Remote Access VPN

Remote Access virtual private network (VPN) allows individual users to connect to your network from a remote location using a computer or other supported iOS or Android device connected to the Internet. This allows mobile workers to connect from their home networks or a public Wi-Fi network, for example.

The following topics explain how to configure remote access VPN for your network.

Remote Access VPN Overview

You can use the device manager to configure remote access VPN over SSL using the AnyConnect Client sofware.

When the AnyConnect Client negotiates an SSL VPN connection with the threat defense device, it connects using Transport Layer Security (TLS) or Datagram Transport Layer Security (DTLS). DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays. The client and the threat defense device negotiate the TLS/DTLS version to use. DTLS is used if the client supports it.

Maximum Concurrent VPN Sessions By Device Model

There is a maximum limit to the number of concurrent remote access VPN sessions allowed on a device based on the device model. This limit is designed so that system performance does not degrade to unacceptable levels. Use these limits for capacity planning.

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firepower 1010	75
Firepower 1120	150

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firepower 1140	400
Firepower 2110	1500
Firepower 2120	3500
Firepower 2130	7500
Firepower 2140	10,000
Secure Firewall 3110	3000
Secure Firewall 3120	6000
Secure Firewall 3130	15,000
Secure Firewall 3140	20,000
Firepower 4100 series, all models	10,000
Firepower 9300 appliance, all models	20,000
Threat Defense Virtual: FTDv5	50
Threat Defense Virtual: FTDv10, FTDv20, FTDv30	250
Threat Defense Virtual: FTDv50	750
Threat Defense Virtual: FTDv100	10,000
ISA 3000	25

Downloading the AnyConnect Client Software

Before you can configure a remote access VPN, you must download the AnyConnect Client software to your workstation. You will need to upload these packages when defining the VPN.

You should download the latest AnyConnect Client version, to ensure that you have the latest features, bug fixes, and security patches. Regularly update the packages on the threat defense device.

Note You can upload one AnyConnect Client package per operating system: Windows, Mac, and Linux. You cannot upload multiple versions for a given OS type.

Obtain the AnyConnect Client software packages from software.cisco.com. You need to download the "Full Installation Package" versions of the clients.

How Users Can Install the AnyConnect Client Software

To complete a VPN connection, your users must install the AnyConnect Client software. You can use your existing software distribution methods to install the software directly. Or, you can have users install the AnyConnect Client directly from the threat defense device.

Users must have Administrator rights on their workstations to install the software.

Once the AnyConnect Client is installed, if you upload new AnyConnect Client versions to the system, the AnyConnect Client will detect the new version on the next VPN connection the user makes. The system will automatically prompt the user to download and install the updated client software. This automation simplifies software distribution for you and your clients.

If you decide to have users initially install the software from the threat defense device, tell users to perform the following steps.

Note Android and iOS users should download the AnyConnect Client from the appropriate App Store.

Procedure

Step 1 Using a web browser, open **https://**ravpn-address, where ravpn-address is the IP address or hostname of the outside interface on which you are allowing VPN connections.

You identify this interface when you configure the remote access VPN. The system prompts the user to log in.

If you changed the port for remote access VPN connections, users must include the custom port in the URL. For example, if you changed the port to 4443: https://ravpn.example.com:4443

Step 2 Log into the site.

Users are authenticated using the directory server configured for the remote access VPN. Log in must be successful to continue.

If log in is successful, the system determines if the user already has the required version of the AnyConnect Client. If the AnyConnect Client is absent from the user's computer, or is down-level, the system automatically starts installing the AnyConnect Client software.

When installation is finished, AnyConnect Client completes the remote access VPN connection.

Controlling User Permissions and Attributes Using RADIUS and Group Policies

You can apply user authorization attributes (also called user entitlements or permissions) to RA VPN connections from an external RADIUS server or from a group policy defined on the threat defense device. If the threat defense device receives attributes from the external AAA server that conflict with those configured on the group policy, then attributes from the AAA server always take precedence.

The threat defense device applies attributes in the following order:

1. User attributes defined on the external AAA server—The server returns these attributes after successful user authentication or authorization.

2. Group policy configured on the threat defense device—If a RADIUS server returns the value of the RADIUS CLASS attribute IETF-Class-25 (OU= group-policy) for the user, the threat defense device places the user in the group policy of the same name and enforces any attributes in the group policy that are not returned by the server.

3. Group policy assigned by the connection profile—The connection profile has the preliminary settings for the connection, and includes a default group policy applied to the user before authentication. All users connecting to the threat defense device initially belong to this group, which provides any attributes that are missing from the user attributes returned by the AAA server, or the group policy assigned to the user.

Threat Defense devices support RADIUS attributes with vendor ID 3076. If the RADIUS server you use does not have these attributes defined, you must manually define them. To define an attribute, use the attribute name or number, type, value, and vendor code (3076).

The following topics explain the supported attributes based on whether the values are defined in the RADIUS server, or whether they are values the system sends to the RADIUS server.

Attributes Sent to the RADIUS Server

RADIUS attributes 146 and 150 are sent from the threat defense device to the RADIUS server for authentication and authorization requests. All of the following attributes are sent from the threat defense device to the RADIUS server for accounting start, interim-update, and stop requests.

Table 1: Attributes Threat Defense Sends to RADIUS

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Client Type	150	Integer	Single	The type of client that is connecting to the VPN: 2 = AnyConnect Client SSL VPN
Session Type	151	Integer	Single	The type of connection: 1 = AnyConnect Client SSL VPN
Tunnel Group Name	146	String	Single	The name of the connection profile that was used to establish the session, as defined on the threat defense device. The name can be 1 - 253 characters.

Attributes Received from the RADIUS Server

The following user authorization attributes are sent to the threat defense device from the RADIUS server.

Table 2: RADIUS Attributes Sent to Threat Defense

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Access-List-Inbound	86	String	Single	Both of the Acess-List attributes take the name of an ACL that is configured on the threat defense device. Create these ACLs using the Smart CLI Extended Access List object type (select **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**).
Access-List-Outbound	87	String	Single	
				These ACLs control traffic flow in the inbound (traffic entering the threat defense device) or outbound (traffic leaving the threat defense device) direction.
Address-Pools	217	String	Single	The name of a network object defined on the threat defense device that identifies a subnet, which will be used as the address pool for clients connecting to the RA VPN. Define the network object on the **Objects** page.
Banner1	15	String	Single	The banner to display when the user logs in.
Banner2	36	String	Single	The second part of the banner to display when the user logs in. Banner2 is appended to Banner1.
Group-Policy	25	String	Single	The group policy to use in the connection. You must create the group policy on the RA VPN **Group Policy** page. You can use one of the following formats: • *group policy name* • OU=*group policy name* • OU=*group policy name*;
Simultaneous-Logins	2	Integer	Single	The number of separate simultaneous connections the user is allowed to establish, 0 - 2147483647.
VLAN	140	Integer	Single	The VLAN on which to confine the user's connection, 0 - 4094. You must also configure this VLAN on a subinterface on the threat defense device.

Two-Factor Authentication

You can configure two-factor authentication for the RA VPN. With two-factor authentication, the user must supply a username and static password, plus an additional item such as an RSA token or a Duo passcode. Two-factor authentication differs from using a second authentication source in that two-factor is configured on a single authentication source, with the relationship to the RSA/Duo server tied to the primary authentication source. The exception is Duo LDAP, where you configure the Duo LDAP server as the secondary authentication source.

The system has been tested with RSA tokens and Duo passcode pushed to mobile for the second factor in conjunction with any RADIUS or AD Server as the first factor in the two-factor authentication process.

RSA Two-Factor Authentication

You can configure RSA using one of the following approaches. See the RSA documentation for information about the RSA-side configuration.

- Define the RSA Server directly in the device manager as a RADIUS server, and use the server as the primary authentication source in the RA VPN.

  When using this approach, the user must authenticate using a username that is configured in the RSA RADIUS server, and concatenate the password with the one-time temporary RSA token, separating the password and token with a comma: *password,token*.

  In this configuration, it is typical to use a separate RADIUS server (such as one supplied in Cisco ISE) to provide authorization services. You would configure the second RADIUS server as the authorization and, optionally, accounting server.

- Integrate the RSA server with a RADIUS or AD server that supports direct integration, and configure the RA VPN to use the non-RSA RADIUS or AD server as the primary authentication source. In this case, the RADIUS/AD server uses RSA-SDI to delegate and orchestrate the two-factor authentication between the client and RSA Server.

  When using this approach, the user must authenticate using a username that is configured in the non-RSA RADIUS or AD server, and concatenate the password with the one-time temporary RSA token, separating the password and token with a comma: *password,token*.

  In this configuration, you would also use the non-RSA RADIUS server as the authorization and, optionally, accounting server.

Duo Two-Factor Authentication Using RADIUS

You can configure the Duo RADIUS server as the primary authentication source. This approach uses the Duo RADIUS Authentication Proxy.

For the detailed steps to configure Duo, please see https://duo.com/docs/cisco-firepower.

You would then configure Duo to forward authentication requests directed to the proxy server to use another RADIUS server, or an AD server, as the first authentication factor, and the Duo Cloud Service as the second factor.

When using this approach, the user must authenticate using a username that is configured on both the Duo Authentication Proxy and the associated RADIUS/AD server, and the password for the username configured in the RADIUS/AD server, followed by one of the following Duo codes:

- *Duo-passcode*. For example, *my-password*,**12345**.

- **push**. For example, *my-password*,**push**. Use **push** to tell Duo to send a push authentication to the Duo Mobile app, which the user must have already installed and registered.

- **sms**. For example, *my-password*,**sms**. Use **sms** to tell Duo to send an SMS message with a new batch of passcodes to the user's mobile device. The user's authentication attempt will fail when using **sms**. The user must then re-authenticate and enter the new passcode as the secondary factor.

- **phone**. For example, *my-password*,**phone**. Use **phone** to tell Duo to perform phone callback authentication.

If the username/password is authenticated, the Duo Authentication Proxy contacts the Duo Cloud Service, which validates that the request is from a valid configured proxy device and then pushes a temporary passcode to the mobile device of the user as directed. When the user accepts this passcode, the session is marked authenticated by Duo and the RA VPN is established.

Duo Two-Factor Authentication using LDAP

You can use the Duo LDAP server as the secondary authentication source in conjunction with a Microsoft Active Directory (AD) or RADIUS server as the primary source. With Duo LDAP, the secondary authentication validates the primary authentication with a Duo passcode, push notification, or phone call.

The threat defense device communicates with Duo LDAP using LDAPS over port TCP/636.

Note that the Duo LDAP server provides authentication services only, it does not provide identity services. Thus, if you use Duo LDAP as a primary authentication source, you will not see usernames associated with RA VPN connections in any dashboards, and you will not be able to write access control rules for these users.

When using this approach, the user must authenticate using a username that is configured on both the RADIUS/AD server and the Duo LDAP server. When prompted to log in by the AnyConnect Client, the user provides the RADIUS/AD password in the primary **Password** field, and for the **Secondary Password**, provides one of the following to authenticate with Duo. For more details, see https://guide.duo.com/anyconnect.

- *Duo passcode*—Authenticate using a passcode, either generated with Duo Mobile, sent via SMS, generated by your hardware token, or provided by an administrator. For example, 1234567.

