



Configure VPN Access

- [Connect and Disconnect to a VPN](#), on page 1
- [Configure Start Before Logon \(PLAP\) on Windows Systems](#), on page 7
- [Use Trusted Network Detection to Connect and Disconnect](#), on page 8
- [Require VPN Connections Using Always-On](#), on page 11
- [Use Captive Portal Hotspot Detection and Remediation](#), on page 16
- [Configure AnyConnect over L2TP or PPTP](#), on page 18
- [Configure AnyConnect Proxy Connections](#), on page 19
- [Select and Exclude VPN Traffic](#), on page 23
- [Manage VPN Authentication](#), on page 27

Connect and Disconnect to a VPN

AnyConnect VPN Connectivity Options

The AnyConnect client provides many options for automatically connecting, reconnecting, or disconnecting VPN sessions. These options provide a convenient way for your users to connect to your VPN, and they also support your network security requirements.

Starting and Restarting AnyConnect Connections

[Configure VPN Connection Servers](#) to provide the names and addresses of the secure gateways your users will manually connect to.

Choose from the following AnyConnect capabilities to provide convenient, automatic VPN connectivity:

- [Automatically Start Windows VPN Connections Before Logon](#)
- [Automatically Start VPN Connections When AnyConnect Starts](#)
- [Automatically Restart VPN Connections](#)

Also, consider using the following Automatic VPN Policy options to enforce greater network security or restrict network access to the VPN only:

- [About Trusted Network Detection](#)
- [Require VPN Connections Using Always-On](#)

- [Use Captive Portal Hotspot Detection and Remediation](#)

Renegotiating and Maintaining the AnyConnect Connection

You can limit how long the ASA keeps an AnyConnect VPN connection available to the user even with no activity. If a VPN session goes idle, you can terminate the connection or re-negotiate the connection.

- **Keepalive**—The ASA sends keepalive messages at regular intervals. These messages are ignored by the ASA, but are useful in maintaining connections with devices between the client and the ASA.

For instructions to configure Keepalive with the ASDM or CLI, see the *Enable Keepalive* section in the [Cisco ASA Series VPN Configuration Guide](#).

- **Dead Peer Detection**—The ASA and AnyConnect client send "R-U-There" messages. These messages are sent less frequently than IPsec's keepalive messages. You can enable both the ASA (gateway) and the AnyConnect client to send DPD messages, and configure a timeout interval.

- If the client does not respond to the ASA's DPD messages, the ASA tries once more before putting the session into "Waiting to Resume" mode. This mode allows the user to roam networks, or enter sleep mode and later recover the connection. If the user does not reconnect before the idle timeout occurs, the ASA will terminate the tunnel. The recommended gateway DPD interval is 300 seconds.
- If the ASA does not respond to the client's DPD messages, the client tries again before terminating the tunnel. The recommended client DPD interval is 30 seconds.

For instructions to configure DPD within the ASDM, refer to *Configure Dead Peer Detection* in the appropriate release of the [Cisco ASA Series VPN Configuration Guide](#).

- **Best Practices:**
 - Set Client DPD to 30 seconds (Group Policy > Advanced > AnyConnect Client > Dead Peer Detection).
 - Set Server DPD to 300 seconds (Group Policy > Advanced > AnyConnect Client > Dead Peer Detection).
 - Set Rekey, for both SSL and IPsec to 1 hour (Group Policy > Advanced > AnyConnect Client > Key Regeneration).

Terminating an AnyConnect Connection

Terminating an AnyConnect connection requires the user to re-authenticate their endpoint to the secure gateway and create a new VPN connection.

The following connection parameters terminate the VPN session based on timeouts:

- **Maximum Connect Time**—Sets the maximum user connection time in minutes. At the end of this time, the system terminates the connection. You can also allow unlimited connection time(default).
- **VPN Idle Timeout**—Terminates any user's session when the session is inactive for the specified time. If the VPN idle timeout is not configured, then the default idle timeout is used.
- **Default Idle Timeout**—Terminates any user's session when the session is inactive for the specified time. The default value is 30 minutes. The default is 1800 second.

See the *Specify a VPN Session Idle Timeout for a Group Policy* section in the appropriate release of the [Cisco ASA Series VPN Configuration Guide](#) to set these parameters.

Configure VPN Connection Servers

The AnyConnect VPN server list consists of host name and host address pairs identifying the secure gateways that your VPN users will connect to. The host name can be an alias, an FQDN, or an IP address.

The hosts added to the server list display in the Connect to drop-down list in the AnyConnect GUI. The user can then select from the drop-down list to initiate a VPN connection. The host at the top of the list is the default server, and appears first in the GUI drop-down list. If the user selects an alternate server from the list, the selected server becomes the new default server.

Once you add a server to the server list, you can view its details and edit or delete the server entry. To add a server to the server list, follow this procedure.

Procedure

-
- Step 1** Open the VPN Profile Editor and choose **Server List** from the navigation pane.
- Step 2** Click **Add**.
- Step 3** Configure the server's host name and address:
- Enter a **Host Display Name**, an alias used to refer to the host, an FQDN, or an IP address. Do not use "&" or "<" characters in the name. If you enter an FQDN or an IP address, you do not need to enter the **FQDN** or **IP Address** in the next step.

If you enter an IP address, use the Public IPv4 or the Global IPv6 address of the secure gateway. Use of the link-local secure gateway address is not supported.
 - (Optional) Enter the host's **FQDN** or **IP Address** if not entered in the Host Display Name.
 - (Optional) Specify a **User Group**.

AnyConnect uses the FQDN or IP Address in conjunction with User Group to form the Group URL.
- Step 4** Enter the server to fall back to as the backup server in the **Backup Server List**. Do not use "&" or "<" characters in the name.
- Note** Conversely, the Backup Server tab on the Server menu is a global entry for all connection entries. Any entries put in that Backup Server location are overwritten with what is entered here for an individual server list entry. This setting takes precedence and is the recommended practice.
- Step 5** (Optional) Add load balancing servers to the **Load Balancing Server List**. Do not use "&" or "<" characters in the name.
- If the host for this server list entry specifies a load balancing cluster of security appliances, and the Always-On feature is enabled, add the load balancing devices in the cluster to this list. If you do not, Always-On blocks access to the devices in the load balancing cluster.
- Step 6** Specify the **Primary Protocol** for the client to use for this ASA:
- Choose SSL (default) or IPsec.

If you specify IPsec, the User Group must be the exact name of the connection profile (tunnel group). For SSL, the user group is the group-url or group-alias of the connection profile.

- b) If you specify IPsec, select **Standard Authentication Only** to disable the default authentication method (proprietary AnyConnect EAP), and choose a method from the drop-down list.

Note Changing the authentication method from the proprietary AnyConnect EAP to a standards-based method disables the ability of the ASA to configure session timeout, idle timeout, disconnected timeout, split tunneling, split DNS, MSIE proxy configuration, and other features.

Step 7 (Optional) Configure SCEP for this server:

- Specify the URL of the SCEP CA server. Enter an FQDN or IP Address. For example, `http://ca01.cisco.com`.
- Check **Prompt For Challenge PW** to enable the user to make certificate requests manually. When the user clicks **Get Certificate**, the client prompts the user for a username and one-time password.
- Enter the certificate thumbprint of the CA. Use SHA1 or MD5 hashes. Your CA server administrator can provide the CA URL and thumbprint and should retrieve the thumbprint directly from the server and not from a “fingerprint” or “thumbprint” attribute field in a certificate it issued.

Step 8 Click **OK**.

Related Topics

[AnyConnect Profile Editor, Server List](#)

[AnyConnect Profile Editor, Add/Edit a Server List](#)

Automatically Start Windows VPN Connections Before Logon

About Start Before Logon

This feature called Start Before Logon (SBL) allows users to establish their VPN connection to the enterprise infrastructure before logging onto Windows.



Note When using Start Before Logon (SBL) and HostScan, you must install the AnyConnect/HostScan posture predeploy module on the endpoints to achieve full HostScan functionality, since SBL is pre-login.

After SBL is installed and enabled, the Network Connection button launches the AnyConnect VPN and Network Access Manager UI.

SBL also includes the Network Access Manager tile and allows connections using user configured home network profiles. Network profiles allowed in SBL mode include all media types employing non-802.1X authentication modes, such as open WEP, WPA/WPA2 Personal, and static key (WEP) networks.

SBL is available on Windows systems only, and is implemented using different mechanisms depending on the version of Windows:

- On Windows, the Pre-Login Access Provider (PLAP) is used to implement AnyConnect SBL.

With PLAP, the Ctrl+Alt+Del key combination opens a window where the user can choose either to log in to the system or activate Network Connections (PLAP components) using the Network Connect button in the lower-right corner of the window.

PLAP supports 32-bit and 64-bit versions of the Windows.

Reasons you might consider enabling SBL for your users include:

- The user's computer is joined to an Active Directory infrastructure.
- A user has network-mapped drives that require authentication with the Microsoft Active Directory infrastructure.
- The user cannot have cached credentials on the computer (the group policy disallows cached credentials). In this scenario, users must be able to communicate with a domain controller on the corporate network for their credentials to be validated before gaining access to the computer.
- The user must run logon scripts that execute from a network resource or need access to a network resource. With SBL enabled, the user has access to the local infrastructure and logon scripts that would normally run when a user is in the office. This includes domain logon scripts, group policy objects and other Active Directory functionality that normally occurs when users log on to their system.
- Networking components (such as MS NAP/CS NAC) exist that might require connection to the infrastructure.

Limitations on Start Before Logon

- AnyConnect is not compatible with fast user switching.
- AnyConnect cannot be started by third-party Start Before Logon applications.
- Because SBL is pre-login and will not have access to the user store, you cannot do multiple certificate authentication (MCA) with it. MCA requires a machine certificate and a user certificate, or two user certificates.

Configure Start Before Logon

Procedure

-
- Step 1** [Install the AnyConnect Start Before Logon Module.](#)
- Step 2** [Enable SBL in the AnyConnect Profile.](#)
-

Install the AnyConnect Start Before Logon Module

The AnyConnect installer detects the underlying operating system and places the appropriate AnyConnect DLL from the AnyConnect SBL module in the system directory. On Windows 7, or the Windows 2008 server, the installer determines whether the 32-bit or 64-bit version of the operating system is in use and installs the appropriate PLAP component, vpnplap.dll or vpnplap64.dll.



Note If you uninstall AnyConnect while leaving the VPNGINA or PLAP component installed, the VPNGINA or PLAP component is disabled and not visible to the remote user.

You can predeploy the SBL module or configure the ASA to download it. When predeploying AnyConnect, the Start Before Logon module requires that the core client software is installed first. If you are predeploying AnyConnect Core and the Start Before Logon components using MSI files, you must get the order right.

Procedure

- Step 1** In ASDM go to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**.
 - Step 2** Select a group policy and click **Edit** or **Add** a new group policy.
 - Step 3** Select **Advanced > AnyConnect Client** in the left navigation pane.
 - Step 4** Uncheck **Inherit** for the Optional Client Module for Download setting.
 - Step 5** Select the **AnyConnect SBL** module in the drop-down list.
-

Enable SBL in the AnyConnect Profile

Before you begin

- SBL requires a network connection to be present at the time it is invoked. In some cases, this might not be possible, because a wireless connection might depend on credentials of the user to connect to the wireless infrastructure. Since SBL mode precedes the credential phase of a logon, a connection would not be available in this scenario. In this case, the wireless connection needs to be configured to cache the credentials across logon, or another wireless authentication needs to be configured, for SBL to work.
- If the Network Access Manager is installed, you must deploy device connection to ensure that an appropriate connection is available.

Procedure

- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 1)** from the navigation pane.
- Step 2** Select **Use Start Before Logon**.
- Step 3** (Optional) To give the remote user control over SBL, select **User Controllable**.

Note The user must reboot the remote computer before SBL takes effect.

