does not discover SSL VPN device configurations. As a result, it does not make changes to these configurations unless and until you define and deploy SSL VPN policies using the Security Manager interface. However, **group-policy** (which is modeled in Security Manager as ASA user group objects) is an exception, because it is used by both SSL VPNs and IPsec remote access VPNs, as follows:

- SSL VPNs—User Groups policy, Connection Profiles policy
- Remote access VPNs—Tunnel Group policy (General tab)

A device configuration can use the same group-policy definition (that is, the same ASA user group) in both policies. When you discover that configuration, only the remote access VPN attributes are imported into Security Manager. As a result, on the next deployment, the remote access VPN attributes are deployed to the device and the SSL VPN attributes are removed.

Therefore, if the device configuration uses the same group-policy definition for remote access VPN as well as for SSL VPN, you must define an SSL user groups policy to compensate for the fact that it was not defined as a result of the discovery process.