- **push**—Push a login request to your phone, if you have installed and activated the Duo Mobile app. Review the request and tap **Approve** to log in.

- **phone**—Authenticate using a phone callback.

- **sms**—Request a Duo passcode in a text message. The login attempt will fail. Log in again using the new passcode.

For a detailed explanation and example of using Duo LDAP, see How to Configure Two-Factor Authentication using Duo LDAP, on page 39.

Licensing Requirements for Remote Access VPN

Your base device license must meet export requirements before you can configure remote access VPN. When you register the device, you must do so with a Smart Software Manager account that is enabled for export-controlled features. You also cannot configure the feature using the evaluation license.

In addition, you need to purchase and enable a remote access VPN license, any of the following: AnyConnect Plus, AnyConnect Apex, or AnyConnect VPN Only. These licenses are treated the same for threat defense devices, even though they are designed to allow different feature sets when used with ASA Software-based headends.

To enable the license, select **Device** > **Smart License** > **View Configuration**, then select the appropriate license in the RA VPN License group. You need to have the license available in your Smart Software Manager account. For more information about enabling licenses, see Enabling or Disabling Optional Licenses.

For more information, see the *Cisco AnyConnect Ordering Guide*, http://www.cisco.com/c/dam/en/us/products/collateral/security/anyconnect-og.pdf. There are also other data sheets available on http://www.cisco.com/c/en/us/products/security/anyconnect-secure-mobility-client/datasheet-listing.html.

Guidelines and Limitations for Remote Access VPN

Please keep the following guidelines and limitations in mind when configuring RA VPN.

- You cannot configure both the device manager access (HTTPS access in the management access list) and remote access SSL VPN on the same interface for the same TCP port. For example, if you configure remote access SSL VPN on the outside interface, you cannot also open the outside interface for HTTPS connections on port 443. If you configure both features on the same interface, ensure that you change the HTTPS port for at least one of these services to avoid a conflict.

- The RA VPN outside interface is a global setting. You cannot configure separate connection profiles on different interfaces.

- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.

- If you configure two-factor authentication using RADIUS and RSA tokens, the default authentication timeout of 12 seconds is too quick to allow successful authentication in most cases. You can increase the authentication timeout value by creating a custom AnyConnect Client profile and applying it to the RA VPN connection profile, as described in Configure and Upload Client Profiles, on page 10. We recommend an authentication timeout of at least 60 seconds, so that users have enough time to authenticate and then paste the RSA token, and for the round-trip verification of the token.

- Issuing commands such as **curl** against the RA VPN headend is not directly supported, and might not have desirable results. For example, the headend does not respond to HTTP HEAD requests.

Configuring Remote Access VPN

To enable remote access VPN for your clients, you need to configure a number of separate items. The following procedure provides the end to end process.

Procedure

Step 1   Configure licenses.

You need to enable two licenses:

- When you register the device, you must do so with a Smart Software Manager account that is enabled for export-controlled features. The base license must meet export control requirements before you can configure remote access VPN. You also cannot configure the feature using the evaluation license. For the procedure to register the device, see Registering the Device.

- A remote access VPN license. For details, see Licensing Requirements for Remote Access VPN, on page 7. To enable the license, see Enabling or Disabling Optional Licenses.

Step 2   Configure Certificates.

Certificates are required to authenticate SSL connections between the clients and the device. You can use the pre-defined DefaultInternalCertificate for the VPN, or create your own.

If you use an encrypted connection for the directory realm used for authentication, you must upload a trusted CA certificate.

For more information on certificates and how to upload them, see Configuring Certificates.

Step 3   (Optional.) Configure TLS/SSL settings.

By default, the system will allow remote users to connect to the remote access VPN using any TLS version and encryption cipher supported by the system. However, you can limit the TLS/DTLS versions, ciphers, and Diffie-Hellman groups allowed to enforce a more secure connection. See Configuring TLS/SSL Cipher Settings.

Step 4   (Optional.) Configure and Upload Client Profiles, on page 10.

Step 5   Configure the identity source used for authenticating remote users.

You can use the following sources for user accounts that are allowed to log into the remote access VPN. Alternatively, you can use client certificates for authentication, either alone or in conjunction with an identity source.

- Active Directory identity realm—As a primary authentication source. The user accounts are defined in your Active Directory (AD) server. See Configuring AD Identity Realms.

- RADIUS server group—As a primary or secondary authentication source, and for authorization and accounting. See Configure RADIUS Server Groups.

- LocalIdentitySource (the local user database)—As a primary or fallback source. You can define users directly on the device and not use an external server. If you use the local database as a fallback source, ensure that you define the same usernames/passwords as the ones defined in the external server. See Configure Local Users.

- Duo LDAP server—As a primary or secondary authentication source. Although you can use a Duo LDAP server as the primary source, this is not the normal configuration. You would normally use it as the secondary source to provide two-factor authentication in conjunction with a primary Active Directory or RADIUS server. For details, see How to Configure Two-Factor Authentication using Duo LDAP, on page 39.

Step 6   (Optional.) Configure Group Policies for RA VPN, on page 22

The group policy defines user-related attributes. You can configure group policies to provide differential access to resources based on group membership. Alternatively, you can use the default policy for all connections.

Step 7   Configure an RA VPN Connection Profile, on page 15.

Step 8   Allow Traffic Through the Remote Access VPN, on page 12.

Step 9   Verify the Remote Access VPN Configuration, on page 13.

If you encounter problems completing a connection, see Troubleshooting Remote Access VPNs, on page 28.

Step 10   (Optional.) Enable the identity policy and configure a rule for passive authentication.

If you enable passive user authentication, users who logged in through the remote access VPN will be shown in the dashboards, and they will also be available as traffic-matching criteria in policies. If you do not enable passive authentication, RA VPN users will be available only if they match an active authentication policy. You must enable the identity policy to get any username information in the dashboards or for traffic matching.

Configure and Upload Client Profiles

AnyConnect Client profiles are downloaded to clients along with the AnyConnect Client software. These profiles define many client-related options, such as auto connect on startup and auto reconnect, and whether the end user is allowed to change the option from the AnyConnect Client preferences and advanced settings.

If you configure a fully-qualified hostname (FQDN) for the outside interface when configuring the remote access VPN connection, the system creates a client profile for you. This profile enables the default settings. You need to create and upload client profiles only if you want non-default behavior. Note that client profiles are optional: if you do not upload one, AnyConnect Client will use default settings for all profile-controlled options.

Note You must include the threat defense device's outside interface in the VPN profile's server list in order for the AnyConnect Client to display all user controllable settings on the first connection. If you do not add the address or FQDN as a host entry in the profile, then filters do not apply for the session. For example, if you create a certificate match and the certificate properly matches the criteria, but you do not add the device as a host entry in that profile, the certificate match is ignored.

You can create profiles for the AnyConnect Client as well as for a variety of modules that you can optionally use with AnyConnect Client, such as the AMP Enabler. Although you can upload profiles for any of these modules, the device manager supports the creation of the AnyConnect Client Profile only. However, you can upload any kind of profile through the device manager, then use the threat defense API (from API Explorer) to change the profile type for the object. The profiles page shows all profiles of any type, although the list does not indicate the profile type. The procedure explains how to do this.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create AnyConnect Client profile objects while editing a profile property by clicking the **Create New Secure Client Profile** link shown in the object list.

Before you begin

Before you can upload client profiles, you must do the following.

- Download and install the stand-alone AnyConnect Client "Profile Editor - Windows / Standalone installer (MSI)." The installation file is for Windows only, and has the file name anyconnect-profileeditor-win-<version>-k9.msi, where <version> is the AnyConnect Client version (the file name is subject to change). For example, anyconnect-profileeditor-win-4.3.04027-k9.msi. You must also install Java JRE 1.6 (or higher) before installing the profile editor. Obtain the AnyConnect Client profile editor from software.cisco.com. Note that this package contains all of the profile editors, not just the one for the VPN client.

- Use the profile editor to create the profiles you need. You should specify the hostname or IP address of the outside interface in the profile. For detailed information, see the editor's online help.

Procedure

Step 1 Select **Objects**, then select **Secure Client Profiles** from the table of contents.

Step 2 Do one of the following:

- To create an object, click the + button.

• To edit an object, click the edit icon (🖊) for the object.

• To download the profile associated with an object, click the download icon (⬇) for the object.

To delete an unreferenced object, click the trash can icon (🗑) for the object.

Step 3   Enter a name and optionally, a description, for the object.

If you are uploading a module profile, use the object name to indicated the module type to make it easier for you to distinguish it from the AnyConnect Client Profiles.

Step 4   Click **Upload** and select the file you created using the Profile Editor.

Step 5   Click **Open** to upload the profile.

Step 6   Click **OK** to add the object.

Step 7   If the profile you created is actually a different type than AnyConnect Client Profile, complete the following steps to alter the profile type of the object.

a) Click the more options button (⋮) and choose **API Explorer**.

The system opens the API Explorer in a separate tab or window, depending on your browser settings.

b) Open the AnyConnectClientProfile resource.

c) Select the GET /object/anyconnectclientprofiles method and click the **Try It Out!** button.

Each profile object will be represented like the following. The highlighted attribute is the one you need to change.

```
{
        "version": "oiwtsaoxbmip7",
        "name": "amp-install-profile",
        "md5Checksum": "12f18388580d3bb2eb0a9dcd8f9a7150",
        "description": null,
        "diskFileName": "bad3506d-9440-11ea-97d2-4d3296494e7b.xml",
        "anyConnectModuleType": "ANY_CONNECT_CLIENT_PROFILE",
        "id": "bba6cd0e-9440-11ea-97d2-7b74302649a4",
        "type": "anyconnectclientprofile",
        "links": {
          "self": "https://10.89.5.38/api/fdm/v6/object/
anyconnectclientprofiles/bba6cd0e-9440-11ea-97d2-7b74302649a4"
          }
```

d) Find your object in the output, select the code and use Ctrl+click to copy it to the clipboard.

e) Select the PUT /object/anyconnectclientprofiles/{objId} method, and paste the contents into the **body** field.

f) Copy the **id** value and paste it into the **objId** edit box above the body. You can also find the object ID at the end of the "self" URL.

Parameters

Parameter	Value
objId	bba6cd0e-9440-11ea-97d2-7b74302649a4

body	{
	        "version": "oiwtsaoxbmip7",
	        "name": "amp-install-profile",
	        "md5Checksum":
	"12f18388580d3bb2eb0a9dcd8f9a7150",
	        "description": null,
	        "diskFileName": "bad3506d-9440-11ea-

Parameter content type: application/json ▼

g) In the object body, find the **anyConnectModuleType** field and replace the value with the one for your profile type. Choose from DART, FEEDBACK, WEB_SECURITY, ANY_CONNECT_CLIENT_PROFILE, AMP_ENABLER, NETWORK_ACCESS_MANAGER, NETWORK_VISIBILITY, START_BEFORE_LOGIN, ISE_POSTURE, UMBRELLA.

h) Again in the **body**, delete the **links** attribute, including the comma after the **type** value.