Troubleshoot Start Before Logon

Procedure

- Step 1** Ensure that the AnyConnect profile is loaded on the ASA, ready to be deployed.
- Step 2** Delete prior profiles (search for them on the hard drive to find the location, *.xml).
- Step 3** Using Windows Add/Remove Programs, uninstall the SBL Components. Reboot the computer and retest.
- Step 4** Clear the user's AnyConnect log in the Event Viewer and retest.

- Step 5** Browse back to the security appliance to install AnyConnect again.
- Step 6** Reboot once. On the next reboot, you should be prompted with the Start Before Logon prompt.
- Step 7** Collect a DART bundle and send it to your AnyConnect Administrator.
- Step 8** If you see the following error, delete the user's AnyConnect profile:

```
Description: Unable to parse the profile C:\Documents and Settings\All Users\Application Data
\Cisco\Cisco AnyConnect Secure Mobility Client\Profile\VABaseProfile.xml. Host data not
available.
```

- Step 9** Go back to the .tmpl file, save a copy as an.xml file, and use that XML file as the default profile.

Automatically Start VPN Connections When AnyConnect Starts

This feature called Auto Connect On Start, automatically establishes a VPN connection with the secure gateway specified by the VPN client profile when AnyConnect starts.

Auto Connect On Start is disabled by default, requiring the user to specify or select a secure gateway.

Procedure

-
- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 1)** from the navigation pane.
- Step 2** Select **Auto Connect On Start**.
- Step 3** (Optional) To give the user control over Auto Connect on Start, select **User Controllable**.
-

Configure Start Before Logon (PLAP) on Windows Systems

The Start Before Logon (SBL) feature starts a VPN connection before the user logs in to Windows. This ensures that users connect to their corporate infrastructure before logging on to their computers.

The SBL AnyConnect feature is known as the Pre-Login Access Provider (PLAP), which is a connectable credential provider. This feature lets programmatic network administrators perform specific tasks, such as collecting credentials or connecting to network resources before logon. PLAP provides SBL functions on all of the supported Windows operating systems. PLAP supports 32-bit and 64-bit versions of the operating system with vpnplap.dll and vpnplap64.dll, respectively. The PLAP functions supports x86 and x64.

Automatically Restart VPN Connections

When Auto Reconnect is enabled (default), AnyConnect recovers from VPN session disruptions and reestablishes a session, regardless of the media used for the initial connection. For example, it can reestablish a session on wired, wireless, or 3G. When Auto Reconnect is enabled, you also specify the reconnect behavior upon system suspend or system resume. A system suspend is a low-power standby, such as Windows "hibernation" or macOS or Linux "sleep." A system resume is a recovery following a system suspend.

If you disable Auto Reconnect, the client does not attempt to reconnect regardless of the cause of the disconnection. Cisco highly recommends using the default setting (enabled) for this feature. Disabling this setting can cause interruptions in VPN connectivity over unstable connections.

Procedure

- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 1)** from the navigation pane.
- Step 2** Select **Auto Reconnect**.
- Step 3** Choose the Auto Reconnect Behavior:
- **Disconnect On Suspend**—(Default) AnyConnect releases the resources assigned to the VPN session upon a system suspend and does not attempt to reconnect after the system resume.
 - **Reconnect After Resume**—The client retains resources assigned to the VPN session during a system suspend and attempts to reconnect after the system resume.
-

Use Trusted Network Detection to Connect and Disconnect

About Trusted Network Detection

Trusted Network Detection (TND) gives you the ability to have AnyConnect automatically disconnect a VPN connection when the user is inside the corporate network (the trusted network) and start the VPN connection when the user is outside the corporate network (the untrusted network).

TND does not interfere with the ability of the user to manually establish a VPN connection. It does not disconnect a VPN connection that the user starts manually in the trusted network. TND only disconnects the VPN session if the user first connects in an untrusted network and moves into a trusted network. For example, TND disconnects the VPN session if the user makes a VPN connection at home and then moves into the corporate office.

You configure TND in the AnyConnect VPN Client profile. No changes are required to the ASA configuration. You need to specify the action or policy AnyConnect takes when recognizing it is transitioning between trusted and untrusted networks, and identify your trusted networks and servers.

Guidelines for Trusted Network Detection

- Because the TND feature controls the AnyConnect GUI and automatically starts connections, the GUI should run at all times. If the user exits the GUI, TND does not automatically start the VPN connection.
- If AnyConnect is also running Start Before Logon (SBL), and the user moves into the trusted network, the SBL window displayed on the computer automatically closes.
- Trusted Network Detection with or without Always-On configured is supported on IPv6 and IPv4 VPN connections to the ASA over IPv4 and IPv6 networks.
- Multiple profiles on a user computer may present problems if the TND configuration is different.

If the user has received a TND-enabled profile in the past, upon system restart, AnyConnect attempts to connect to the security appliance it was last connected to, which may not be the behavior you desire. To connect to a different security appliance, they must manually disconnect and re-connect to that headend. The following workarounds will help you prevent this problem:

- Enable TND in the client profiles loaded on all the ASAs on your corporate network.
 - Create one profile listing all the ASAs in the host entry section, and load that profile on all your ASAs.
 - If users do not need to have multiple, different profiles, use the same profile name for the profiles on all the ASAs. Each ASA overrides the existing profile.
- To use TND on Linux, you must have the Network Manager installed and running properly on the target (RHEL/Ubuntu) device, and the network manager must be maintaining the network interfaces.

Configure Trusted Network Detection

Procedure

Step 1 Open the VPN profile editor and choose **Preferences (Part 2)** from the navigation pane.

Step 2 Select **Automatic VPN Policy**.

Step 3 Choose a **Trusted Network Policy**.

This is the action the client takes when the user is inside the corporate network (the trusted network). The options are:

- **Disconnect**—(Default) The client terminates the VPN connection in the trusted network.
- **Connect**—The client starts a VPN connection in the trusted network.
- **Do Nothing**—The client takes no action in the trusted network. Setting both the Trusted Network Policy and Untrusted Network Policy to Do Nothing disables Trusted Network Detection (TND).
- **Pause**—AnyConnect suspends the VPN session (instead of disconnecting it) if a user enters a network configured as trusted after establishing a VPN session outside the trusted network. When the user goes outside the trusted network again, AnyConnect resumes the session. This feature is for the user's convenience because it eliminates the need to establish a new VPN session after leaving a trusted network.

Step 4 Choose an **Untrusted Network Policy**.

This is the action the client takes when the user is outside the corporate network. The options are:

- **Connect**—The client starts a VPN connection upon the detection of an untrusted network.
- **Do Nothing**—The client takes no action upon detection of an untrusted network. This option disables Always-On VPN. Setting both the Trusted Network Policy and Untrusted Network Policy to **Do Nothing** disables Trusted Network Detection.

Step 5 Specify **Trusted DNS Domains**.

Specify the DNS suffixes (a string separated by commas) that a network interface may have when the client is in the trusted network. You can assign multiple DNS suffixes if you add them to the split-dns list and specify a default domain on the ASA.

The AnyConnect client builds the DNS suffix list in the following order:

- The domain passed by the head end.
- The split-DNS suffix list passed by the head end.
- The public interface's DNS suffixes, if configured. If not, the primary and connection-specific suffixes, along with the parent suffixes of the primary DNS suffix (if the corresponding box is checked in the Advanced TCP/IP Settings).

To Match This DNS Suffix:	Use This Value for TrustedDNSDomains:
example.com (only)	*example.com
example.com AND anyconnect.example.com	*.example.com OR example.com, anyconnect.example.com
asa.example.com AND anyconnect.example.com	*.example.com OR asa.example.com, anyconnect.example.com

Step 6 Specify Trusted DNS Servers.

All DNS server addresses (a string separated by commas) that a network interface may have when the client is in the trusted network. For example: 203.0.113.1,2001:DB8::1. Wildcards (*) are supported for IPv4 and IPv6 DNS server addresses.

You must have a DNS entry for the headend server that is resolvable via DNS. If your connections are by IP address, you need a DNS server that can resolve mus.cisco.com. If mus.cisco.com is not resolvable via DNS, captive portal detection will not work as expected.

Note You can configure either TrustedDNSDomains, TrustedDNSServers, or both. If you configure TrustedDNSServers, be sure to enter all your DNS servers, so your site(s) will all be part of the Trusted Network.

An active interface will be considered as an In-Trusted-Network if it matches *all* the rules in the VPN profile.

Step 7 Specify a host URL that you want to add as trusted. You must have a secure web server that is accessible with a trusted certificate to be considered trusted. After you click **Add**, the URL is added and the certificate hash is pre-filled. If the hash is not found, an error message prompts the user to enter the certificate hash manually and click **Set**.

Note You can configure this parameter only when at least one of the Trusted DNS Domains or Trusted DNS Servers is defined. If Trusted DNS Domains or Trusted DNS Servers are not defined, this field is disabled.

Require VPN Connections Using Always-On

About Always-On VPN

Always-On operation prevents access to Internet resources when the computer is not on a trusted network, unless a VPN session is active. Enforcing the VPN to always be on in this situation protects the computer from security threats.

When Always-On is enabled, it establishes a VPN session automatically after the user logs in and upon detection of an untrusted network. The VPN session remains open until the user logs out of the computer, or the session timer or idle session timer (specified in the ASA group policy) expires. AnyConnect continually attempts to reestablish the connection to reactivate the session if it is still open; otherwise, it continually attempts to establish a new VPN session.

When Always-On is enabled in the VPN Profile, AnyConnect protects the endpoint by deleting all the other downloaded AnyConnect profiles and ignores any public proxies configured to connect to the ASA.

The following AnyConnect options also need to be considered when enabling Always-On:

- Allowing the user to disconnect the Always-On VPN session: AnyConnect provides the ability for the user to disconnect Always-On VPN sessions. If you enable **Allow VPN Disconnect**, AnyConnect displays a Disconnect button upon the establishment of a VPN session. By default, the profile editor enables the Disconnect button when you enable Always-On VPN.

Pressing the disconnect button locks all interfaces to prevent data from leaking out and to protect the computer from internet access except for establishing a VPN session. Users of Always-On VPN sessions may want to click Disconnect so they can choose an alternative secure gateway due to performance issues with the current VPN session, or reconnection issues following the interruption of a VPN session.

- Setting a connect failure policy: The connect failure policy determines whether the computer can access the internet if Always-On VPN is enabled and AnyConnect cannot establish a VPN session. See [Set a Connect Failure Policy for Always-On](#).
- Handling captive portal hotspots: See [Use Captive Portal Hotspot Detection and Remediation](#).

Limitations of Always-On VPN

- Always On is available only on Windows and macOS
- If Always-On is enabled, but the user does not log on, AnyConnect does not establish the VPN connection. AnyConnect starts the VPN connection only post-login.
- Always-On VPN does not support connecting through a proxy.

Guidelines for Always-On VPN

To enhance protection against threats, we recommend the following additional protective measures if you configure Always-On VPN:

- We strongly recommend purchasing a digital certificate from a certificate authority (CA) and enrolling it on the secure gateways. The ASDM provides an **Enroll ASA SSL VPN with Entrust** button on the

Configuration > Remote Access VPN > Certificate Management > Identity Certificates panel to facilitate enrollment of a public certificate.

- If you are using always-on VPN, external SAML IdP is not supported (however, with internal SAML IdP, the ASA proxies all traffic to IdP and is supported)
- Predeploy a profile configured with Always-On to the endpoints to limit connectivity to the pre-defined ASAs. Predeployment prevents contact with a rogue server.
- Restrict administrator rights so that users cannot terminate processes. A PC user with admin rights can bypass an Always-On policy by stopping the agent. If you want to ensure fully-secure Always-On, you must deny local admin rights to users.
- Restrict access to the Cisco sub-folders on Windows computers, typically `C:\ProgramData`.
- Users with limited or standard privileges may sometimes have write access to their program data folders. They could use this access to delete the AnyConnect profile file and thereby circumvent the Always-On feature.
- Predeploy a group policy object (GPO) for Windows users to prevent users with limited rights from terminating the GUI. Predeploy equivalent measures for macOS users.