The object body should look similar to the following:

```
{
        "version": "oiwtsaoxbmip7",
        "name": "amp-install-profile",
        "md5Checksum": "12f18388580d3bb2eb0a9dcd8f9a7150",
        "description": null,
        "diskFileName": "bad3506d-9440-11ea-97d2-4d3296494e7b.xml",
        "anyConnectModuleType": "AMP_ENABLER",
        "id": "bba6cd0e-9440-11ea-97d2-7b74302649a4",
        "type": "anyconnectclientprofile"
        }
```

i) Click **Try It Out!** Examine the response to verify that the object was modified correctly. You should get a response code of 200, and a response body that echos your changes. You can use the GET method to further verify the results.

Allow Traffic Through the Remote Access VPN

You can use one of the following techniques to enable traffic flow in the remote access VPN tunnel.

- Configure the **sysopt connection permit-vpn** command, which exempts traffic that matches the VPN connection from the access control policy. The default for this command is **no sysopt connection permit-vpn**, which means VPN traffic must also be allowed by the access control policy.

  This is the more secure method to allow traffic in the VPN, because external users cannot spoof IP addresses in the remote access VPN address pool. The downside is that the VPN traffic will not be inspected, which means that intrusion and file protection, URL filtering, or other advanced features will not be applied to the traffic. This also means that no connection events will be generated for the traffic, and thus statistical dashboards will not reflect VPN connections.

  To configure this command, select the **Bypass Access Control policy for decrypted traffic** option in your RA VPN connection profiles.

- Create access control rules to allow connections from the remote access VPN address pool. This method ensures that VPN traffic is inspected and advanced services can be applied to the connections. The downside is that it opens the possibility for external users to spoof IP addresses and thus gain access to your internal network.

Verify the Remote Access VPN Configuration

After you configure the remote access VPN, and deploy the configuration to the device, verify that you can make remote connections.

If you encounter problems, read through the troubleshooting topics to help isolate and correct the problems. See Troubleshooting Remote Access VPNs, on page 28.

Procedure

Step 1 From an external network, establish a VPN connection using the AnyConnect Client.

Using a web browser, open **https://**_ravpn-address_, where _ravpn-address_ is the IP address or hostname of the outside interface on which you are allowing VPN connections. If necessary, install the client software and complete the connection. See How Users Can Install the AnyConnect Client Software, on page 3.

If you changed the port for remote access VPN connections, you must include the custom port in the URL. For example, if you changed the port to 4443: https://ravpn.example.com:4443

If you configured group URLs, also try those URLs.

Step 2 Log into the device CLI as explained in Logging Into the Command Line Interface (CLI). Alternatively, open the CLI Console.

Step 3 Use the **show vpn-sessiondb** command to view summary information about current VPN sessions.

The statistics should show your active AnyConnect Client session, and information on cumulative sessions, the peak concurrent number of sessions, and inactive sessions. Following is sample output from the command.

```
> show vpn-sessiondb
--------------------------------------------------------------------------
VPN Session Summary
--------------------------------------------------------------------------
                            Active : Cumulative : Peak Concur : Inactive
                            ----------------------------------------------
AnyConnect Client          :     1 :        49 :          3 :        0
  SSL/TLS/DTLS             :     1 :        49 :          3 :        0
Clientless VPN             :     0 :         1 :          1
  Browser                  :     0 :         1 :          1
--------------------------------------------------------------------------
Total Active and Inactive  :     1              Total Cumulative :    50
Device Total VPN Capacity  : 10000
Device Load                :    0%
--------------------------------------------------------------------------


--------------------------------------------------------------------------
Tunnels Summary
--------------------------------------------------------------------------
                            Active : Cumulative : Peak Concurrent
                            ----------------------------------------------
Clientless                 :     0 :         1 :              1
AnyConnect-Parent          :     1 :        49 :              3
```

```
SSL-Tunnel                        :      1 :        46 :              3
DTLS-Tunnel                       :      1 :        46 :              3
-------------------------------------------------------------------------
Totals                            :      3 :       142
-------------------------------------------------------------------------


-------------------------------------------------------------------------
IPv6 Usage Summary
-------------------------------------------------------------------------
                                  Active : Cumulative : Peak Concurrent
                                  ---------------------------------------
AnyConnect SSL/TLS/DTLS           :        :              :
  Tunneled IPv6                   :      1 :        20 :              2
-------------------------------------------------------------------------
```

Step 4 Use the **show vpn-sessiondb anyconnect** command to view detailed information about current VPN sessions.

Detailed information includes encryption used, bytes transmitted and received, and other statistics. If you use your VPN connection, you should see the bytes transmitted/received numbers change as you re-issue this command.

```
> show vpn-sessiondb anyconnect

Session Type: AnyConnect

Username     : priya                    Index        : 4820
Assigned IP  : 172.18.0.1               Public IP    : 192.168.2.20
Assigned IPv6: 2009::1
Protocol     : AnyConnect-Parent SSL-Tunnel DTLS-Tunnel
License      : AnyConnect Premium
Encryption   : AnyConnect-Parent: (1)none  SSL-Tunnel: (1)AES-GCM-256  DTLS-Tunnel: (1)AES256
Hashing      : AnyConnect-Parent: (1)none  SSL-Tunnel: (1)SHA384  DTLS-Tunnel: (1)SHA1
Bytes Tx     : 27731                    Bytes Rx     : 14427
Group Policy : MyRaVpn|Policy       Tunnel Group : MyRaVpn
Login Time   : 21:58:10 UTC Mon Apr 10 2017
Duration     : 0h:51m:13s
Inactivity   : 0h:00m:00s
VLAN Mapping : N/A                      VLAN         : none
Audt Sess ID : c0a800fd012d400058ebfff2
Security Grp : none                     Tunnel Zone  : 0
```

Managing the Remote Access VPN Configuration

Remote access VPN connection profiles define the characteristics that allow external users to make a VPN connection to the system using the AnyConnect Client. Each profile defines the AAA servers and certificates used to authenticate users, the address pool for assigning users IP addresses, and the group policies that define a variety of user-oriented attributes.

You would create multiple profiles if you need to provide variable services to different user groups, or if you have different authentication sources. For example, if your organization merges with a different organization that uses different authentication servers, you can create a profile for the new group that uses those authentication servers.

Procedure

Step 1 Click **View Configuration** in the **Device** > **Remote Access VPN** group.

The group shows summary information on how many connection profiles and group policies are currently configured.

Step 2 Click **Connection Profiles** in the table of contents if it is not already selected.

Step 3 Do any of the following:

- Click the + button to create a new connection profile. For detailed instructions, see Configure an RA VPN Connection Profile, on page 15.

- Click the view button ( ) to open a summary of the connection profile and connection instructions. Within the summary, you can click **Edit** to make changes.

- Click the delete button ( ) to delete a connection profile that you no longer need.

- Select **Group Policies** in the table of contents to define the user-oriented attributes for the connection profiles. See Configure Group Policies for RA VPN, on page 22.

Configure an RA VPN Connection Profile

You can create a remote access VPN connection profile to allow your users to connect to your inside networks when they are on external networks, such as their home network. Create separate profiles to accommodate different authentication methods.

Before you begin

Before configuring the remote access (RA) VPN connection:

- Download the required AnyConnect Client software packages from software.cisco.com to your workstation.

- The outside interface, the one that terminates remote access VPN connections, cannot also have a management access list that allows HTTPS connections on the same port. Either configure a different port for management access (see Configuring the HTTPS Port for Management Access on Data Interfaces), or configure a different port for the connection profile. Both services use port 443 by default, so one must change.

Procedure

Step 1 Click **View Configuration** in the **Device** > **Remote Access VPN** group.

The group shows summary information on how many connection profiles and group policies are currently configured.

Step 2 Click **Connection Profiles** in the table of contents if it is not already selected.

Step 3   Do one of the following:

- Click the + button to create a new connection profile.

- Click the view button (👁) to open a summary of the connection profile and connection instructions. Within the summary, you can click **Edit** to make changes.

Step 4   Configure the basic connection attributes.

- **Connection Profile Name**—The name for this connection, up to 50 characters without spaces. For example, MainOffice. You cannot use an IP address as the name.

  Note   The name you enter here is what users will see in the connection list in the AnyConnect Client client. Choose a name that will make sense to your users.

- **Group Alias**, **Group URL**—Aliases contain alternate names or URLs for a specific connection profile. VPN users can choose an alias name in the AnyConnect Client client in the list of connections when they connect to the threat defense device. The connection profile name is automatically added as a group alias. Aliases can be up to 31 characters.

  You can also configure the list of group URLs, which your endpoints can select while initiating the Remote Access VPN connection. If users connect using the group URL, the system will automatically use the connection profile that matches the URL. This URL would be used by clients who do not yet have the AnyConnect Client client installed.

  Add as many group aliases and URLs as required. These aliases and URLs must be unique across all connection profiles defined on the device. Group URLs must start with **https://**.

  For example, you might have the alias Contractor and the group URL https://ravpn.example.com/contractor. Once the AnyConnect Client client is installed, the user would simply select the group alias in the AnyConnect Client VPN drop-down list of connections.

Step 5   Configure the primary and optionally, secondary identity sources.

These options determine how remote users authenticate to the device to enable the remote access VPN connection. The simplest approach is to use AAA only and then select an AD realm or use the LocalIdentitySource. You can use the following approaches for **Authentication Type**:

- **AAA Only**—Authenticate and authorize users based on username and password. For details, see Configure AAA for a Connection Profile, on page 18.

- **Client Certificate Only**—Authenticate users based on client device identity certificate. For details, see Configure Certificate Authentication for a Connection Profile, on page 21.

- **AAA and ClientCertificate**—Use both username/password and client device identity certificate.

- **SAML**—Use a SAML server at the primary authentication. When using SAML, you cannot configure a fallback or a secondary authentication source. For details, see Configure AAA for a Connection Profile, on page 18.

Step 6   Configure the address pool for clients.

The address pool defines the IP addresses that the system can assign to remote clients when they establish a VPN connection. For more information, see Configure Client Addressing for RA VPN, on page 22.

Step 7   Click **Next**.

Step 8   Select the **Group Policy** to use for this profile.

The group policy sets terms for user connections after the tunnel is established. The system includes a default group policy named DfltGrpPolicy. You can create additional group policies to provide the services you require.

When you select a group policy, you are shown a summary of the group characteristics. Click **Edit** in the summary to make changes.

If the group policy you need does not yet exist, click **Create New Group Policy** in the drop-down list.

For detailed information about group policies, see .

Step 9   Click **Next**.

Step 10   Configure the global settings.

These options apply to every connection profile. After you create the first connection profile, these options are pre-configured for each subsequent profile. If you make changes, you are changing every configured connection profile.

- **Certificate of Device Identity**—Select the internal certificate used to establish the identity of the device. Clients must accept this certificate to complete a secure VPN connection. If you do not already have a certificate, click **Create New Internal Certificate** in the drop-down list. You must configure a certificate.

- **Outside Interface**—The interface to which users connect when making the remote access VPN connection. Although this is normally the outside (Internet-facing) interface, choose whichever interface is between the device and the end users you are supporting.

- **Fully-qualified Domain Name for the Outside Interface**—The name of the interface, for example, ravpn.example.com. If you specify a name, the system can create a client profile for you.

  Note   You are responsible for ensuring that the DNS servers used in the VPN and by clients can resolve this name to the outside interface's IP address. Add the FQDN to the relevant DNS servers.

- **Port**—The TCP port to use for RA VPN connections. The default is 443. If you need to connect to the device manager on the same interface used for RA VPN, you must change the port number either for the connection profile or for the device manager. Both services use 443 by default. Note that users will have to include the port number in the URL if you change the port for remote access VPN connections.

- **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)**—Whether to subject VPN traffic to the access control policy. Decrypted VPN traffic is subjected to access control policy inspection by default. Enabling the **Bypass Access Control policy for decrypted traffic** option bypasses the access control policy, but for remote access VPN, the VPN Filter ACL and the authorization ACL downloaded from the AAA server are still applied to VPN traffic.

  Note that if you select this option, the system configures the **sysopt connection permit-vpn** command, which is a global setting. This will also impact the behavior of site-to-site VPN connections. Also, you cannot make different selections for this option across your connection profiles: the feature is either on or off for all profiles.

  If you do not select this option, it might be possible for external users to spoof IP addresses in your remote access VPN address pool, and thus gain access to your network. This can happen because you will need to create access control rules that allow your address pool to have access to internal resources. If you use access control rules, consider using user specifications to control access, rather than source IP address alone.

The downside of selecting this option is that the VPN traffic will not be inspected, which means that intrusion and file protection, URL filtering, or other advanced features will not be applied to the traffic. This also means that no connection events will be generated for the traffic, and thus statistical dashboards will not reflect VPN connections.

- **NAT Exempt**—Enable NAT Exempt to exempt traffic to and from the remote access VPN endpoints from NAT translation. If you do not exempt VPN traffic from NAT, ensure that the existing NAT rules for the outside and inside interfaces do not apply to the RA VPN pool of addresses. NAT exempt rules are manual static identity NAT rules for a given source/destination interface and network combination, but they are not reflected in the NAT policy, they are hidden. If you enable NAT Exempt, you must also configure the following.

  Note that this is a global option; it applies to all connection profiles. Thus, simply add interfaces and inside networks, do not replace them, or you will be changing the NAT exempt settings for all the other connection profiles that you have already defined.

    - **Inside Interfaces**—Select the interfaces for the internal networks remote users will be accessing. NAT rules are created for these interfaces.

    - **Inside Networks**—Select the network objects that represent internal networks remote users will be accessing. The networks list must contain the same IP types as the address pools you are supporting.

- **Secure Client Package**—The AnyConnect Client full installation software images that you will support on RA VPN connections. For each package, the filename, including extensions, can be no more than 60 characters. You can upload separate packages for Windows, Mac, and Linux endpoints. However, you cannot configure different packages for different connection profiles. If you already configured a package for another profile, the package is pre-selected. Changing it will change it for all profiles.

  Download the packages from software.cisco.com. If the endpoint does not already have the right package installed, the system prompts the user to download and install the package after the user authenticates.

Step 11   Click **Next**.

Step 12   Review the summary.

First, verify that the summary is correct.

Then, click **Instructions** to see what end users need to do to initially install the AnyConnect Client software and test that they can complete a VPN connection. Click **Copy** to copy these instructions to the clipboard, and then distribute them to your users.

Step 13   Click **Finish**.

What to do next

Ensure that traffic is allowed in the VPN tunnel, as explained in .

Configure AAA for a Connection Profile

Authentication, Authorization, and Accounting (AAA) servers use username and password to determine if a user is allowed access to the remote access VPN. If you use RADIUS servers, you can distinguish authorization levels among authenticated users, to provide differential access to protected resources. You can also use RADIUS accounting services to keep track of usage.

When configuring AAA, you must configure a primary identity source. Secondary and fallback sources are optional. Use a secondary source if you want to implement dual authentication, for example, using RSA tokens or DUO.

Primary Identity Source Options

- **Primary Identity Source for User Authentication**—The primary identity source used for authenticating remote users. End users must be defined in this source or the optional fallback source to complete a VPN connection. Select one of the following:

  - An Active Directory (AD) identity realm. If the realm you need does not yet exist, click **Create New Identity Realm**.

  - A RADIUS server group.

  - LocalIdentitySource (the local user database)—You can define users directly on the device and not use an external server.

  - A Duo LDAP server. However, this is best used as a secondary authentication source to provide two-factor authentication, as described in How to Configure Two-Factor Authentication using Duo LDAP, on page 39. If you use it as a primary source, you will not get user identity information, and you will not see user information in the dashboards, nor will you be able to write user-based access control rules.

  - A SAML server. If you use a SAML server, you cannot configure a fallback or secondary authentication source. You can use RADIUS as an authorization server, but you must configure the RADIUS server so that authentication is not required. That is, so the RADIUS server will provide authorization information after the connection is authenticated by SAML.

- **SAML Login Experience**—If you select SAML as the primary authentication source, you need to select which client browser to use to complete the web authentication:

  - **VPN Client embedded browser**—The VPN client uses its embedded browser for web authentication, so the authentication applies to the VPN connection only. This is the default and requires no further configuration.

  - **Default OS Browser**—The VPN client uses the system's default browser for web authentication. This option enables single sign-on (SSO) between your VPN authentication and other corporate logins. Also choose this option if you want to support web authentication methods, such as biometric authentication, that cannot be performed in the embedded browser.

    You must upload a package that enables web authentication in the browser. Obtain packages from software.cisco.com. Note that the package you upload is used by all connection profiles that use SAML with the default OS browser; the packages are global, not connection-profile specific.

- **Fallback Local Identity Source**—If the primary source is an external server, you can select the LocalIdentitySource as a fallback in case the primary server is unavailable. If you use the local database as a fallback source, ensure that you define the same local usernames/passwords as the ones defined in the external server.

Advanced Options—Click the **Advanced** link and configure the following options:

- **Strip options**—A realm is an administrative domain. Enabling the following options allows the authentication to be based on the username alone. You can enable any combination of these options. However, you must select both check boxes if your server cannot parse delimiters.

- **Strip Identity Source Server from Username**—Whether to remove the identity source name from the username before passing the username on to the AAA server. For example, if you select this option and the user enters domain\username as the username, the domain is stripped off from the username and sent to AAA server for authentication. By default this option is unchecked.

- **Strip Group from Username**—Whether to remove the group name from the username before passing the username on to the AAA server. This option applies to names given in the username@domain format; the option strips the domain and @ sign. By default this option is unchecked.

- **Enable Password Management**—Whether to allow the user to change the password when it expires. If you do not select this option, when the user's password expires, the AnyConnect Client will refuse the connection and the user must go and change the password on the AAA server. If you select this option, AnyConnect Client prompts the user to change the password when it expires, which is much more convenient for the user. Select one of the following options. Also, ensure that you enable MSCHAPv2 on the AAA server.

  - **Notify user x days prior to password expiration** (LDAP only)—Starting the number of days you specify, warn the user of the upcoming password expiration. You can set the warning from 1-180 days, with 14 being the default.

  - **Notify user on the day of password expiration**—The user is not warned, but is still prompted to change the password when the password expires. Even if you set a warning period, RADIUS users always get this behavior.

Secondary Identity Source

- **Secondary Identity Source for User Authorization**—The optional second identity source. If the user successfully authenticates with the primary source, the user is prompted to authenticate with the secondary source. You can select an AD realm, RADIUS server group, Duo LDAP server, or the local identity source.

- **Advanced options**—Click the **Advanced** link and configure the following options:

  - **Fallback Local Identity Source for Secondary**—If the secondary source is an external server, you can select the LocalIdentitySource as a fallback in case the secondary server is unavailable. If you use the local database as a fallback source, ensure that you define the same local usernames/passwords as the ones defined in the secondary external server.

  - **Use Primary Username for Secondary Login**—By default, when using a secondary identity source, the system will prompt for both username and password for the secondary source. If you select this option, the system prompts for the secondary password only, and uses the same username for the secondary source that was authenticated against the primary identity source. Select this option if you configure the same usernames in both the primary and secondary identity sources.

  - **Username for Session Server**—After successful authentication, the username is shown in events and statistical dashboards, is used to determine matches for user- or group-based SSL decryption and access control rules, and is used for accounting. Because you are using two authentication sources, you need to tell the system whether to use the Primary or Secondary username as the user identity. By default, the primary name is used.

  - **Password Type**—How to obtain the password for the secondary server. This field applies only if you select **AAA and Client Certificate** for the authentication type, and for the certificate options,

you select both **Prefill username from certificate on user login window** and **Hide username in login window**. The default is **Prompt**, which means the user is asked to enter the password.

Select **Primary Identity Source Password** to automatically use the password entered when the user authenticated to the primary server.

Select **Common Password** to use the same password for every user, then enter that password in the **Common Password** field.

Additional Options

- **Authorization Server**—The RADIUS server group that has been configured to authorize remote access VPN users.

  After authentication is complete, authorization controls the services and commands available to each authenticated user. Authorization works by assembling a set of attributes that describe what the user is authorized to perform, their actual capabilities, and restrictions. Were you not to use authorization, authentication alone would provide the same access to all authenticated users. For information on configuring RADIUS for authorization, see Controlling User Permissions and Attributes Using RADIUS and Group Policies, on page 3.

  Note that if the system obtains authorization attributes from the RADIUS server that overlap those defined in the group policy, the RADIUS attributes override the group policy attributes.

- **Accounting Server**—(Optional.) The RADIUS server group to use to account for the remote access VPN session.

  Accounting tracks the services users are accessing as well as the amount of network resources they are consuming. The threat defense device reports user activity to the RADIUS server. Accounting information includes when sessions start and stop, usernames, the number of bytes that pass through the device for each session, the service used, and the duration of each session. You can then analyze the data for network management, client billing, or auditing. You can use accounting alone or together with authentication and authorization.

Configure Certificate Authentication for a Connection Profile

You can use certificates installed on the client device to authenticate remote access VPN connections. When using certificate authentication, ensure that the trusted CA certificate used to validate the remote access user connections includes the **SSL Client** option for **Validation Usage**.

When using client certificates, you can still configure a secondary identity source, fallback source, and authorization and accounting servers. These are AAA options; for details, see Configure AAA for a Connection Profile, on page 18.

Following are the certificate-specific attributes. You can configure these attributes separately for the primary and secondary identity sources. Configuring a secondary source is optional.

- **Username from Certificate**—Select one of the following:

  - **Map Specific Field**—Use the certificate elements in the order of **Primary Field** and **Secondary Field**. The defaults are CN (Common Name) and OU (Organizational Unit). Select the options that work for your organization. The fields are combined to provide the username, and this is the name used in events, dashboards, and for matching purposes in SSL decryption and access control rules.

- **Use entire DN (distinguished name) as username**—The system automatically derives the username from the DN fields.

- **Advanced options**—Click the **Advanced** link and configure the following options:

    - **Prefill username from certificate on user login window**—Whether to fill in the username field with the retrieved username when prompting the user to authenticate.

    - **Hide username in login window**—If you select the **Prefill** option, you can hide the username, which means the user cannot edit the username in the password prompt.

Configure Client Addressing for RA VPN

There must be a way for the system to provide an IP address to endpoints that connect to the remote access VPN. These addresses can be provided by the AAA server, a DHCP server, an IP address pool configured in the group policy, or an IP address pool configured in the connection profile. The system tries these resources in that order, and stops when it obtains an available address, which it then assigns to the client. Thus, you can configure multiple options to create a failsafe in case of an unusual number of concurrent connections.

Use one or more of the following methods to configure the address pool for a connection profile.