Configure Always-On VPN

Procedure

- Step 1** [Configure Always-On in the AnyConnect VPN Client Profile, on page 12.](#)
 - Step 2** (Optional) [Add Load-Balancing Backup Cluster Members to the Server List.](#)
 - Step 3** (Optional) [Exempt Users from Always-On VPN.](#)
-

Configure Always-On in the AnyConnect VPN Client Profile

Before you begin

Always-On VPN requires that a valid, trusted server certificate be configured on the ASA; otherwise, it fails and logs an event indicating the certificate is invalid. In addition, ensuring that the server certificate can pass Strict Certificate Trust mode prevents the download of an Always-On VPN profile that locks a VPN connection to a rogue server.

Procedure

- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 2)** from the navigation pane.
- Step 2** Select **Automatic VPN Policy**.
- Step 3** [Configure Trusted Network Detection, on page 9](#)
- Step 4** Select **Always On**.
- Step 5** (Optional) Select or un-select **Allow VPN Disconnect**.

- Step 6** (Optional) [Configure a Connect Failure Policy](#).
- Step 7** (Optional) [Configure Captive Portal Remediation](#).

Add Load-Balancing Backup Cluster Members to the Server List

Always-On VPN affects the load balancing of AnyConnect VPN sessions. With Always-On VPN disabled, when the client connects to a primary device within a load balancing cluster, the client complies with a redirection from the primary device to any of the backup cluster members. With Always-On enabled, the client does not comply with a redirection from the primary device unless the address of the backup cluster member is specified in the server list of the client profile. Therefore, be sure to add any backup cluster members to the server list.

To specify the addresses of backup cluster members in the client profile, use ASDM to add a load-balancing backup server list by following these steps:

Procedure

- Step 1** Open the VPN Profile Editor and choose **Server List** from the navigation pane.
- Step 2** Choose a server that is a primary device of a load-balancing cluster and click **Edit**.
- Step 3** Enter an FQDN or IP address of any load-balancing cluster member.
-

Exempt Users from Always-On VPN

You can configure exemptions to override an Always-On policy. For example, you might want to let certain individuals establish VPN sessions with other companies or exempt the Always-On policy for noncorporate assets.

Exemptions set in group policies and dynamic access policies on the ASA override the Always-On policy. You specify exceptions according to the matching criteria used to assign the policy. If an AnyConnect policy enables Always-On and a dynamic access policy or group policy disables it, the client retains the disable setting for the current and future VPN sessions as long as its criteria match the dynamic access policy or group policy on the establishment of each new session.

This procedure configures a dynamic access policy that uses AAA endpoint criteria to match sessions to noncorporate assets.

Procedure

- Step 1** Choose **Configuration > Remote Access VPN > Network (Client) Access > Dynamic Access Policies > Add** or **Edit**.
- Step 2** Configure criteria to exempt users from Always-On VPN. For example, use the Selection Criteria area to specify AAA attributes to match user logon IDs.
- Step 3** Click the **AnyConnect** tab on the bottom half of the Add or Edit Dynamic Access Policy window.

Add Dynamic Access Policy

Policy Name: Description: ACL Priority:

Selection Criteria

Define the AAA and endpoint attributes used to select this access policy. A policy is used when a user's authorization attributes match the AAA attribute criteria below and every endpoint attribute has been satisfied. These attributes can be created using the tables below and/or by expanding the Advanced option to specify the logical expression text.

User has ANY of the following AAA Attributes values... and the following endpoint attributes are satisfied.

AAA Attribute	Operation/Value	Endpoint ID	Name/Operation/Value
cisco.username	=		jsmith

Advanced

Access/Authorization Policy Attributes

Configure access/authorization attributes for this policy. Attribute values specified here will override those values obtained from the AAA system and the group-policy hierarchy. The resulting VPN authorization policy is an aggregation of DAP attributes, AAA attributes, and group-policy hierarchy attributes (those that are not specified in DAP).

Always-On VPN for AnyConnect client:
 Unchanged
 Use AnyConnectProfile setting
 Disable

Step 4 Click **Disable** next to “Always-On VPN for AnyConnect client.”

Set a Connect Failure Policy for Always-On

About the Connect Failure Policy

The connect failure policy determines whether the computer can access the internet if Always-On VPN is enabled and AnyConnect cannot establish a VPN session. This can occur when a secure gateway is unreachable, or when AnyConnect fails to detect the presence of a captive portal hotspot.

An open policy permits full network access, letting users continue to perform tasks where access to the Internet or other local network resources is needed.

A closed policy disables all network connectivity until the VPN session is established. AnyConnect does this by enabling packet filters that block all traffic from the endpoint that is not bound for a secure gateway to which the computer is allowed to connect.

Regardless of the connect failure policy, AnyConnect continues to try to establish the VPN connection.

Guidelines for Setting the Connect Failure Policy

Consider the following when using an open policy which permits full network access:

- Security and protection are not available until the VPN session is established; therefore, the endpoint device may get infected with web-based malware or sensitive data may leak.
- An open connect failure policy does not apply if you enable the Disconnect button and the user clicks **Disconnect**.

Consider the following when using a closed policy which disables all network connectivity until the VPN session is established:

- A closed policy can halt productivity if users require Internet access outside the VPN.
- The purpose of closed is to help protect corporate assets from network threats when resources in the private network that protect the endpoint are not available. The endpoint is protected from web-based malware and sensitive data leakage at all times because all network access is prevented except for local resources such as printers and tethered devices permitted by split tunneling.
- This option is primarily for organizations where security persistence is a greater concern than always-available network access.
- A closed policy prevents captive portal remediation unless you specifically enable it.
- You can allow the application of the local resource rules imposed by the most recent VPN session if **Apply Last VPN Local Resources** is enabled in the client profile. For example, these rules could determine access to active sync and local printing.
- The network is unblocked and open during an AnyConnect software upgrade when Always-On is enabled regardless of a closed policy.
- If you deploy a closed connection policy, we highly recommend that you follow a phased approach. For example, first deploy Always-On with a connect failure open policy and survey users for the frequency with which AnyConnect does not connect seamlessly. Then deploy a small pilot deployment of a connect failure closed policy among early-adopter users and solicit their feedback. Expand the pilot program gradually while continuing to solicit feedback before considering a full deployment. As you deploy a connect failure closed policy, be sure to educate the VPN users about the network access limitation as well as the advantages of a connect failure closed policy.



Caution A connect failure closed policy prevents network access if AnyConnect fails to establish a VPN session. Use extreme caution when implementing a connect failure closed policy.

Configure a Connect Failure Policy

You configure a Connect Failure Policy only when the Always-On feature is enabled. By default, the connect failure policy is closed, preventing Internet access if the VPN is unreachable. To allow Internet access in this situation the connect failure policy must be set to open.

Procedure

- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 2)** from the navigation pane.
- Step 2** Set the **Connect Failure Policy** parameter to one of the following settings:
- Closed—(Default) Restricts network access when the secure gateway is unreachable.
 - Open—Permits network access by browsers and other applications when the client cannot connect to the secure gateway.
- Step 3** If you specified a closed policy:
- a) [Configure Captive Portal Remediation](#).
 - b) Select **Apply Last VPN Local Resources** if you would like to retain the last VPN session's local device rules while network access is disabled.
-

Use Captive Portal Hotspot Detection and Remediation

About Captive Portals

Many facilities that offer Wi-Fi and wired access, such as airports, coffee shops, and hotels, require the user to pay before obtaining access, to agree to abide by an acceptable use policy, or both. These facilities use a technique called captive portal to prevent applications from connecting until the user opens a browser and accepts the conditions for access. Captive portal detection is the recognition of this restriction, and captive portal remediation is the process of satisfying the requirements of a captive portal hotspot in order to obtain network access.

Captive portals are detected automatically by AnyConnect when initiating a VPN connection requiring no additional configuration. Also, AnyConnect does not modify any browser configuration settings during captive portal detection and does not automatically remediate the captive portal. It relies on the end user to perform the remediation. AnyConnect reacts to the detection of a captive portal depending on the current configuration:

- If Always-On is disabled, or if Always-On is enabled and the Connect Failure Policy is open, the following message is displayed on each connection attempt:

```
The service provider in your current location is restricting access to the Internet.  
You need to log on with the service provider before you can establish a VPN session.  
You can try this by visiting any website with your browser.
```

The end user must perform captive portal remediation by meeting the requirements of the provider of the hotspot. These requirements could be paying a fee to access the network, signing an acceptable use policy, both, or some other requirement defined by the provider.

- If Always-On is enabled and the connect failure policy is closed, captive portal remediation needs to be explicitly enabled. If enabled, the end user can perform remediation as described above. If disabled, the following message is displayed upon each connection attempt, and the VPN cannot be connected.

```
The service provider in your current location is restricting access to the Internet.
```


The AnyConnect protection settings must be lowered for you to log on with the service provider. Your current enterprise security policy does not allow this.

Configure Captive Portal Remediation

You configure captive portal remediation only when the Always-On feature is enabled and the Connect Failure Policy is set to closed. In this situation, configuring captive portal remediation allows AnyConnect to connect to the VPN when a captive portal is preventing it from doing so.



Note Configuration of captive portal remediation is not applicable to Linux, since Always On is not supported on this platform. Therefore, regardless of the *Allow Captive Portal Remediation Always On* setting in the profile editor, the Linux user can remediate a captive portal.

If the Connect Failure Policy is set to open or Always-On is not enabled, your users are not restricted from network access and are capable of remediating a captive portal without any specific configuration in the AnyConnect VPN client profile.

By default, captive portal remediation is disabled on platforms supporting Always on (Windows and macOS) to provide the greatest security. AnyConnect does not provide data leakage protection capabilities during the captive portal remediation phase. If data loss protection is desired, you should employ a relevant endpoint security product.

Procedure

Step 1 Open the VPN Profile Editor and choose **Preferences (Part 1)** from the navigation pane.

Step 2 Select **Allow Captive Portal Remediation**.

This setting lifts the network access restrictions imposed by the closed connect failure policy.

Step 3 Specify the Remediation Timeout.

Enter the number of minutes for which AnyConnect lifts the network access restrictions. The user needs enough time to satisfy the captive portal requirements.

Troubleshoot Captive Portal Detection and Remediation

AnyConnect can falsely assume that it is in a captive portal in the following situations.

- If AnyConnect attempts to contact an ASA with a certificate containing an incorrect server name (CN), then the AnyConnect client will think it is in a “captive portal” environment.

To prevent this, make sure the ASA certificate is properly configured. The CN value in the certificate must match the name of the ASA server in the VPN client profile.

- If there is another device on the network before the ASA, and that device responds to the client's attempt to contact an ASA by blocking HTTPS access to the ASA, then the AnyConnect client will think it is in a “captive portal” environment. This situation can occur when a user is on an internal network, and connects through a firewall to connect to the ASA.

If you need to restrict access to the ASA from inside the corporation, configure your firewall such that HTTP and HTTPS traffic to the ASA's address does not return an HTTP status. HTTP/HTTPS access to the ASA should either be allowed or completely blocked to ensure that HTTP/HTTPS requests sent to the ASA will not return an unexpected response.

If users cannot access a captive portal remediation page, ask them to try the following:

- Terminate any applications that use HTTP, such as instant messaging programs, e-mail clients, IP phone clients, and all but one browser to perform the remediation.

The captive portal may be actively inhibiting DoS attacks by ignoring repetitive attempts to connect, causing them to time out on the client end. The attempt by many applications to make HTTP connections exacerbates this problem.

- Disable and re-enable the network interface. This action triggers a captive portal detection retry.
- Restart the computer.

Configure AnyConnect over L2TP or PPTP

ISPs in some countries require support of the Layer 2 Tunneling Protocol (L2TP) and Point-to-Point Tunneling Protocol (PPTP).

To send traffic destined for the secure gateway over a Point-to-Point Protocol (PPP) connection, AnyConnect uses the point-to-point adapter generated by the external tunnel. When establishing a VPN tunnel over a PPP connection, the client must exclude traffic destined for the ASA from the tunneled traffic intended for destinations beyond the ASA. To specify whether and how to determine the exclusion route, use the PPP Exclusion setting in the AnyConnect profile. The exclusion route appears as a non-secured route in the Route Details display of the AnyConnect GUI.