- **AAA Server**—First, configure a network object on the threat defense device that specifies a subnet for the address pool. Then, in the RADIUS server, configure the Address-Pools (217) attribute for the user with the object name. Also, specify the RADIUS server for authentication in the connection profile.

- **DHCP**—First, configure a DHCP server with one or more IPv4 address ranges for the RA VPN (you cannot configure IPv6 pools using DHCP). Then, create a host network object with the IP address of the DHCP server. You can then select this object in the **DHCP Servers** attribute of the connection profile. You can configure up to 10 DHCP servers.

  If the DHCP server has multiple address pools, you can use the **DHCP Scope** attribute in the group policy that you attach to the connection profile to select which pool to use. Create a host network object with the network address of the pool. For example, if the DHCP pool contains 192.168.15.0/24 and 192.168.16.0/24, setting the DHCP scope to 192.168.16.0 will ensure that an address from the 192.168.16.0/24 subnet will be selected.

- **Local IP address pools**—First, create up to six network objects that specify subnets. You can configure separate pools for IPv4 and IPv6. Then, select these objects in the **IPv4 Address Pool** and **IPv6 Address Pool** options, either in the group policy, or in the connection profile. You do not need to configure both IPv4 and IPv6, just configure the address scheme you want to support.

  You also do not need to configure the pool in both the group policy and the connection profile. The group policy overrides the connection profile settings, so if you configure the pools in the group policy, leave the options empty in the connection profile.

  Note that the pools are used in the order in which you list them.

Configure Group Policies for RA VPN

A group policy is a set of user-oriented attribute/value pairs for remote access VPN connections The connection profile uses a group policy that sets terms for user connections after the tunnel is established. Group policies let you apply whole sets of attributes to a user or a group of users, rather than having to specify each attribute individually for each user.

The system includes a default group policy named DfltGrpPolicy. You can create additional group policies to provide the services you require.

Procedure

Step 1   Click **View Configuration** in the **Device** > **Remote Access VPN** group.

The group shows summary information on how many connection profiles and group policies are currently configured.

Step 2   Click **Group Policies** in the table of contents.

Step 3   Do any of the following:

- Click the + button to create a new group. See the following topics for explanations of the attributes on the pages of the group policy:

- Click the edit button () to edit an existing group policy.

- Click the delete button () to delete a group that you no longer need. The group cannot be currently used in a connection profile.

General Attributes

The general attributes of a group policy define the name of the group and some other basic settings. The Name attribute is the only required attribute.

- **Name**—The name of the group policy. The name can be up to 64 characters, spaces are allowed.

- **Description**—A description of the group policy. The description can be up to 1,024 characters.

- **DNS Server**—Select the DNS server group that defines the DNS servers clients should use for domain name resolution when connected to the VPN. If the group you need is not yet defined, click **Create DNS Group** and create it now.

- **Banner**—The banner text, or welcome message, to present to users at login. The default is no banner. The length can be up to 496 characters. The AnyConnect Client supports partial HTML. To ensure that the banner displays properly to remote users, use the <BR> tag to indicate line breaks.

- **Default Domain**—The default domain name for users in the RA VPN. For example, example.com. This domain is added to hostnames that are not fully-qualified, for example, serverA instead of serverA.example.com.

- **Secure Client Profiles**—Click + and select the AnyConnect Client Profiles to use for this group. If you configure a fully-qualified domain name for the outside interface (in the connection profile), a default profile will be created for you. Alternatively, you can upload your own client profile. Create these profiles using the standalone AnyConnect Client Profile Editor, which you can download and install from software.cisco.com. If you do not select a client profile, the AnyConnect Client uses default values for all options. The items in this list are AnyConnect Client Profile objects rather than the profiles themselves. You can create (and upload) new profiles by clicking **Create New Secure Client Profile** in the drop-down list.

  You can select AnyConnect Client module profiles, such as AMP Enabler, in addition to the AnyConnect Client Profile. You can select one profile per module type.

Session Settings Attributes

The session settings of a group policy control how long users can connect through the VPN and how many separate connections they can establish.

- **Maximum Connection Time**—The maximum length of time, in minutes, that users are allowed to stay connected to the VPN without logging out and reconnecting, from 1- 4473924 or blank. The default is unlimited (blank), but the idle timeout still applies.

- **Connection Time Alert Interval**—If you specify a maximum connection time, the alert interval defines the amount of time before the maximum time is reached to display a warning to the user about the upcoming automatic disconnect. The user can choose to end the connection and reconnect to restart the timer. The default is 1 minute. You can specify 1 to 30 minutes.

- **Idle Time**—The length of time, in minutes, that the VPN connection can be idle before it is automatically closed, from 1-35791394. If there is no communication activity on the connection for this consecutive number of minutes, the system stops the connection. The default is 30 minutes.

- **Idle Time Alert Interval**—The amount of time before the idle time is reached to display a warning to the user about the upcoming automatic disconnect due to an idle session. Any activity resets the timer. The default is 1 minute. You can specify 1 to 30 minutes.

- **Simultaneous Login Per User**—The maximum number of simultaneous connections allowed for a user. The default is 3. You can specify 1 to 2147483647 connections. Allowing a large number of simultaneous connections might compromise security and affect performance.

Address Assignment Attributes

The address assignment attributes of a group policy define the IP address pool for the group. The pool defined here overrides the pool defined in any connection profile that uses this group. Leave these settings blank if you want to use the pool defined in the connection profile.

- **IPv4 Address Pool**, **IPv6 Address Pool**—These options define the address pools for the remote endpoints. Clients are assigned an address from these pools based on the IP version they use to make the VPN connection. Select a network object that defines a subnet for each IP type you want to support. Leave the list empty if you do not want to support that IP version. For example, you could define an IPv4 pool as 10.100.10.0/24. The address pool cannot be on the same subnet as the IP address for the outside interface.

You can specify a list of up to six address pools to use for local address allocation. The order in which you specify the pools is significant. The system allocates addresses from these pools in the order in which the pools appear.

- **DHCP Scope**—If you configure DHCP servers for the address pool in the connection profile, the DHCP scope identifies the subnets to use for the pool for this group. The DHCP server must also have addresses in the same subnet identified by the scope. The scope allows you to select a subset of the address pools defined in the DHCP server to use for this specific group.

  If you do not define a network scope, the DHCP server assigns IP addresses in the order of the address pools configured. It goes through the pools until it identifies an unassigned address.

  To specify a scope, select a network object that contains a routeable address on the same subnet as the desired pool, but not within the pool. The DHCP server determines which subnet this IP address belongs to and assigns an IP address from that pool.

  We recommend using the IP address of an interface whenever possible for routing purposes. For example, if the pool is 10.100.10.2-10.100.10.254, and the interface address is 10.100.10.1/24, use 10.100.10.1 as the DHCP scope. Do not use the network number. Click **Create New Network** if the object does not yet exist. You can use DHCP for IPv4 addressing only. If the address you choose is not an interface address, you might need to create a static route for the scope address.

Split Tunneling Attributes

The split tunneling attributes of a group policy define how the system should handle traffic meant for the internal network vs. externally-directed traffic. Split tunneling directs some network traffic through the VPN tunnel (encrypted) and the remaining network traffic outside the VPN tunnel (unencrypted or in clear text).

- **IPv4 Split Tunneling**, **IPv6 Split Tunneling**—You can specify different options based on whether the traffic uses IPv4 or IPv6 addressing, but the options for each are the same. If you want to enable split tunneling, specify one of the options that requires you to select network objects.

  - **Allow all traffic over tunnel**—Do no split tunneling. Once the user makes an RA VPN connection, all the user's traffic goes through the protected tunnel. This is the default. It is also considered the most secure option.

  - **Allow specified traffic over the tunnel**—Select the network objects that define destination network and host addresses. Any traffic to these destinations goes through the protected tunnel. Traffic to any other destination is routed by the client to connections outside the tunnel (such as a local Wi-Fi or network connection).

  - **Exclude networks specified below**—Select the network objects that define destination network or host addresses. Any traffic to these destinations is routed by the client to connections outside the tunnel. Traffic to any other destination goes through the tunnel.

- **Split DNS**—You can configure the system to send some DNS requests through the secure connection, while allowing the client to send other DNS requests to the DNS servers configured on the client. You can configure the following DNS behavior:

  - **Send DNS Request as per split tunnel policy**—With this option, DNS requests are handled the same way as the split tunnel options are defined. If you enable split tunneling, DNS requests are sent based on the destination addresses. If you do not enable split tunneling, all DNS requests go over the protected connection.

- **Always send DNS requests over tunnel**—Select this option if you enable split tunneling, but you want all DNS requests sent through the protected connection to the DNS servers defined for the group.

- **Send only specified domains over tunnel**—Select this option if you want your protected DNS servers to resolve addresses for certain domains only. Then, specify those domains, separating domain names with commas. For example, example.com, example1.com. Use this option if you want your internal DNS servers to resolve names for internal domains, while external DNS servers handle all other Internet traffic.

AnyConnect Client Attributes

The AnyConnect Client attributes of a group policy define some SSL and connection settings used by the AnyConnect Client for a remote access VPN connection.

SSL Settings

- **Enable Datagram Transport Layer Security (DTLS)**—Whether to allow the AnyConnect Client to use two simultaneous tunnels: an SSL tunnel and a DTLS tunnel. Using DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays. If you do not enable DTLS, AnyConnect Client users establishing SSL VPN connections connect with an SSL tunnel only.

- **DTLS Compression**—Whether to compress Datagram Transport Layer Security (DTLS) connections for this group using LZS. DTLS Compression is disabled by default.

- **SSL Compression**—Whether to enable data compression, and if so, the method of data compression to use, **Deflate**, or **LZS**. SSL Compression is **Disabled** by default. Data compression speeds up transmission rates, but also increases the memory requirement and CPU usage for each user session. Therefore, SSL compression decreases the overall throughput of the device.

- **SSL Rekey Method**, **SSL Rekey Interval**—The client can rekey the VPN connection, renegotiating the crypto keys and initialization vectors, to increase the security of the connection. Disable rekeying by selecting **None**. To enable rekey, select **New Tunnel** to create a new tunnel each time. (The **Existing Tunnel** option results in the same action as **New Tunnel**.) If you enable rekeying, also set the rekey interval, which is 4 minutes by default. You can set the interval to 4-10080 minutes (1 week).

Connection Settings

- **Ignore the DF (Don't Fragment) bit**—Whether to ignore the Don't Fragment (DF) bit in packets that need fragmentation. Select this option to allow the forced fragmentation of packets that have the DF bit set, so that these packets can pass through the tunnel.

- **Client Bypass Protocol**—Allows you to configure how the secure gateway manages IPv4 traffic (when it is expecting only IPv6 traffic), or how it manages IPv6 traffic (when it is expecting only IPv4 traffic).

  When the AnyConnect Client makes a VPN connection to the headend, the headend assigns it an IPv4, IPv6, or both an IPv4 and IPv6 address. If the headend assigns the AnyConnect Client connection only an IPv4 address or only an IPv6 address, you can configure the Client Bypass Protocol to drop network traffic for which the headend did not assign an IP address (default, disabled, not checked), or allow that traffic to bypass the headend and be sent from the client unencrypted or "in the clear" (enabled, checked).

For example, assume that the secure gateway assigns only an IPv4 address to the AnyConnect Client 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; however, if Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear.