Procedure

- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 2)** from the navigation pane.
- Step 2** Choose a **PPP Exclusion** method. Also, check **User Controllable** for this field to let users view and change this setting:
- **Automatic**—Enables PPP exclusion. AnyConnect automatically determines the IP address of the PPP server.
 - **Override**—Enables PPP Exclusion using a predefined server IP address specified in the *PPP Exclusion Server IP* field. The *PPP Exclusion Server IP* field is only applicable to this Override method and should only be used when the Automatic options fails to detect the IP address of the PPP server.

Checking **User Controllable** for the PPP Exclusion Server IP field allows the end user to manually update the IP address via the preferences.xml file. Refer to the [Instruct Users to Override PPP Exclusion, on page 19](#) section.
 - **Disabled**—PPP exclusion is not applied.
-

Instruct Users to Override PPP Exclusion

If automatic detection does not work and you configured the PPP Exclusion fields as user controllable, the user can override the setting by editing the AnyConnect preferences file on the local computer.

Procedure

- Step 1** Use an editor such as Notepad to open the preferences XML file.
- This file is at one of the following paths on the user's computer:
- Windows: %LOCAL_APPDATA%\Cisco\Cisco AnyConnect Secure Mobility Client\preferences.xml.
For example,
 - macOS: /Users/username/.anyconnect
 - Linux: /home/username/.anyconnect
- Step 2** Insert the PPPEXclusion details under <ControllablePreferences>, while specifying the Override value and the IP address of the PPP server. The address must be a well-formed IPv4 address. For example:
- ```
<AnyConnectPreferences>
<ControllablePreferences>
<PPPEXclusion>Override
<PPPEXclusionServerIP>192.168.22.44</PPPEXclusionServerIP></PPPEXclusion>
</ControllablePreferences>
</AnyConnectPreferences>
```
- Step 3** Save the file.
- Step 4** Exit and restart AnyConnect.
- 

## Configure AnyConnect Proxy Connections

### About AnyConnect Proxy Connections

AnyConnect supports VPN sessions through Local, Public, and Private proxies:

- Local Proxy Connections:

A local proxy runs on the same PC as AnyConnect, and is sometimes used as a transparent proxy. Some examples of a transparent proxy service include acceleration software provided by some wireless data cards, or a network component on some antivirus software, such as Kaspersky.

The use of a local proxy is enabled or disabled in the AnyConnect VPN client profile, see [Allow a Local Proxy Connection](#).

- Public Proxy Connections:

Public proxies are usually used to anonymize web traffic. When Windows is configured to use a public proxy, AnyConnect uses that connection. Public proxy is supported on macOS and Linux for both native and override.

Configuring a public proxy is described in [Public Proxy, on page 21](#).

- Private Proxy Connections:

Private proxy servers are used on a corporate network to prevent corporate users from accessing certain Web sites based on corporate usage policies, for example, pornography, gambling, or gaming sites.

You configure a group policy to download private proxy settings to the browser after the tunnel is established. The settings return to their original state after the VPN session ends. See [Configure a Private Proxy Connection, on page 22](#).




---

**Note** AnyConnect SBL connections through a proxy server are dependent on the Windows operating system version and system (machine) configuration or other third-party proxy software capabilities; therefore, refer to system wide proxy settings as provided by Microsoft or whatever third-party proxy application you use.

---

### Control Client Proxy with VPN Client Profile

The VPN Client profile can block or redirect the client system's proxy connection. For Windows and Linux, you can configure, or you can allow the user to configure, the address of a public proxy server.

For more information about configuring the proxy settings in the VPN client profile, see [AnyConnect Profile Editor, Preferences \(Part 2\)](#).

### Proxy Auto-Configuration File Generation for Clientless Support

Some versions of the ASA require AnyConnect configuration to support clientless portal access through a proxy server after establishing an AnyConnect session. AnyConnect uses a proxy auto-configuration (PAC) file to modify the client-side proxy settings to let this occur. AnyConnect generates this file only if the ASA does not specify private-side proxy settings.

## Requirements for AnyConnect Proxy Connections

OS support of proxy connections varies as shown:

Proxy Connection Type	Windows	macOS	Linux
Local Proxy	Yes	Yes (Override& Native)	Yes
Private Proxy	Yes (on Internet Explorer)	Yes (set as system proxy settings)	No
Public Proxy	Yes (IE and Override)	Yes (Override & Native)	Yes (Override& Native)

## Limitations on Proxy Connections

- IPv6 proxies are not supported for any type of proxy connection.
- Connecting through a proxy is not supported with the Always-On feature enabled.

- A VPN client profile is required to allow access to a local proxy.

## Allow a Local Proxy Connection

### Procedure

---

- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 2)** from the navigation pane.
- Step 2** Select (default) or unselect **Allow Local Proxy Connections**. Local proxy is disabled by default.
- 

## Public Proxy

Public proxies are supported on Windows and Linux platforms. Proxy servers are chosen based on preferences set in the client profile. In case of proxy override, AnyConnect extracts proxy servers from the profile. With release 4.1 we added proxy support on macOS along with Native-proxy configuration on Linux and macOS.

On Linux, native-proxy settings are exported before AnyConnect runs. If you change the settings, a restart must happen.

Authenticating Proxy Servers requires a username and password. AnyConnect supports Basic and NTLM authentication when the proxy server is configured to require authentication. AnyConnect dialogs manage the authentication process. After successfully authenticating to the proxy server, AnyConnect prompts for the ASA username and password.

## Configure a Public Proxy Connection, Windows

Follow these steps to configure a public proxy connection on Windows.

### Procedure

---

- Step 1** Open **Internet Options** from Internet Explorer or the Control Panel.
- Step 2** Select the **Connections** Tab, and click the **LAN Settings** button.
- Step 3** Configure the LAN to use a proxy server, and enter the IP address of the proxy server.
- 

## Configure a Public Proxy Connection, macOS

### Procedure

---

- Step 1** Go to system preferences and choose the appropriate interface on which you are connected.
- Step 2** Click **Advanced**.
- Step 3** Choose **Proxies** tab from the new window.
- Step 4** Enable HTTPS proxy.

- Step 5** Enter the proxy server address in the Secure Proxy Server field on the right panel.
- 

## Configure a Public Proxy Connection, Linux

To configure a public proxy connection in Linux, you must set an environment variable.

## Configure a Private Proxy Connection

### Procedure

---

- Step 1** Configure the private proxy information in the ASA group policy. See the [Configuring a Browser Proxy for an Internal Group Policy](#) section in the *Cisco ASA Series VPN Configuration Guide*.

**Note** In a macOS environment, the proxy information that is pushed down from the ASA (upon a VPN connection) is not viewed in the browser until you open up a terminal and issue a `scutil --proxy`.

- Step 2** (Optional) [Configure the Client to Ignore Browser Proxy Settings](#).

- Step 3** (Optional) [Lock Down the Internet Explorer Connections Tab](#).
- 

## Configure the Client to Ignore Browser Proxy Settings

You can specify a policy in the AnyConnect profile to bypass the Microsoft Internet Explorer or Safari proxy configuration settings on the user's PC. This prevents the user from establishing a tunnel from outside the corporate network, and prevents AnyConnect from connecting through an undesirable or illegitimate proxy server.

### Procedure

---

- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 2)** from the navigation pane.

- Step 2** In the Proxy Settings drop-down list, choose **IgnoreProxy**. Ignore Proxy causes the client to ignore all proxy settings. No action is taken against proxies that are downloaded from the ASA.
- 

## Lock Down the Internet Explorer Connections Tab

Under certain conditions, AnyConnect hides the Internet Explorer Tools > Internet Options > Connections tab. When exposed, this tab lets the user set proxy information. Hiding this tab prevents the user from intentionally or unintentionally circumventing the tunnel. The tab lockdown is reversed on disconnect, and it is superseded by any administrator-defined policies applied to that tab. The conditions under which this lock down occurs are the following:

- The ASA configuration specifies Connections tab lockdown.
- The ASA configuration specifies a private-side proxy.

- A Windows group policy previously locked down the Connections tab (overriding the no lockdown ASA group policy setting).

You can configure the ASA to allow or not allow proxy lockdown, in the group policy. To do this using ASDM, follow this procedure:

#### Procedure

---

- Step 1** In ASDM go to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**.
  - Step 2** Select a group policy and click **Edit** or **Add** a new group policy.
  - Step 3** In the navigation pane, go to **Advanced > Browser Proxy**. The Proxy Server Policy pane displays.
  - Step 4** Click **Proxy Lockdown** to display more proxy settings.
  - Step 5** Uncheck **Inherit** and select **Yes** to enable proxy lockdown and hide the Internet Explorer Connections tab for the duration of the AnyConnect session or; select **No** to disable proxy lockdown and expose the Internet Explorer Connections tab for the duration of the AnyConnect session.
  - Step 6** Click **OK** to save the Proxy Server Policy changes.
  - Step 7** Click **Apply** to save the Group Policy changes.
- 

## Verify the Proxy Settings

- For Windows: Find the proxy settings in the registry under:

```
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Internet Settings
```

- For macOS: Open a terminal window, and type:

```
scutil --proxy
```

## Select and Exclude VPN Traffic

### Configure IPv4 or IPv6 Traffic to Bypass the VPN

You can configure how the AnyConnect client manages IPv4 traffic when the ASA is expecting only IPv6 traffic or how AnyConnect manages IPv6 traffic when the ASA is only expecting IPv4 traffic using the Client Bypass Protocol setting.

When the AnyConnect client makes a VPN connection to the ASA, the ASA can assign the client an IPv4, IPv6, or both an IPv4 and IPv6 address.

If Client Bypass Protocol is enabled for an IP protocol and an address pool is not configured for that protocol (in other words, no IP address for that protocol was assigned to client by the ASA), any IP traffic using that protocol will not be sent through the VPN tunnel. It will be sent outside the tunnel.

If Client Bypass Protocol is disabled, and an address pool is not configured for that protocol, the client drops all traffic for that IP protocol once the VPN tunnel is established.

For example, assume that the ASA assigns only an IPv4 address to an AnyConnect connection and the endpoint is dual stacked. When the endpoint attempts to reach an IPv6 address, if Client Bypass Protocol is disabled, the IPv6 traffic is dropped. If Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear.

If establishing an IPsec tunnel (as opposed to an SSL connection), the ASA is not notified whether or not IPv6 is enabled on the client, so ASA always pushes down the client bypass protocol setting.

You configure the Client Bypass Protocol on the ASA in the group policies.

### Procedure

- 
- Step 1** In ASDM go to **Configuration > Remote Access VPN > Network (Client) Access > Group Policies**.
- Step 2** Select a group policy and click **Edit** or **Add** a new group policy.
- Step 3** Select **Advanced > AnyConnect**.
- Step 4** Next to **Client Bypass Protocol**, uncheck **Inherit** if this is a group policy other than the default group policy.
- Step 5** Choose one of these options:
- Click **Disable** to drop IP traffic for which the ASA did not assign an address.
  - Click **Enable** to send that IP traffic in the clear.
- Step 6** Click **OK**.
- Step 7** Click **Apply**.
- 

## Configure a Client Firewall with Local Printer and Tethered Device Support

See the [Client Firewall with Local Printer and Tethered Device Support](#) section in the *Cisco ASA Series Configuration Guide*.

## About Dynamic Split Tunneling

Dynamic split tunneling was designed to enhance the current split tunneling option, which is configured with the Exclude Network List Below option in ASDM group policy configuration. Beyond the static exclusions typically used to define split tunneling, the dynamic split tunneling exclusions address scenarios when traffic pertaining to a certain service needs to be excluded from the VPN tunnel dynamically, at run time. Multiple cloud-based services may be hosted on the same IP pool and may resolve to different IP addresses based on the location of the user or the load of cloud-hosted compute resources. Administrators who only want to exclude a single such service from the VPN tunnel would have a difficult time defining such a policy using static exclusions, especially when ISP NAT, 6to4, 4to6, and other network translation schemes are also considered.