- **MTU**—The maximum transmission unit (MTU) size for SSL VPN connections established by the AnyConnect Client. The default is 1406 bytes. The range is 576 to 1462 bytes.

- **Keepalive Messages Between Secure Client and VPN Gateway**—Whether to exchange keepalive messages between peers to demonstrate that they are available to send and receive data in the tunnel. Keepalive messages transmit at set intervals. The default interval is 20 seconds, the valid range is 15 to 600 seconds.

- **DPD on Gateway Side Interval**, **DPD on Client Side Interval**—Enable Dead Peer Detection (DPD) to ensure that the VPN gateway or VPN client quickly detects when the peer is no longer responding. You can separately enable gateway or client DPD. The default interval is 30 seconds for sending DPD messages. The interval can be 5-3600 seconds.

Traffic Filters Attributes

The traffic filter attributes of a group policy define restrictions you want to place on users assigned to the group. You can use these attributes instead of creating access control policy rules to restrict RA VPN users to specific resources, based on host or subnet address and protocol, or on VLAN.

By default, RA VPN users are not restricted by the group policy from accessing any destination on your protected network.

- **Access List Filter**—Restrict access using an extended access control list (ACL). Select the Smart CLI Extended ACL object, or click **Create Extended Access List** and create it now.

  The extended ACL lets you filter based on source address, destination address, and protocol (such as IP or TCP). ACLs are evaluated on a top-down, first-match basis, so ensure that you place specific rules before more general rules. There is an implicit "deny any" at the end of the ACL, so if your intention is to simply deny access to a few subnets while allowing all other access, ensure that you include a "permit any" rule at the end of the ACL. The VPN filter applies to initial connections only. It does not apply to secondary connections, such as a SIP media connection, that are opened due to the action of application inspection.

  Because you cannot create network objects while editing an extended ACL Smart CLI object, you should create the ACL before editing the group policy. Otherwise, you might need to simply create the object, then go back later to create the network objects and then all the access control entries that you need. To create the ACL, go to **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**, create an object, and select **Extended Access List** as the object type. For an example, see .

- **Restrict VPN to VLAN**—Also called "VLAN mapping," this attribute specifies the egress VLAN interface for sessions to which this group policy applies. The system forwards all traffic from this group to the selected VLAN.

  Use this attribute to assign a VLAN to the group policy to simplify access control. Assigning a value to this attribute is an alternative to using an ACL to filter traffic on a session. Ensure that you specify a VLAN number that is defined on a subinterface on the device. Values range from 1 to 4094.

Windows Browser Proxy Attributes

The Windows browser proxy attributes of a group policy determine how, and whether, a proxy defined on the user's browser operates.

You can select one of the following values for **Browser Proxy During VPN Session**:

- **No change in endpoint settings**—Allow the user to configure (or not configure) a browser proxy for HTTP, and use the proxy if it is configured.

- **Disable browser proxy**—Do not use the proxy defined for the browser, if any. No browser connections will go through the proxy.

- **Auto detect settings**—Enable the use of automatic proxy server detection in the browser for the client device.

- **Use custom settings**—Define a proxy that should be used by all client devices for HTTP traffic. Configure the following settings:

    - **Proxy Server IP or Hostname**, **Port**—The IP address, or hostname, of the proxy server, and the port used for proxy connections by the proxy server. The host and port combined cannot exceed 100 characters.

    - **Browser Exemption List**—Connections to the hosts/ports in the exemption list do not go through the proxy. Add all of the host/port values for destinations that should not use the proxy. For example, www.example.com port 80. Click the **Add** link to add items to the list. Click the trash can icon to delete items. The entire proxy exception list, combining all addresses and ports, cannot be longer than 255 characters.

Monitoring Remote Access VPN

To monitor and troubleshoot remote access VPN connections, open the CLI console or log into the device CLI and use the following commands.

- **show vpn-sessiondb** displays information about VPN sessions. You can reset these statistics using the **clear vpn-sessiondb** command.

- **show webvpn** *keyword* displays information about the remote access VPN configuration, including statistics and the AnyConnect images installed. Enter **show webvpn ?** to see the available keywords.

- **show aaa-server** displays statistics about the directory server used with remote access VPN.

Troubleshooting Remote Access VPNs

Remote access VPN connection issues can originate in the client or in the threat defense device configuration. The following topics cover the main troubleshooting problems you might encounter.

Troubleshooting SSL Connection Problems

If the user cannot make the initial, non-AnyConnect Client, SSL connection to the outside IP address to download the AnyConnect Client, do the following:

1. If you configured a non-default port for the remote access VPN connection profile, ensure the user is including the port number in the URL. For example: https://ravpn.example.com:4443

2. From the client workstation, verify that you can ping the IP address of the outside interface. If you cannot, determine why there is no route from the user's workstation to the address.

3. From the client workstation, verify that you can ping the fully-qualified domain name (FQDN) of the outside interface, the one defined in the remote access (RA) VPN connection profile. If you can ping the IP address but not the FQDN, then you need to update the DNS servers used by the client and RA VPN connection profile to add the FQDN-to-IP-address mapping.

4. Verify that the user is accepting the certificate presented by the outside interface. The user should accept it permanently.

5. Examine the RA VPN connection configuration and verify that you selected the correct outside interface. A common mistake is to select an inside interface, the one facing the internal networks, rather than the outside interface, which faces the RA VPN users.

6. If SSL encryption is properly configured, use an external sniffer to verify whether the TCP three-way handshake is successful.

Troubleshooting AnyConnect Client Download and Installation Problems

If the user can make an SSL connection to the outside interface, but cannot download and install the AnyConnect Client package, consider the following:

- Ensure that you uploaded an AnyConnect Client package for the client's operating system. For example, if the user's workstation runs Linux, but you did not upload a Linux AnyConnect Client image, there is no package that can be installed.

- For Windows clients, the user must have Administrator rights to install software.

- For Windows clients, the workstation must enable ActiveX or install Java JRE 1.5 or higher, with JRE 7 recommended.

- For Safari browsers, Java must be enabled.

- Try different browsers, one might fail where another succeeds.

Troubleshooting AnyConnect Client Connection Problems

If the user was able to connect to the outside interface, download, and install the AnyConnect Client, but could not then complete a connection using AnyConnect Client, consider the following:

- If authentication fails, verify that the user is entering the correct username and password, and that the username is defined correctly in the authentication server. The authentication server must also be available through one of the data interfaces.

Note If the authentication server is on an external network, you need to configure a site-to-site VPN connection to the external network, and include the remote access VPN interface address within the VPN. For details, see How to Use a Directory Server on an Outside Network with Remote Access VPN, on page 49.

- If you configured a fully-qualified domain name (FQDN) for the outside interface in the remote access (RA) VPN connection profile, verify that you can ping the FQDN from the client device. If you can ping the IP address but not the FQDN, then you need to update the DNS servers used by the client and RA VPN connection profile to add the FQDN-to-IP-address mapping. If you are using the default AnyConnect Client profile that is generated when you specify an FQDN for the outside interface, the user will need to edit the server address to use the IP address until DNS is updated.

- Verify that the user is accepting the certificate presented by the outside interface. The user should accept it permanently.

- If the user's AnyConnect Client includes multiple connection profiles, that they are selecting the right one.

- If everything seems right on the client end, make an SSH connection to the threat defense device, and enter the **debug webvpn** command. Examine the messages issued during a connection attempt.

Troubleshooting RA VPN Traffic Flow Problems

If the user can make a secure remote access (RA) VPN connection, but cannot send and receive traffic, do the following:

1. Have the client disconnect, then reconnect. Sometimes this eliminates the problem.

2. In the AnyConnect Client, check the traffic statistics to determine whether both the sent and received counters are increasing. If the received packet count stays at zero, the threat defense device is not returning any traffic. There is likely a problem in the threat defense configuration. Common problems include the following:

   - Access rules are blocking traffic. Check the access control policy for rules that prevent traffic between the inside networks and the RA VPN address pool. You might need to create an explicit Allow rule if your default action is to block traffic.

   - The VPN filter is blocking traffic. Check the ACL traffic filter or VLAN filter configured in the group policy for the connection profile. You might need to make adjustments in the ACL or change the VLAN, depending on how (or if) you are filtering traffic based on group policy.

   - NAT rules are not being bypassed for the RA VPN traffic. Ensure that NAT exempt is configured for the RA VPN connection for every inside interface. Alternatively, ensure that the NAT rules do not prevent communication between the inside networks and interfaces and the RA VPN address pool and outside interface.

   - Routes are misconfigured. Ensure that all defined routes are valid and functioning correctly. For example, if you have a static IP address defined for the outside interface, ensure that the routing table includes a default route (for 0.0.0.0/0 and ::/0).

   - Ensure that the DNS server and domain name configured for the RA VPN are correct, and that the client system is using the correct ones. Verify that the DNS servers are reachable.

   - If you enable split tunneling in the RA VPN, check whether traffic to the specified inside networks is going through the tunnel, while all other traffic is bypassing the tunnel (so that the threat defense device does not see it).

3. Make an SSH connection to the threat defense device and verify that traffic is being sent and received for the remote access VPN. Use the following commands.

> • **show webvpn anyconnect**
>
> • **show vpn-sessiondb**

Examples for Remote Access VPN

The following are examples of configuring remote access VPN.

How to Implement RADIUS Change of Authorization

RADIUS Change of Authorization (CoA), also known as dynamic authorization, provides end-point security for the threat defense remote access VPN. A key challenge for RA VPNs is to secure the internal network against compromised end points and to secure the end point itself when it is affected by viruses or malware, by remediating the attack on the endpoint. There is a need to secure the endpoint and the internal network in all phases, that is, before, during, and after the RA VPN session. The RADIUS CoA feature helps in achieving this goal.

If you use Cisco Identity Services Engine (ISE) RADIUS servers, you can configure Change of Authorization policy enforcement.

The ISE Change of Authorization feature provides a mechanism to change the attributes of an authentication, authorization, and accounting (AAA) session after it is established. When a policy changes for a user or user group in AAA, ISE sends CoA messages to the threat defense device to reinitialize authentication and apply the new policy. An Inline Posture Enforcement Point (IPEP) is not required to apply access control lists (ACLs) for each VPN session established with the threat defense device.

The following topics explain how CoA works, and how to configure it.

System Flow for Change of Authorization

Cisco ISE has a client posture agent that assesses an endpoint's compliance for criteria such as processes, files, registry entries, antivirus protection, antispyware protection, and firewall software installed on the host. Administrators can then restrict network access until the endpoint is in compliance or can elevate local user privileges so they can establish remediation practices. ISE Posture performs a client-side evaluation. The client receives the posture requirement policy from ISE, performs the posture data collection, compares the results against the policy, and sends the assessment results back to ISE.

Following is the system flow between the threat defense device, ISE, and the RA VPN client for Change of Authorization (CoA) processing.

1. The remote user starts an RA VPN session, using the AnyConnect Client, with the threat defense device.

2. The threat defense device sends a RADIUS Access-Request message for that user to the ISE server.

3. Because the client posture is unknown at this point, ISE matches the user to the authorization policy that is configured for unknown posture. This policy defines the following cisco-av-pair options, which ISE sends to the threat defense in a RADIUS Access-Accept response.

   • **url-redirect-acl**=*acl_name*, where *acl_name* is the name of an extended ACL that is configured on the threat defense device. This ACL defines which user traffic should be redirected to the ISE server, which is HTTP traffic. For example:

```
url-redirect-acl=redirect
```

- **url-redirect**=*url*, where the URL is the one to which traffic should be redirected. For example:

```
url-redirect=https://ise2.example.com:8443/guestportal/gateway?sessionId=xx&action=cpp
```

You must configure DNS for data interfaces so that the hostname can be resolved. If you also configure traffic filtering in the group policy for the connection profile, ensure that the client pool can reach the ISE server through the port (TCP/8443 in the example).

4. The threat defense device sends a RADIUS Accounting-Request start packet and receives a response from ISE. The accounting request includes all of the details of the session, including the session ID, the external IP address of the VPN client, and the IP address of the threat defense device. ISE uses the session ID to identify that session. The threat defense device also sends periodic interim account information, where the most important attribute is the Framed-IP-Address with the IP address that is assigned to the client by the threat defense device.

5. While in an unknown posture state, the threat defense device redirects traffic from the client that matches the redirect ACL to the redirect URL. ISE determines if the client has the required posture compliance module, and prompts the user to install it if necessary.

6. After the agent is installed on the client device, it automatically performs the checks that are configured in the ISE posture policy. The client communicates directly with ISE. It sends a posture report to ISE, which can include multiple exchanges using the SWISS protocol and ports TCP/UDP 8905.

7. When ISE receives the posture report from the agent, it processes the authorization rules once again. This time, the posture result is known and a different rule now matches the client. ISE sends a RADIUS CoA packet, which includes the downloadable ACL (DACL) for either compliant or non-compliant endpoints. For example, the compliant DACL might permit all access, while the non-compliant DACL denies all access. The contents of the DACL are up to the ISE administrator.

8. The threat defense device removes the redirection. If it does not have the DACLs cached, it must send an Access-Request in order to download them from ISE. The specific DACL is attached to the VPN session; it does not become part of the device configuration.

9. The next time that the RA VPN user tries to access the web page, the user can access the resources that are permitted by the DACL that is installed on the threat defense device for the session.

Note If the endpoint fails to satisfy any mandatory requirement and if a manual remediation is required, then a remediation window opens in the AnyConnect Client, displaying the items that require action. The remediation window runs in the background so that the updates on network activity do not pop up and interfere or cause disruption. A user can click **Details** in the ISE Posture tile portion of the AnyConnect Client to see what has been detected and what updates are needed before the user can join the network.

Configure Change of Authorization on the Threat Defense Device

Most of the Change of Authorization policy is configured in the ISE server. However, you must configure the threat defense device to connect to ISE correctly. The following procedure explains how to configure the threat defense side of the configuration.

Before you begin

If you use hostnames in any object, ensure that you configure DNS servers for use with the data interfaces, as explained in Configuring DNS for Data and Management Traffic. You typically need to configure DNS anyway to have a fully-functional system.

Procedure

Step 1 Configure the extended access control list (ACL) for redirecting initial connections to ISE.

The purpose of the redirect ACL is to send initial traffic to ISE so that ISE can assess the client posture. The ACL should send HTTPS traffic to ISE, but not traffic that is already destined for ISE, or traffic that is directed to a DNS server for name resolution. A sample redirect ACL might look like the following:

```
access-list redirect extended deny ip any host <ISE server IP>
access-list redirect extended deny ip any host <DNS server IP>
access-list redirect extended deny icmp any any
access-list redirect extended permit tcp any any eq www
```

However, note that ACLs have an implicit "deny any any" as the last access control entry (ACE). In this example, the last ACE, which matches TCP port www (that is, port 80), will not match any traffic that matches the first 3 ACEs, so those are redundant. You could simply create an ACL with the last ACE and get the same results.

Note that in a redirect ACL, the permit and deny actions simply determine which traffic matches the ACL, with permit matching and deny not matching. No traffic is actually dropped, denied traffic is simply not redirected to ISE.

To create the redirect ACL, you need to configure a Smart CLI object.

a) Choose **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**.
b) Click + to create a new object.
c) Enter a name for the ACL. For example, **redirect**.
d) For **CLI Template**, select **Extended Access List**.
e) Configure the following in the **Template** body:

- configure access-list-entry action = permit

- source-network = any-ipv4

- destination-network = any-ipv4

- configure permit port = any-source

- destination-port = HTTP

- configure logging = disabled

The ACE should look like the following:

f) Click **OK**.

This ACL will be configured the next time you deploy changes. You do not need to use the object in any other policy to force deployment.

Note	This ACL applies to IPv4 only. If you also want to support IPv6, simply add a second ACE with all the same attributes, except select any-ipv6 for the source and destination networks. You can also add the other ACEs to ensure traffic to the ISE or DNS server is not redirected. You will first need to create host network objects to hold the IP addresses of those servers.

Step 2   Configure a RADIUS server group for dynamic authorization.

There are several critical options that you must select correctly in the RADIUS server and server group objects to enable Change of Authorization, also known as dynamic authorization. The following procedure focuses on these attributes. For more details on these objects, see RADIUS Servers and Groups.

a) Choose **Objects** > **Identity Sources**.

b) Click + > **RADIUS Server**.

c) Enter a name for the server, and the hostname/IP address of the ISE RADIUS server, authentication port, and secret key configured on the server. Adjust the timeout if desired. These options are not directly related to dynamic authorization.

d) Click the RA VPN Only link and configure the following options:

- **Redirect ACL**—Select the extended ACL you created for redirection. In this example, the ACL named redirect.

- **Interface used to connect to Radius server**—Select **Manually Choose Interface**, and select the interface through which the server can be reached. You must select a specific interface so that the system can correctly enable the CoA listener on the interface.

  If the server is on the same network as the management address, which means you will select the diagnostic interface, you must also configure an IP address on the diagnostic interface. Having a management IP address is not sufficient. Go to **Device** > **Interfaces**, and configure an IP address on the diagnostic interface that is on the same subnet as the management IP address.

If you also use this server for the device manager administrative access, this interface is ignored. Administrative access attempts are always authenticated through the management IP address.

The following example shows the options configured for the inside interface.

e) Click **OK** to save the server object.

If you have a redundant setup, with multiple duplicate ISE RADIUS servers, create server objects for each of these servers.

f) Click + > **RADIUS Server Group**.

g) Enter a name for the server group, and adjust the dead time and maximum attempts if desired.

h) Select the **Dynamic Authorization** option, and change the port number if your ISE server is configured to use a different port. Port 1700 is the default port used for listening for CoA packets.

i) If the RADIUS server is configured to use an AD server for authenticating users, select the **Realm that Supports the RADIUS Server** that specifies the AD server used in conjunction with this RADIUS server. If the realm does not already exist, click **Create New Identity Realm** at the bottom of the list and configure it now.

j) Under **RADIUS Server**, click + and select the server object you created for RA VPN.

k) Click **OK** to save the server group object.

Step 3 Choose **Device** > **RA VPN** > **Connection Profiles**, and create a connection profile that uses this RADIUS server group.

Use **AAA Authentication** (either only or with certificates), and select the server group in the **Primary Identity Source for User Authentication**, **Authorization**, and **Accounting** options.

Configure all other options as needed for your organization.

Note If the DNS servers are reached through the VPN network, edit the group policy used in the connection profile to configure the **Split DNS** option on the Split Tunneling Attributes page.

Configure Change of Authorization in ISE

Most of the Change of Authorization configuration is done in the ISE server. ISE has a posture assessment agent that runs on the endpoint device, and ISE communicates directly with the device to determine posture stance. The threat defense device essentially waits for instructions from ISE on how to handle a given end user.

A full discussion of configuring posture assessment policies is outside the scope of this document. However, the following procedure explains some of the basics. Use this as a starting point for configuring ISE. Note that the exact command paths, page names, and attribute names can change from release to release. The version of ISE you are using might use different terminology or organization.

The minimum supported ISE release is 2.2 patch 1.

Before you begin

This procedure assumes you have already configured users in the ISE RADIUS server.

Procedure

Step 1   Choose **Administration** > **Network Resources** > **Network Devices** > **Network Devices**, add the threat defense device to the ISE Network Device inventory, and configure the RADIUS settings.

Select the **RADIUS Authentication Settings**, and configure the same **Shared Secret** that is configured in the threat defense RADIUS server object. If desired, change the **CoA Port** number and ensure you configure the same port in the threat defense RADIUS server group object.

Step 2   Choose **Policy** > **Policy Elements** > **Results** > **Authorization** > **Downloadable ACLs**.

Create 2 downloadable ACLs (DACL), one for use by compliant endpoints, one for non-compliant endpoints.

For example, you might allow all access for compliant endpoints (permit ip any any), while denying all access to non-compliant endpoints (deny ip any any). You can make these DACLs as complex as you require, to provide the exact access users should have based on their compliance state. You will use these DACLs in authorization profiles.

Step 3   Choose **Policy** > **Policy Elements** > **Results** > **Authorization** > **Authorization Profile** and configure the required profiles.

You need profiles for the following states. Minimum attributes for each are listed.

- **Unknown**—The unknown posture profile is the default posture profile. Every endpoint is matched to this policy when they initially establish the RA VPN connection. The point of this rule is to apply the redirect ACL and URL, and to download the posture agent if it is not already on the endpoint. Endpoints can remain attached to this profile if the agent is not installed, or if installation fails. Otherwise, after assessing the posture, endpoints move to the compliant or non-compliant profiles.

  Minimum attributes include the following:

  - **Name**—For example, PRE_POSTURE.

  - **Access Type**—Select **ACCESS_ACCEPT**.

  - **Common Tasks**—Select **Web Redirection (CWA, MDM, NSP, CPP)**, then select **Client Provisioning (Posture)**, and enter the name of the redirect ACL you configured on the threat defense device. In **Value**, select **Client Provisioning Portal** if it is not already selected.

  - The **Attribute Details** should show two cisco-av-pair values, for url-redirect-acl and url-redirect. ISE will send this data to the threat defense device, which will apply the criteria to the RA VPN user session.

- **Compliant**—After the posture assessment completes, if the endpoint meets all requirements configured for the endpoint, the client is considered compliant and gets this profile. You would typically give this client full access.

  Minimum attributes include the following:

    - **Name**—For example, FULL_ACCESS.

    - **Access Type**—Select **ACCESS_ACCEPT**.

    - **Common Tasks**—Select **DACL Name**, and select the downloadable ACL for compliant users, for example, PERMIT_ALL_TRAFFIC. ISE will send the ACL to the threat defense device, which will apply it to the user session. This DACL will replace the initial redirect ACL for the user session.

- **Non-compliant**—If the posture assessment determines that the endpoint does not meet all requirements, there is a countdown during which the client can bring the endpoint into compliance, for example, by installing required updates. The AnyConnect Client informs the user of the compliance issues. During the countdown, the endpoint remains in the unknown compliance state. If the endpoint remains non-compliant after the countdown expires, the session is marked non-compliant and it gets the non-compliant profile. You would typically prevent all access for this endpoint, or at least restrict access in some way.

  Minimum attributes include the following:

    - **Name**—For example, Non_Compliant.

    - **Access Type**—Select **ACCESS_ACCEPT**.

    - **Common Tasks**—Select **DACL Name**, and select the downloadable ACL for non-compliant users, for example, DENY_ALL_TRAFFIC. ISE will send the ACL to the threat defense device, which will apply it to the user session. This DACL will replace the initial redirect ACL for the user session.