The limits also vary from static split tunneling to dynamic split tunneling. For static split tunneling, the limit is 2500 networks/ACEs per IP protocol. With dynamic split tunneling, the limit goes to 5000 characters (about 400 domain names) and is only enforced via truncation on the client.



## Interoperability Between Static Split Tunneling and Dynamic Split Tunneling

Both static and dynamic tunnel exclusions (or inclusions) can coexist. While static split tunneling is applied when the tunnel is established, dynamic split tunneling is applied when the traffic to the excluded (or included) domain occurs, while the tunnel is already connected. Dynamic split tunneling applies to "tunnel all," "split include," and "split exclude" tunneling:

- Tunnel All Networks—All exclusions from the VPN tunnel are dynamic.
- Exclude Specific Networks—Dynamic exclusions are added to preconfigured static ones.
- Include Specific Networks—Dynamic exclusions are only relevant if at least one IP address of the excluded host names overlaps with a split include network. Otherwise, the traffic is already excluded from the VPN tunnel, and no dynamic exclusion is performed.

Enhanced dynamic split exclude tunneling applies to "tunnel all" and "split exclude" tunneling. If both dynamic split exclude and dynamic split include domains, as well as split include tunneling, are configured, the resulting configuration is enhanced dynamic split include tunneling.

Dynamic exclusion covers only IP addresses not already excluded. When both static and dynamic exclusions are applied and a new dynamic exclusion needs to be enforced, a collision with an already applied exclusion may occur. When a dynamic exclusion is enforced (which contains all IP addresses that are part of a DNS response matching an excluded domain name), only those addresses not already excluded are considered for exclusion.

Umbrella Roaming Security protection is active when either static or dynamic split tunneling is enabled. You may have to statically exclude the Umbrella cloud resolvers from the VPN tunnel, unless they are reachable and can be probed by the VPN tunnel.

## Notifications of Dynamic Split Tunneling Usage

While the VPN tunnel is connected, you can see what is set for dynamic split tunneling in several ways:

- Statistics tab—Displays Dynamic Tunnel Exclusions, containing the domain names excluded from the VPN tunnel, as configured in the ASA group policy.
- Export Stats—Produces a file that contains the domain names excluded from the VPN tunneling, along with the tunnel modes for both IPv4 and IPv6.
- Route Details tab—Shows the IPv4 and IPv6 dynamic split exclude routes with the host names that correspond to each excluded IP address.




---

**Note** The AnyConnect UI only displays up to 200 per IP protocol of the secured or non-secured routes enforced by AnyConnect VPN. In excess of 200 routes, truncation occurs, and you can run either `route print` on Windows or `onnetstat -rn` on Linux or macOS to view all routes.

---

- VPN configuration log message—Shows the number of domains excluded from the VPN tunnel.

## Configure Dynamic Split Tunneling

### Before you begin

Refer to [About Dynamic Split Tunneling, on page 24](#).

With dynamic split tunneling, you can dynamically provision split exclude tunneling after tunnel establishment based on the host DNS domain name. Dynamic split tunneling is configured by creating custom attribute and adding it to a group policy on ASA. Refer to [Configure Dynamic Split Tunneling](#) in the *Cisco ASA Series VPN ASDM Configuration Guide* for GUI steps.

## Split DNS

When split DNS is configured in the Network (Client) Access group policy, AnyConnect tunnels specific DNS queries to the private DNS server (also configured in the group policy). All other DNS queries go to the DNS resolver on the client operating system, in the clear, for DNS resolution. If split DNS is not configured, AnyConnect tunnels all DNS queries.

If split DNS is not configured, AnyConnect tunnels all DNS queries.

## Requirements for Split DNS

Split DNS supports standard and update queries (including A, AAAA, NS, TXT, MX, SOA, ANY, SRV, PTR, and CNAME). PTR queries matching any of the tunneled networks are allowed through the tunnel.

Split DNS is supported on Windows and macOS platforms.

- Limited support is available on Linux, namely only tunneled DNS requests are subject to the split DNS policy. Consequently, some DNS requests sent outside the tunnel may not comply with the split DNS policy.

For macOS, AnyConnect can use true split-DNS for a certain IP protocol only if one of the following conditions is met:

- Split-DNS is configured for one IP protocol (such as IPv4), and Client Bypass Protocol is configured for the other IP protocol (such as IPv6) in the group policy (with no address pool configured for the latter IP protocol).
- Split-DNS is configured for both IP protocols.

## Configure Split DNS for Split Include Tunneling

To configure split DNS for split include tunneling in the group policy, do the following:

### Procedure

---

**Step 1** Configure at least one DNS server.

See the *Configure Server Attributes for an Internal Group Policy* section in the [Cisco ASA Series VPN Configuration Guide](#).

Ensure the private DNS servers specified do not overlap with the DNS servers configured for the client platform. If they do, name resolution may not function properly.

**Step 2** Configure split-include tunneling:

On the **Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Advanced > Split Tunneling** pane, choose the **Tunnel Network List Below** policy, and specify a **Network List** of addresses to be tunneled.

Split-DNS does not support the Exclude Network List Below split-tunneling policy. You must use the Tunnel Network List Below split-tunneling policy to configure split-DNS.

**Step 3** On the **Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Advanced > Split Tunneling** pane, uncheck **Send All DNS lookups through tunnel**, and specify the names of the domains whose queries will be tunneled in **DNS Names**.

---

#### What to do next

After making changes to the group policy in ASDM, be sure the group policy is associated with a Connection Profile in **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles > Add/Edit > Group Policy**.

## Verify Split DNS Using AnyConnect Logs

### Check Which Domains Use Split DNS

You can use any tool or application that relies on the operating system's DNS resolver for domain name resolution. For example, you can use a ping or web browser to test the split DNS solution. Other tools such as nslookup or dig circumvent the OS DNS resolver.

To use the client to check which domains are used for split DNS, follow these steps:

#### Procedure

---

**Step 1** Run **ipconfig/all** and record the domains listed next to DNS Suffix Search List.

**Step 2** Establish a VPN connection and again check the domains listed next to DNS Suffix Search List.

Those extra domains added after establishing the tunnel are the domains used for split DNS.

**Note** This process assumes that the domains pushed from the ASA do not overlap with the ones already configured on the client host.

---

## Manage VPN Authentication

### Important Security Considerations

We do not recommend using a self-signed certificate on your secure gateway

- because of the possibility that a user could inadvertently configure a browser to trust a certificate on a rogue server, and
- because of the inconvenience to users of having to respond to a security warning when connecting to your secure gateway.

We strongly recommend that you enable Strict Certificate Trust for the AnyConnect client. To configure **Strict Certificate Trust**, see the *Local Policy Parameters and Values* section: [Local Policy Preferences](#).

## Configure Server Certificate Handling

### Server Certificate Verification

- (Windows only) For both SSL and IPsec VPN connections, you have the option to perform Certificate Revocation List (CRL) checking. When enabled in the profile editor, AnyConnect retrieves the updated CRL for all certificates in the chain. It then verifies whether the certificate in question is among those revoked certificates which should no longer be trusted; and if found to be a certificate revoked by the Certificate Authority, it does not connect. Refer to [Local Policy Preferences](#) for further information.
- When a user connects to an ASA that is configured with a server certificate, the checkbox to trust and import that certificate will still display, even if there is a problem with the trust chain (Root, Intermediate, etc.) If there are any other certificate problems, that checkbox will not display.
- SSL connections being performed via FQDN do not make a secondary server certificate verification with the FQDN's resolved IP address for name verification if the initial verification using the FQDN fails.
- IPsec and SSL connections require that if a server certificate contains Key Usage, the attributes must contain DigitalSignature AND (KeyAgreement OR KeyEncipherment). If the server certificate contains an EKU, the attributes must contain serverAuth (for SSL and IPsec) or ikeIntermediate (for IPsec only). Note that server certificates are not required to have a KU or an EKU to be accepted.
- IPsec and SSL connections perform name verification on server certificates. The following rules are applied for the purposes of IPsec and SSL name verification:
  - If a Subject Alternative Name extension is present with relevant attributes, name verification is performed solely against the Subject Alternative Name. Relevant attributes include DNS Name attributes for all certificates, and additionally include IP address attributes if the connection is being performed to an IP address.
  - If a Subject Alternative Name extension is not present, or is present but contains no relevant attributes, name verification is performed against any Common Name attributes found in the Subject of the certificate.
  - If a certificate uses a wildcard for the purposes of name verification, the wildcard must be in the first (left-most) subdomain only, and additionally must be the last (right-most) character in the subdomain. Any wildcard entry not in compliance is ignored for the purposes of name verification.
- For OSX, expired certificates are displayed only when Keychain Access is configured to “Show Expired Certificates.” Expired certificates are hidden by default, which may confuse users.

## Invalid Server Certificate Handling

In response to the increase of targeted attacks against mobile users on untrusted networks, we have improved the security protections in the client to help prevent serious security breaches. The default client behavior has been changed to provide an extra layer of defense against Man-in-the-middle attacks.

### User Interaction


When the user tries to connect to a secure gateway, and there is a certificate error (due to expired, invalid date, wrong key usage, or CN mismatch), the user sees a red-colored dialog with Change Settings and Keep Me Safe buttons.



**Note** The dialogs for Linux may look different from the ones shown in this document.



- Clicking **Keep Me Safe** cancels the connection.
- Clicking **Change Settings** opens AnyConnect's Advanced > VPN > Preferences dialog, where the user can enable connections to untrusted servers. The current connection attempt is canceled.



If the user un-checks **Block connections to untrusted servers**, and the only issue with the certificate is that the CA is untrusted, then the next time the user attempts to connect to this secure gateway, the user will not see the Certificate Blocked Error Dialog dialog; they only see the following dialog:



If the user checks **Always trust this VPN server and import the certificate**, then future connections to this secure gateway will not prompt the user to continue.



---

**Note** If the user checks **Block connections to untrusted servers** in **AnyConnect Advanced > VPN > Preferences**, or if the user's configuration meets one of the conditions in the list of the modes described under the guidelines and limitations section, then AnyConnect rejects invalid server certificates and connections to untrusted servers, regardless of whether the Strict Certificate Trust option in the Profile Editor is enabled.

---

### Improved Security Behavior

When the client accepts an invalid server certificate, that certificate is saved in the client's certificate store. Previously, only the thumbprint of the certificate was saved. Note that invalid certificates are saved only when the user has elected to always trust and import invalid server certificates.

There is no administrative override to make the end user less secure automatically. To completely remove the preceding security decisions from your end users, enable **Strict Certificate Trust** in the user's local policy file. When Strict Certificate Trust is enabled, the user sees an error message, and the connection fails; there is no user prompt.

For information about enabling Strict Certificate Trust in the local policy file, see the *AnyConnect Local Policy Parameters and Values* section: [Local Policy Preferences](#).

### Guidelines and Limitations

Invalid server certificates are rejected when:

- Always On is enabled in the AnyConnect VPN client profile and is not turned off by an applied group policy or DAP.
- The client has a Local Policy with Strict Certificate Trust enabled.
- AnyConnect is configured to start before logon.
- A client certificate from the machine certificate store is used for authentication.

## Configure Certificate-Only Authentication

You can specify whether you want users to authenticate using AAA with a username and password or using a digital certificate (or both). When you configure certificate-only authentication, users can connect with a digital certificate and are not required to provide a user ID and password.

To support certificate-only authentication in an environment where multiple groups are used, you may provision more than one group-url. Each group-url would contain a different client profile with some piece of customized data that would allow for a group-specific certificate map to be created. For example, the Department\_OU value of Engineering could be provisioned on the ASA to place the user in this group when the certificate from this process is presented to the ASA.



---

**Note** The certificate used to authenticate the client to the secure gateway must be valid and trusted (signed by a CA). A self-signed client certificate will not be accepted.

---

## Procedure

---

- Step 1** Go to **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles**. Select a connection profile and click Edit. The Edit AnyConnect Connection Profile window opens.
- Step 2** If it is not already, click the **Basic** node of the navigation tree on the left pane of the window. In the right pane of the window, in the **Authentication** area, enable the method **Certificate**.
- Step 3** Click **OK** and apply your changes.
- 

## Configure Certificate Enrollment

The Cisco AnyConnect Secure Mobility Client uses the Simple Certificate Enrollment Protocol (SCEP) to provision and renew a certificate as part of client authentication. Certificate enrollment using SCEP is supported by AnyConnect IPsec and SSL VPN connections to the ASA in the following ways:

- **SCEP Proxy:** The ASA acts as a proxy for SCEP requests and responses between the client and the Certificate Authority (CA).
  - The CA must be accessible to the ASA, not the AnyConnect client, since the client does not access the CA directly.
  - Enrollment is always initiated automatically by the client. No user involvement is necessary.

### Related Topics

[AnyConnect Profile Editor, Certificate Enrollment](#)

## SCEP Proxy Enrollment and Operation

The following steps describe how a certificate is obtained and a certificate-based connection is made when AnyConnect and the ASA are configured for SCEP Proxy.

1. The user connects to the ASA headend using a connection profile configured for both certificate and AAA authentication. The ASA requests a certificate and AAA credentials for authentication from the client.
2. The user enters his/her AAA credentials, but a valid certificate is not available. This situation triggers the client to send an automatic SCEP enrollment request after the tunnel has been established using the entered AAA credentials.
3. The ASA forwards the enrollment request to the CA and returns the CA's response to the client.
4. If SCEP enrollment is successful, the client presents a (configurable) message to the user and disconnects the current session. The user can now connect using certificate authentication to an ASA tunnel group.

If SCEP enrollment fails, the client displays a (configurable) message to the user and disconnects the current session. The user should contact his/her administrator.

Other SCEP Proxy operational considerations:

- If configured to do so, the client automatically renews the certificate before it expires, without user intervention.
- SCEP Proxy enrollment uses SSL for both SSL and IPsec tunnel certificate authentication.



## Certificate Authority Requirements

- All SCEP-compliant CAs, including IOS CS, Windows Server 2003 CA, and Windows Server 2008 CA, are supported.
- The CA must be in auto-grant mode; polling for certificates is not supported.
- You can configure some CAs to email users an enrollment password for an additional layer of security. The CA password is the challenge password or token that is sent to the certificate authority to identify the user. The password can then be configured in the AnyConnect client profile, which becomes part of SCEP request that the CA verifies before granting the certificate.

## Guidelines for Certificate Enrollment

- Clientless (browser-based) VPN access to the ASA does not support SCEP proxy, but WebLaunch (clientless-initiated AnyConnect) does.
- ASA Load balancing is supported with SCEP enrollment.
- The ASA does not indicate why an enrollment failed, although it does log the requests received from the client. Connection problems must be debugged on the CA or the client.
- Certificate-Only Authentication and Certificate Mapping on the ASA:  
To support certificate-only authentication in an environment where multiple groups are used, you may provision more than one group-url. Each group-url would contain a different client profile with some piece of customized data that would allow for a group-specific certificate map to be created. For example, the Department\_OU value of Engineering could be provisioned on the ASA to place the user in this tunnel group when the certificate from this process is presented to the ASA.
- Identifying Enrollment Connections to Apply Policies:  
On the ASA, the aaa.cisco.sceprequired attribute can be used to catch the enrollment connections and apply the appropriate policies in the selected DAP record.
- Windows Certificate Warning:  
When Windows clients first attempt to retrieve a certificate from a certificate authority they may see a warning. When prompted, users must click Yes. This allows them to import the root certificate. It does not affect their ability to connect with the client certificate.

## Configure SCEP Proxy Certificate Enrollment

### Configure a VPN Client Profile for SCEP Proxy Enrollment

#### Procedure

---

- Step 1** Open the VPN Profile Editor and choose **Certificate Enrollment** from the navigation pane.
- Step 2** Select **Certificate Enrollment**.
- Step 3** Configure the **Certificate Contents** to be requested in the enrollment certificate. For definitions of the certificate fields, see [AnyConnect Profile Editor, Certificate Enrollment](#).

- Note**
- If you use %machineid%, then Hostscan/Posture must be loaded for the desktop client.
  - For mobile clients, at least one certificate field must be specified.

---

## Configure the ASA to Support SCEP Proxy Enrollment

For SCEP Proxy, a single ASA connection profile supports certificate enrollment and the certificate authorized VPN connection.

### Procedure

---

- Step 1** Create a group policy, for example, cert\_group. Set the following fields:
- On General, enter the URL to the CA in **SCEP Forwarding URL**.
  - On the Advanced > AnyConnect Client pane, uncheck **Inherit** for Client Profiles to Download and specify the client profile configured for SCEP Proxy. For example, specify the ac\_vpn\_scep\_proxy client profile.
- Step 2** Create a connection profile for certificate enrollment and certificate authorized connection, for example, cert\_tunnel.
- Authentication: Both (AAA and Certificate).
  - Default Group Policy: cert\_group.
  - On Advanced > General, check **Enable SCEP Enrollment for this Connction Profile**.
  - On Advanced > GroupAlias/Group URL, create a Group URL containing the group (cert\_group) for this connection profile.
- 

## Set Up a Windows 2008 Server Certificate Authority for SCEP

If your Certificate Authority software is running on a Windows 2008 server, you may need to make one of the following configuration changes to the server to support SCEP with AnyConnect.

### Disable the SCEP Password on the Certificate Authority

The following steps describe how to disable the SCEP challenge password, so that clients will not need to provide an out-of-band password before SCEP enrollment.

### Procedure

---

- Step 1** On the Certificate Authority server, launch the Registry Editor. You can do this by selecting **Start > Run**, typing **regedit**, and clicking **OK**.
- Step 2** Navigate to HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Cryptography\MSCEP\EnforcePassword.

If the EnforcePassword key does not exist, create it as a new Key.

- Step 3** Edit EnforcePassword, and set it to '0'. If it does not exist, create it as a REG-DWORD.
- Step 4** Exit regedit, and reboot the certificate authority server.
- 

## Setting the SCEP Template on the Certificate Authority

The following steps describe how to create a certificate template, and assign it as the default SCEP template.

### Procedure

---

- Step 1** Launch the Server Manager. You can do this by selecting Start > Admin Tools > Server Manager.
- Step 2** Expand Roles > Certificate Services (or AD Certificate Services).
- Step 3** Navigate to CA Name > Certificate Templates.
- Step 4** Right-click **Certificate Templates > Manage**.
- Step 5** From the Cert Templates Console, right-click User template and choose **Duplicate**
- Step 6** Choose **Windows Server 2008 version** for new template, and click **OK**.
- Step 7** Change the template display name to something descriptive, such as NDES-IPSec-SSL.
- Step 8** Adjust the Validity Period for your site. Most sites choose three or more years to avoid expired certificates.
- Step 9** On the Cryptography tab, set the minimum key size for your deployment.
- Step 10** On the Subject Name tab, select **Supply in Request**.
- Step 11** On the Extensions tab, set the Application Policies to include at least:
- Client Authentication
  - IP security end system
  - IP security IKE intermediate
  - IP security tunnel termination
  - IP security user
- These values are valid for SSL or IPsec.
- Step 12** Click **Apply**, then **OK** to save new template.
- Step 13** From Server manager > Certificate Services-CA Name, right-click Certificate Templates. Select New > Certificate Template to Issue, select the new template you created (in this example, NDES-IPSec-SSL), and click **OK**.
- Step 14** Edit the registry. You can do this by selecting Start > Run, regedit, and clicking **OK**.
- Step 15** Navigate to HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Cryptography\MSCEP.
- Step 16** Set the value of the following three keys to **NDES-IPSec-SSL**.
- EncryptionTemplate
  - GeneralPurposeTemplate
  - SignatureTemplate

**Step 17** Click **Save**, and reboot the certificate authority server.

---

## Configure a Certificate Expiration Notice

Configure AnyConnect to warn users that their authentication certificate is about to expire. The **Certificate Expiration Threshold** setting specifies the number of days before the certificate's expiration date that AnyConnect warns users that their certificate is expiring. AnyConnect warns the user upon each connect until the certificate has actually expired or a new certificate has been acquired.



**Note** The Certificate Expiration Threshold feature cannot be used with RADIUS.

---

### Procedure

---

**Step 1** Open the VPN Profile Editor and choose **Certificate Enrollment** from the navigation pane.

**Step 2** Select **Certificate Enrollment**.

**Step 3** Specify a **Certificate Expiration Threshold**.

This is the number of days before the certificate expiration date, that AnyConnect warns users that their certificate is going to expire.

The default is 0 (no warning displayed). The range is 0 to 180 days.

**Step 4** Click **OK**.

---

## Configure Certificate Selection

The following steps show all the places in the AnyConnect profiles where you configure how certificates are searched for and how they are selected on the client system. None of the steps are required, and if you do not specify any criteria, AnyConnect uses default key matching.

AnyConnect reads the browser certificate stores on Windows. For Linux, you must create a Privacy Enhanced Mail (PEM) formatted file store. For macOS, you may use a Privacy Enhanced Mail (PEM) formatted file store or the Keychain.

### Procedure

---

**Step 1** Windows and macOS: [Configure Which Certificate Stores to Use, on page 37](#)

Specify which certificate stores are used by AnyConnect in the VPN client profile.

**Step 2** Windows Only: [Prompt Windows Users to Select Authentication Certificate, on page 39](#)

Configure AnyConnect to present a list of valid certificates to users and let them choose the certificate to authenticate the session.

- Step 3** For macOS and Linux environments: [Create a PEM Certificate Store for macOS and Linux, on page 40](#)
- Step 4** For macOS and Linux environments: Select which certificate stores to exclude in the VPN Local Policy profile.
- Step 5** [Configure Certificate Matching, on page 41](#)  
 Configure keys that AnyConnect tries to match, when searching for a certificate in the store. You can specify keys, extended keys, and add custom extended keys. You can also specify a pattern for the value of an operator in a distinguished name for AnyConnect to match.

## Configure Which Certificate Stores to Use

For Windows and macOS, separate certificate stores are provided for AnyConnect to use in the VPN client profile. You can have single or multiple certificate authentication combinations and can configure the secure gateway to dictate to the client which one of the multiple certificate authentication choices is acceptable for a particular VPN connection. For example, if you set `ExcludeMacNativeCertStore` to `true` in the local policy file (to force AnyConnect to use only file certificate stores such as user and system file certificate stores) and also set the profile-based certificate store to `Login` (to force AnyConnect to use only certificate stores such as login and dynamic smartcard keychains, plus the user file store), the combined filtering results in AnyConnect using strictly the user file certificate store.

Users with administrative privileges on the computer have access to both certificate stores. Users without administrative privileges only have access to the user certificate store. Usually, Windows users do not have administrative privileges. Selecting **Certificate Store Override** allows AnyConnect to access the machine store, even when the user does not have administrative privileges.



**Note** Access-control for the machine store can vary depending on the Windows version and security settings. Because of this, the user may be unable to use certificates in the machine store even though they have administrative privileges. In this case, select **Certificate Store Override** to allow machine store access.

The following table describes how AnyConnect searches for certificates on a client based on what **Certificate Store** is searched, and whether **Certificate Store Override** is checked.

Certificate Store Setting	Certificate Store Override Setting	AnyConnect Search Strategy
All (for Windows)	cleared	AnyConnect searches all certificate stores. AnyConnect is not allowed to access the machine store when the user does not have administrative privileges.  This setting is the default. This setting is appropriate for most cases. Do not change this setting unless you have a specific reason or scenario requirement to do so.

Certificate Store Setting	Certificate Store Override Setting	AnyConnect Search Strategy
All (for Windows)	checked	AnyConnect searches all certificate stores. AnyConnect is allowed to access the machine store when the user does not have administrative privileges.
Machine (for Windows)	true	AnyConnect searches in machine certificate store only. AnyConnect is allowed to access the machine store when the user does not have administrative privileges.
All (for macOS)	checked	AnyConnect uses certificates from all available macOS keychains and file stores.
User (for Windows)	does not apply	AnyConnect searches in the user certificate store only. The certificate store override is not applicable because users without administrative rights can have access to this certificate store.
System (for macOS)	checked	AnyConnect uses certificates only from the macOS system keychain and system file/PEM store. Uses certificates only from the macOS system keychain and system file/PEM store.
Login (for macOS)	checked	AnyConnect uses certificates only from the macOS login and dynamic smartcard keychains, as well as the user file/PEM store.
All (for Linux)	does not apply	AnyConnect uses client certificates from both system and user PEM file stores, as well as the user Firefox NSS store.
Machine (for Linux)	does not apply	AnyConnect uses client certificate stores only from the system PEM file store.
User (for Linux)	does not apply	AnyConnect uses client certificates only from the user PEM file store, as well as the user Firefox NSS store.

## With Multiple Certificate Authentication

### Before you begin

- Only supported on desktop platforms (Windows, macOS, and Linux).
- You must have *AutomaticCertSelection* enabled in the VPN profile.
- The certificate matching configuration you set in the VPN profile limits the certificates available for multiple certificate authentication.



---

**Note** SCEP is not supported.

---

### Procedure

---

- Step 1** Set **Certificate Store**:
- For one machine and one user certificate, set *CertificateStore* to **All** in the VPN profile and enable *CertificateStoreOverride* as described in Step 2.
  - For two user certificates, set *CertificateStore* to either **All** or **User** in the VPN profile but keep *CertificateStoreOverride* as described in Step 2.
- Step 2** Choose **Certificate Store Override** if you want to allow AnyConnect to search the machine certificate store when users do not have administrative privileges.
- 

## With Basic Certificate Authentication

### Procedure

---

- Step 1** Set **Certificate Store**.
- **All**—(Default) Directs the AnyConnect client to use all certificate stores for locating certificates.
  - **Machine**—Directs the AnyConnect client to restrict certificate lookup to the Windows local machine certificate store.
  - **User**—Directs the AnyConnect client to restrict certificate lookup to the local user certificate stores.
- Step 2** Choose **Certificate Store Override** if you want to allow AnyConnect to search the machine certificate store when users do not have administrative privileges.
- 

## Prompt Windows Users to Select Authentication Certificate

You can configure the AnyConnect to present a list of valid certificates to users and let them choose the certificate to authenticate the session. An expired certificate is not necessarily considered invalid. For example,

if you are using SCEP, the server might issue a new certificate to the client. Eliminating expired certificates might keep a client from connecting at all; thus requiring manual intervention and out-of-band certificate distribution. AnyConnect only restricts the client certificate based on security-related properties, such as key usage, key type and strength, and so on, based on configured certificate matching rules. This configuration is available only for Windows. By default, user certificate selection is disabled.

### Procedure

- 
- Step 1** Open the VPN Profile Editor and choose **Preferences (Part 2)** from the navigation pane.
- Step 2** To enable certificate selection, uncheck **Disable Certificate Selection**.
- Step 3** Uncheck **User Controllable**, unless you want users to be able to turn automatic certificate selection on and off in the **Advanced > VPN > Preferences** pane.
- 

## Create a PEM Certificate Store for macOS and Linux

AnyConnect supports certificate retrieval from a Privacy Enhanced Mail (PEM) formatted file store. AnyConnect reads PEM-formatted certificate files from the file system on the remote computer, verifies, and signs them.

### Before you begin

In order for the client to acquire the appropriate certificates under all circumstances, ensure that your files meet the following requirements:

- All certificate files must end with the extension `.pem`.
- All private key files must end with the extension `.key`.
- A client certificate and its corresponding private key must have the same filename. For example: `client.pem` and `client.key`.




---

**Tip** Instead of keeping copies of the PEM files, you can use soft links to PEM files.

---

To create the PEM file certificate store, create the paths and folders listed below. Place the appropriate certificates in these folders:

PEM File Certificate Store Folders	Type of Certificates Stored
<code>~/.cisco/certificates/ca</code> <b>Note</b> <code>.cisco/</code> is located in the home directory.	Trusted CA and root certificates
<code>~/.cisco/certificates/client</code>	Client certificates
<code>~/.cisco/certificates/client/private</code>	Private keys

Machine certificates are the same as PEM file certificates, except for the root directory. For machine certificates, substitute `/opt/.cisco` for `~/.cisco`. Otherwise, the paths, folders, and types of certificates listed apply.



## Configure Certificate Matching

AnyConnect can limit its search of certificates to those certificates that match a specific set of keys. Certificate matchings are global criteria that are set in an AnyConnect VPN client profile, in the **Certificate Matching** pane. The criteria are:

- Key Usage
- Extended Key Usage
- Distinguished Name

### Related Topics

[AnyConnect Profile Editor, Certificate Matching](#)

## Configure Key Usage

Selecting the **Key Usage** keys limits the certificates that AnyConnect can use to those certificates that have at least one of the selected keys. The supported set is listed in the **Key Usage** list on the VPN client profile, and it includes:

- DECIPHER\_ONLY
- ENCIPHER\_ONLY
- CRL\_SIGN
- KEY\_CERT\_SIGN
- KEY\_AGREEMENT
- DATA\_ENCIPHERMENT
- KEY\_ENCIPHERMENT
- NON\_REPUDIATION
- DIGITAL\_SIGNATURE

If one or more criteria are specified, a certificate must match at least one to be considered a matching certificate.

## Configure Extended Key Usage

Selecting the **Extended Key Usage** keys limits the certificates that AnyConnect can use to the certificates that have these keys. The following table lists the well-known set of constraints with their corresponding object identifiers (OIDs).

Constraint	OID
ServerAuth	1.3.6.1.5.5.7.3.1
ClientAuth	1.3.6.1.5.5.7.3.2
CodeSign	1.3.6.1.5.5.7.3.3
EmailProtect	1.3.6.1.5.5.7.3.4
IPSecEndSystem	1.3.6.1.5.5.7.3.5

Constraint	OID
IPSecTunnel	1.3.6.1.5.5.7.3.6
IPSecUser	1.3.6.1.5.5.7.3.7
TimeStamp	1.3.6.1.5.5.7.3.8
OCSPSign	1.3.6.1.5.5.7.3.9
DVCS	1.3.6.1.5.5.7.3.10
IKE Intermediate	1.3.6.1.5.5.8.2.2

### Configure Custom Extended Match Key

All other OIDs (such as 1.3.6.1.5.5.7.3.11, used in some examples in this document) are considered “custom.” As an administrator, you can add your own OIDs if the OID that you want is not in the well-known set.

### Configure Certificate Distinguished Name

The **Distinguished Name** table contains certificate identifiers that limit the certificates that the client can use to the certificates that match the specified criteria and criteria match conditions. Click the **Add** button to add criteria to the list and to set a value or wildcard to match the contents of the added criteria.

Identifier	Description
CN	SubjectCommonName
SN	SubjectSurName
GN	SubjectGivenName
N	SubjectUnstructName
I	SubjectInitials
GENQ	SubjectGenQualifier
DNQ	SubjectDnQualifier
C	SubjectCountry
L	SubjectCity
SP	SubjectState
ST	SubjectState
O	SubjectCompany
OU	SubjectDept
T	SubjectTitle
EA	SubjectEmailAddr

Identifier	Description
DC	DomainComponent
ISSUER-CN	IssuerCommonName
ISSUER-SN	IssuerSurName
ISSUER-GN	IssuerGivenName
ISSUER-N	IssuerUnstructName
ISSUER-I	IssuerInitials
ISSUER-GENQ	IssuerGenQualifier
ISSUER-DNQ	IssuerDnQualifier
ISSUER-C	IssuerCountry
ISSUER-L	IssuerCity
ISSUER-SP	IssuerState
ISSUER-ST	IssuerState
ISSUER-O	IssuerCompany
ISSUER-OU	IssuerDept
ISSUER-T	IssuerTitle
ISSUER-EA	IssuerEmailAddr
ISSUER-DC	IssuerDomainComponent

**Distinguished Name** can contain zero or more matching criteria. A certificate must match all specified criteria to be considered a matching certificate. **Distinguished Name** matching specifies that a certificate must or must not have the specified string, and whether wild carding for the string is allowed.

## VPN Authentication Using SAML

You can use SAML 2.0 integrated with ASA release 9.7.1 for initial session authentication. To provide a seamless reconnect without disruption, AnyConnect intentionally skips the repeating of the SAML authentication process. Additionally, if the user logs out of the IdP using a browser, the AnyConnect session remains intact.

Follow these guidelines when using SAML:

- If you are using always-on VPN in failover mode, external SAML IdP is not supported (however, with internal SAML IdP, the ASA proxies all traffic to IdP and is supported)
- (Mobile only) Single logout is not supported.
- You must synchronize your ASA's Network Time Protocol (NTP) server with the IdP NTP server in order to use the SAML feature.

- The VPN Wizard on ASDM does not currently support SAML configurations.
- The SAML IdP *NameID* attribute determines the user's username and is used for authorization, accounting, and VPN session database.
- You should set Auto Reconnect to *ReconnectAfterResume* in the [AnyConnect Profile Editor, Preferences \(Part 1\)](#) if you want users to re-authenticate with the Identity Provider (IdP) every time they establish a VPN session via SAML.

For SAML to function as expected, you must upgrade to the latest webdeploy versions (or later):

- anyconnect-macos-4.4.01054-webdeploy-k9.pkg
- anyconnect-win-4.4.01054-webdeploy-k9.pkg
- anyconnect-linux64-4.4.01054-webdeploy-k9.pkg

For AnyConnect Mobile the following platforms and versions are supported:

- Chrome OS 4.0.10151

Refer to the *SSO Using SAML 2.0* section in the appropriate release, 9.7 or later, of the [Cisco ASA Series VPN Configuration Guide](#) for additional configuration details.

## VPN Authentication Using SDI Token (SoftID) Integration

AnyConnect integrates support for RSA SecurID client software versions 1.1 and later running on Windows 7 x86 (32-bit) and x64 (64-bit).

RSA SecurID software authenticators reduce the number of items a user has to manage for safe and secure access to corporate assets. RSA SecurID Software Tokens residing on a remote device generate a random one-time-use passcode that changes every 60 seconds. The term SDI stands for Security Dynamics, Inc. technology, which refers to this one-time password generation technology that uses hardware and software tokens.

Typically, users make an AnyConnect connection by clicking the AnyConnect icon in the tools tray, selecting the connection profile with which they wish to connect, and then entering the appropriate credentials in the authentication dialog box. The login (challenge) dialog box matches the type of authentication configured for the tunnel group to which the user belongs. The input fields of the login dialog box clearly indicate what kind of input is required for authentication.

For SDI authentication, the remote user enters a PIN (Personal Identification Number) into the AnyConnect software interface and receives an RSA SecurID passcode. After the user enters the passcode into the secured application, the RSA Authentication Manager validates the passcode and allows the user to gain access.

Users who use RSA SecurID hardware or software tokens see input fields indicating whether the user should enter a passcode or a PIN, a PIN, or a passcode and the status line at the bottom of the dialog box provides further information about the requirements. The user enters a software token PIN or passcode directly into the AnyConnect user interface.

The appearance of the initial login dialog box depends on the secure gateway settings: the user can access the secure gateway either through the main login page, the main index URL, a tunnel-group login page, or a tunnel group URL (URL/tunnel-group). To access the secure gateway via the main login page, the “Allow user to select connection” check box must be set in the Network (Client) Access AnyConnect Connection Profiles page. In either case, the secure gateway sends the client a login page. The main login page contains a drop-down

list in which the user selects a tunnel group; the tunnel-group login page does not, since the tunnel-group is specified in the URL.

In the case of a main login page (with a drop-down list of connection profiles or tunnel groups), the authentication type of the default tunnel group determines the initial setting for the password input field label. For example, if the default tunnel group uses SDI authentication, the field label is “Passcode;” but if the default tunnel group uses NTLM authentication, the field label is “Password.” In Release 2.1 and later, the field label is not dynamically updated with the user selection of a different tunnel group. For a tunnel-group login page, the field label matches the tunnel-group requirements.

The client supports input of RSA SecurID Software Token PINs in the password input field. If the RSA SecurID Software Token software is installed and the tunnel-group authentication type is SDI, the field label is “Passcode” and the status bar states “Enter a username and passcode or software token PIN.” If a PIN is used, subsequent consecutive logins for the same tunnel group and username have the field label “PIN.” The client retrieves the passcode from the RSA SecurID Software Token DLL using the entered PIN. With each successful authentication, the client saves the tunnel group, the username, and authentication type, and the saved tunnel group becomes the new default tunnel group.

AnyConnect accepts passcodes for any SDI authentication. Even when the password input label is “PIN,” the user may still enter a passcode as instructed by the status bar. The client sends the passcode to the secure gateway as is. If a passcode is used, subsequent consecutive logins for the same tunnel group and username have the field label “Passcode.”

The RSA SecureID Integration profile setting has three possible values:

- **Automatic**—The client first attempts one method, and if it fails, the other method is tried. The default is to treat the user input as a token passcode (HardwareToken), and if that fails, treat it as a software token pin (SoftwareToken). When authentication is successful, the successful method is set as the new SDI Token Type and cached in the user preferences file. For the next authentication attempt, the SDI Token Type defines which method is attempted first. Generally, the token used for the current authentication attempt is the same token used in the last successful authentication attempt. However, when the username or group selection is changed, it reverts to attempting the default method first, as shown in the input field label.



---

**Note** The SDI Token Type only has meaning for the automatic setting. You can ignore logs of the SKI Token Type when the authentication mode is not automatic. HardwareToken as the default avoids triggering next token mode.

---

- **SoftwareToken**—The client always interprets the user input as a software token PIN, and the input field label is “PIN:”.
- **HardwareToken**—The client always interprets the user input as a token passcode, and the input field label is “Passcode:”.



---

**Note** AnyConnect does not support token selection from multiple tokens imported into the RSA Software Token client software. Instead, the client uses the default selected via the RSA SecurID Software Token GUI.

---

## Categories of SDI Authentication Exchanges

All SDI authentication exchanges fall into one of the following categories:

- Normal SDI Authentication Login
- New User mode
- New PIN mode
- Clear PIN mode
- Next Token Code mode

### Normal SDI Authentication Login

A normal login challenge is always the first challenge. The SDI authentication user must provide a user name and token passcode (or PIN, in the case of a software token) in the username and passcode or PIN fields, respectively. The client returns the information to the secure gateway (central-site device), and the secure gateway verifies the authentication with the authentication server (SDI or SDI via RADIUS proxy).

If the authentication server accepts the authentication request, the secure gateway sends a success page back to the client, and the authentication exchange is complete.

If the passcode is not accepted, the authentication fails, and the secure gateway sends a new login challenge page, along with an error message. If the passcode failure threshold on the SDI server has been reached, then the SDI server places the token into next token code mode.

### New User, Clear PIN, and New PIN Modes

The PIN can be cleared only on the SDI server and only by the network administrator.

In the New User, Clear PIN, and New PIN modes, AnyConnect caches the user-created PIN or system-assigned PIN for later use in the “next passcode” login challenge.

Clear PIN mode and New User mode are identical from the point of view of the remote user and are both treated the same by the secure gateway. In both cases, the remote user either must enter a new PIN or be assigned a new PIN by the SDI server. The only difference is in the user response to the initial challenge.

For New PIN mode, the existing PIN is used to generate the passcode, as it would be in any normal challenge. For Clear PIN mode, no PIN is used at all for hardware tokens, with the user entering just a token code. A PIN of eight consecutive zeros (00000000) is used to generate a passcode for RSA software tokens. In either case, the SDI server administrator must inform the user of what, if any, PIN value to use.

Adding a new user to an SDI server has the same result as clearing the PIN of an existing user. In both cases, the user must either provide a new PIN or be assigned a new PIN by the SDI server. In these modes, for hardware tokens, the user enters just a token code from the RSA device. In either case, the SDI server administrator must inform the user of what, if any, PIN value to use.

### Creating a New PIN

If there is no current PIN, the SDI server requires that one of the following conditions be met, depending on how the system is configured:

- The system must assign a new PIN to the user (Default)
- The user must create a new PIN

- The user can choose whether to create a PIN or have the system assign it

If the SDI server is configured to allow the remote user to choose whether to create a PIN or have the system assign a PIN, the login screen presents a drop-down list showing the options. The status line provides a prompt message.

For a system-assigned PIN, if the SDI server accepts the passcode that the user enters on the login page, then the secure gateway sends the client the system-assigned PIN. The client sends a response back to the secure gateway, indicating that the user has seen the new PIN, and the system continues with a “next passcode” challenge.

If the user chooses to create a new PIN, AnyConnect presents a dialog box on which to enter that PIN. The PIN must be a number from 4 to 8 digits long. Because the PIN is a type of password, anything the user enters into these input fields is displayed as asterisks.

With RADIUS proxy, the PIN confirmation is a separate challenge, subsequent to the original dialog box. The client sends the new PIN to the secure gateway, and the secure gateway continues with a “next passcode” challenge.

### **“Next Passcode” and “Next Token Code” Challenges**

For a “next passcode” challenge, the client uses the PIN value cached during the creation or assignment of a new PIN to retrieve the next passcode from the RSA SecurID Software Token DLL and return it to the secure gateway without prompting the user. Similarly, in the case of a “next Token Code” challenge for a software token, the client retrieves the next Token Code from the RSA SecurID Software Token DLL.

## **Compare Native SDI with RADIUS SDI**

The network administrator can configure the secure gateway to allow SDI authentication in either of the following modes:

- Native SDI refers to the native ability in the secure gateway to communicate directly with the SDI server for handling SDI authentication.
- RADIUS SDI refers to the process of the secure gateway performing SDI authentication using a RADIUS SDI proxy, which communicates with the SDI server.

Native SDI and RADIUS SDI appear identical to the remote user. Because the SDI messages are configurable on the SDI server, the message text on the ASA must match the message text on the SDI server. Otherwise, the prompts displayed to the remote client user might not be appropriate for the action required during authentication. AnyConnect might fail to respond and authentication might fail.

RADIUS SDI challenges, with minor exceptions, essentially mirror native SDI exchanges. Since both ultimately communicate with the SDI server, the information needed from the client and the order in which that information is requested is the same.

During authentication, the RADIUS server presents access challenge messages to the ASA. Within these challenge messages are reply messages containing text from the SDI server. The message text is different when the ASA is communicating directly with an SDI server from when communicating through the RADIUS proxy. Therefore, in order to appear as a native SDI server to AnyConnect, the ASA must interpret the messages from the RADIUS server.

Also, because the SDI messages are configurable on the SDI server, the message text on the ASA must match (in whole or in part) the message text on the SDI server. Otherwise, the prompts displayed to the remote client user may not be appropriate for the action required during authentication. AnyConnect might fail to respond and authentication might fail.

## Configure the ASA to Support RADIUS/SDI Messages

To configure the ASA to interpret SDI-specific RADIUS reply messages and prompt the AnyConnect user for the appropriate action, you must configure a connection profile (tunnel group) to forward RADIUS reply messages in a manner that simulates direct communication with an SDI server. Users authenticating to the SDI server must connect over this connection profile.

### Procedure

- 
- Step 1** Go to **Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles**.
  - Step 2** Select the connection profile you want to configure to interpret SDI-specific RADIUS reply messages and click **Edit**.
  - Step 3** In the **Edit AnyConnect Connection Profile** window, expand the Advanced node in the navigation pane on the left and select **Group Alias / Group URL**.
  - Step 4** Check **Enable the display of SecurID messages on the login screen**.
  - Step 5** Click **OK**.
  - Step 6** Choose **Configuration > Remote Access VPN > AAA/Local Users > AAA Server Groups**.
  - Step 7** Click **Add** to Add a AAA Server group.
  - Step 8** Configure the AAA server group in the Edit AAA Server Group dialog and click **OK**.
  - Step 9** In the **AAA Server Groups** area, select the AAA server group you just created and then click **Add** in the **Servers in the Selected Group** area.
  - Step 10** In the SDI Messages area, expand the **Message Table** area. Double-click a message text field to edit the message. Configure the RADIUS reply message text on the ASA to match (in whole or in part) the message text sent by the RADIUS server.

The following table shows the message code, the default RADIUS reply message text, and the function of each message:

**Note** The default message text used by the ASA is the default message text used by Cisco Secure Access Control Server (ACS). If you are using Cisco Secure ACS, and it is using the default message text, you do not need to configure the message text on the ASA.

Because the security appliance searches for strings in the order in which they appear in the table, you must ensure that the string you use for the message text is not a subset of another string. For example, “new PIN” is a subset of the default message text for both new-pin-sup and next-ccode-and-reauth. If you configure new-pin-sup as “new PIN,” when the security appliance receives “new PIN with the next card code” from the RADIUS server, it will match the text to the new-pin-sup code instead of the next-ccode-and-reauth code.

Message Code	Default RADIUS Reply Message Text	Function
next-code	Enter Next PASSCODE	Indicates the user must enter the NEXT tokencode without the PIN.
new-pin-sup	Please remember your new PIN	Indicates the new system PIN has been supplied and displays that PIN for the user.



Message Code	Default RADIUS Reply Message Text	Function
new-pin-meth	Do you want to enter your own pin	Requests from the user which new PIN method to use to create a new PIN.
new-pin-req	Enter your new Alpha-Numerical PIN	Indicates a user-generated PIN and requests that the user enter the PIN.
new-pin-reenter	Reenter PIN:	Used internally by the ASA for user-supplied PIN confirmation. The client confirms the PIN without prompting the user.
new-pin-sys-ok	New PIN Accepted	Indicates the user-supplied PIN was accepted.
next-ccode-and-reauth	new PIN with the next card code	Follows a PIN operation and indicates the user must wait for the next tokencode and to enter both the new PIN and next tokencode to authenticate.
ready-for-sys- pin	ACCEPT A SYSTEM GENERATED PIN	Used internally by the ASA to indicate the user is ready for the system-generated PIN.

**Step 11** Click **OK**, then **Apply**, then **Save**.

## About Certificate Pinning

AnyConnect certificate pinning helps to detect if a server certificate chain actually came from the connecting server. This feature is guided by VPN profile settings and is an addition to the AnyConnect server certificate verification policies. The strict certificate trust settings in the AnyConnect local policy file have no influence on Certificate Pinning check. You can configure pins globally or by per host basis in the VPN profile. Those pins configured for primary host are also valid for back up hosts in the server list. The preference to perform certificate pinning checks is not user controllable. A pin verification failure results in the termination of the VPN connection.



**Note** AnyConnect performs pin verification only when the preference is enabled and the connecting server has pins in the VPN profile.

In the VPN profile editor [AnyConnect Profile Editor, Certificate Pin](#), you can enable the preference and configure the global and per host certificate pins.

You must be cautious when configuring and maintaining certificate pinning. Consider these recommendations when setting preferences:

- Pin root and/or intermediate certificates since they are well maintained by CA vendors in the operating system
- Pin multiple root and/or intermediate certificates from a different CA to serve as a backup when any CA is compromised
- Pin multiple root and/or intermediate certificates for ease of CA transitions
- Use the same Certificate Signing Request if a leaf certificate is pinned, to retain the public key upon certificate renewal
- Pin all connection hosts in the server list

## Global and Per Host Pins

You can configure certificate pins on a global or by per host basis. Pins which are valid for most of the connection hosts are configured as global pins. We recommend that you configure root, intermediate certificate authorities, and wild card leaf certificates under global pins in the VPN profile. Pins that are valid only for a connection host are considered as per host pins. We recommend that you configure leaf, self-signed certificates under per host pins in the VPN profile.




---

**Note** AnyConnect checks global pins and per host pins for the corresponding connection server during pin verification.

---




---

**Note** Global pins across multiple VPN profiles are not merged. Pins are strictly considered from the file connection server for VPN connection.

---




---

**Note** You can only pin per host certificates when certificate pinning preference is enabled in the global pins section.

---