Step 4   Choose **Policy** > **Policy Elements** > **Results** > **Client Provisioning** > **Resources** and configure the following resources:

- **AnyConnect package**—The head end package file, which you download from software.cisco.com. You need separate packages for the client platforms you support, so you might need to configure multiple types, such as AnyConnectDesktopWindows.

- **ISE Posture Configuration File (Type: AnyConnectProfile)**—This configuration file defines the settings that the compliance module uses to evaluate the end user's device. This file also defines the length of time the user has to bring a non-compliant device into compliance.

- **Compliance Module Package (Type: ComplianceModule)**—The AnyConnect Client Compliance Module file is the file which will be pushed down to the installed AnyConnect package to check endpoint compliance. Download this file using the **Add Resource from Cisco Site** command. Ensure that you download the correct module based on the AnyConnect Client packages you have configured, or users will get download failures. You can also find these files on software.cisco.com in the AnyConnect Client listings in the ISEComplianceModule folder.

- **AnyConnect Configuration File (Type: AnyConnectConfig)**—These AnyConnect Client release-specific settings define the **AnyConnect Package**, **Compliance Module**, and **ISE Posture** to apply. Because the packages are OS-specific, create separate configuration files for each client OS you will support (for example, Windows, MAC, Linux).

Step 5   Choose **Policy** > **Client Provisioning** and configure the client provisioning policy.

Create new rules, for example, with names like CoA_ClientProvisionWin, for each operating system that should implement CoA. Select the appropriate operating system for the rule, and in **Results**, select the AnyConnect Client configuration file you created for the OS as the **Agent**.

Disable the default OS-specific rules that you are replacing.

Step 6   Configure the posture policy.

In this step, you develop the posture requirements that make sense for your organization.

- Choose **Policy** > **Policy Elements** > **Conditions** > **Posture**, and define the simple posture conditions that need to be met. For example, you might require that the user have certain applications installed.

- Choose **Policy** > **Policy Elements** > **Results** > **Posture** > **Requirements**, and define the compliance module requirement for the endpoint.

- Choose **Policy** > **Posture** > **Posture Policy** and configure the policies for the supported operating systems.

Step 7   Choose **Policy** > **Policy Sets** > **Default** > **Authorization Policy** and create the policy.

Add rules for each of the compliant conditions. These sample values are based on the examples in previous steps.

- Unknown, for pre-posture and posture download.

  - Name—For example, PRE_POSTURE

  - Conditions—"Session-PostureStatus EQUALS Unknown" AND "Radius-NAS-Port-Type EQUALS Virtual"

  - Profiles—For example, PRE_POSTURE

- Compliant, for clients that satisfy posture requirements.

  - Name—For example, FULL_ACCESS

  - Conditions—"Session-PostureStatus EQUALS Compliant" AND "Radius-NAS-Port-Type EQUALS Virtual"

  - Profiles—For example, FULL_ACCESS

- Non-compliance, for clients that fail posture requirements.

  - Name—For example, NON-COMPLIANT

  - Conditions—"Session-PostureStatus EQUALS NonCompliant" AND "Radius-NAS-Port-Type EQUALS Virtual"

  - Profiles—For example, Non_Compliant

Step 8   (Optional.) Choose **Administration** > **Settings** > **Posture** > **Reassessments** and enable posture reassessment.

By default, posture is assessed at connection time only. You can enable posture reassessment to periodically check the posture of connected endpoints. You can set the reassessment interval to determine how often this occurs.

If the system fails reassessment, you can define how the system should respond. You can allow the user to continue (remain connected), log the user off, or ask the user to remediate the system.

How to Configure Two-Factor Authentication using Duo LDAP

You can use the Duo LDAP server as the secondary authentication source in conjunction with a Microsoft Active Directory (AD) or RADIUS server as the primary source. With Duo LDAP, the secondary authentication validates the primary authentication with a Duo passcode, push notification, or phone call.

The following topics explain the configuration in more detail.

System Flow for Duo LDAP Secondary Authentication

The following graphic shows how the threat defense and Duo work together to provide two-factor authentication using LDAP.

Following is an explanation of the system flow:

1. The user makes a remote access VPN connection to the threat defense device and provides username and password.

2. Threat Defense authenticates this primary authentication attempt with the primary authentication server, which might be Active Directory or RADIUS.

3. If the primary authentication works, the threat defense sends a request for secondary authentication to the Duo LDAP server.

4. Duo then authenticates the user separately, through push notification, text message with a passcode, or a telephone call. The user must complete this authentication successfully.

5. Duo responds to the threat defense device to indicate whether the user authenticated successfully.

6. If the secondary authentication was successful, the threat defense device establishes a remote access VPN connection with the user's AnyConnect Client.

Configure Duo LDAP Secondary Authentication

The following procedure explains the end-to-end process of configuring two-factor authentication, using Duo LDAP as the secondary authentication source, for remote access VPN. Note that you must have an account with Duo, and obtain some information from Duo, to complete this configuration.

Procedure

Step 1 Create a Duo account and obtain the integration key, secret key, and API hostname.

Following is an overview of the process. For details, please see the Duo web site, https://duo.com.

a) Sign up for a Duo account.

b) Log in to the Duo Admin Panel and navigate to **Applications**.

c) Click **Protect an Application** and locate Cisco SSL VPN in the applications list. Click **Protect this Application** to get your integration key, secret key, and API hostname. For help, see the Duo *Getting Started* guide, https://duo.com/docs/getting-started.

Step 2 Create a Duo LDAP identity source for the Duo LDAP server.

You must use the threat defense API to create the Duo LDAP object; you cannot create it using the device manager. You can either use the API Explorer, or write your own client application, to create the object. The following procedure explains how to create the object using API Explorer.

a) Log into the device manager, click the more options button (⋮), and choose **API Explorer**.

The system opens the API Explorer in a separate tab or window, depending on your browser settings.

b) (Optional.) Obtain the values needed to identify the interface the system should use to connect to the Duo LDAP server.

If you do not specify an interface, the system uses the routing table. If necessary, you can create a static route for the Duo LDAP server. Alternatively, you can specify the interface to use in the Duo LDAP object. If you want to specify the interface, use the various GET methods in the Interfaces group to obtain the needed values. You can use physical, subinterface, EtherChannel, or VLAN interfaces. For example, to get the values for a physical interface, use the GET /devices/default/interfaces method and find the object for the interface you need to use. You need the following values from the interface object:

- id

- type

- version

- name

c) Click on the **DuoLDAPIdentitySource** heading to open the group.

d) Click on the **POST /object/duoldapidentitysources** method.

e) Under the **Parameters** heading, for the **body** element, click in the **Example Value** display box in the **Data Type** column on the right. This action loads the example into the body value edit box.

f) In the **body value** edit box, do the following:

- Delete the following attribute lines: **version**, **id**. (These attributes are needed for PUT calls but not for POST.)

- For **name**, enter a name for the object, such as Duo-LDAP-server.

- For **description**, either enter a meaningful description of the object for your reference purposes, or delete the attribute line.

- For **apiHostname**, enter the API Hostname that you obtained from your Duo account. The hostname should look like the following, with the X's replaced with your unique value: API-XXXXXXXX.DUOSECURITY.COM. Uppercase is not required.

- For **port**, enter the TCP port to use for LDAPS. This should be 636 unless you have been told by Duo to use a different port. Note that you must ensure that your access control list allows traffic to the Duo LDAP server through this port.

- For **timeout**, enter the timeout, in seconds, to connect to the Duo server. The value can be 1-300 seconds. The default is 120. To use the default, either enter 120 or delete the attribute line.

- For **integrationKey**, enter the integration key that you obtained from your Duo account.

- For **secretKey**, enter the secret key that you obtained from your Duo account. This key will subsequently be masked.

- For **interface**, either enter the id, type, version, and name values of the interface to use to connect to the Duo LDAP server, or delete the 6 lines used to define the interface attribute, including the trailing closing brace.

- For **type**, leave the value as duoldapidentitysource.

For example, the object body might look like the following, where the apiHostname and integrationKey are obfuscated, but the intentionally faked secret key is shown:

```
{
  "name": "Duo-LDAP-server",
  "description": "Duo LDAP server for RA VPN",
  "apiHostname": "API-XXXXXXXX.DUOSECURITY.COM",
  "port": 636,
  "timeout": 120,
  "integrationKey": "XXXXXXXXXXXXXXXXXXXX",
  "secretKey": "123456789",
  "type": "duoldapidentitysource"
}
```

g) Click the **Try It Out!** button.

The system will issue the **curl** command to post the object to the device configuration. You are shown the curl command, the response body, and the response code. If you created a valid body, you should see **200** in the **Response Code** field.

If you made an error, look at the response body for an error message. You can correct the body value and try again.

h) Return to the device manager by clicking **Device** in the top menu.

i) Click **Objects**, then click **Identity Sources** in the table of contents.

Your Duo LDAP object should appear in the list. If it does not, go back to the API Explorer and try to create the object again. You can use the GET method to check whether it was actually created.

Note that you can delete the object using the device manager, but you cannot edit it or see its contents. You must use the API for those actions. The relevant methods are shown in the **DuoLDAPIdentitySource** group.

Step 3 Upload the trusted CA certificate for the Duo web site to the device manager.

The threat defense system must have the certificate needed to validate the connection to the Duo LDAP server. You can obtain and upload the certificate using this procedure, which was done with the Google Chrome browser. The exact steps for your browser might differ. Alternatively, you can go directly to https://www.digicert.com/digicert-root-certificates.htm and download the certificate, but the following procedure is generic and you can use it to obtain root trusted CA certificates for any site.

a) Open https://duo.com in your browser.
b) Click the site information link in the browser's URL field, then click the **Certificate** link. This action opens the certificate information dialog box.
c) Click the **Certificate Path** tab, and select the root (top) level of the path. In this case, DigiCert.
d) With DigiCert selected, click **View Certificate**. This action will open a new Certificate dialog box, and the General tab should indicate that it was issued to DigiCert High Assurance EV Root CA. This is the root CA certificate that you need to upload to the device manager.
e) Click the **Details** tab, then click the **Copy to File** button to start the certificate download wizard.
f) Use the wizard to download the certificate to your workstation. Download using the default DER format.
g) In the device manager, choose to **Objects** > **Certificates**.
h) Click + > **Add Trusted CA Certificate**.
i) Enter a name for the certificate, for example, DigiCert_High_Assurance_EV_Root_CA. (Spaces are not allowed.)
j) Click **Upload Certificate** and select the file you downloaded.

k) Click **OK**.

Step 4 Use the AnyConnect Client Profile Editor to create a profile that specifies 60 seconds or more for authentication timeout.

You need to give users extra time to obtain the Duo passcode and complete the secondary authentication. We recommend at least 60 seconds.

For details on creating AnyConnect Client profiles and uploading them, see Configure and Upload Client Profiles, on page 10. The following procedure explains how to configure the authentication timeout only, and then upload the profile to the threat defense. If you want to change other settings, you can do so now.

a) If you have not already done so, download and install the AnyConnect Client profile editor package. You can find this in the Cisco Software center (software.cisco.com) in the folder for your AnyConnect Client version.

b) Open the AnyConnect Client **VPN Profile Editor**.

c) Select **Preferences (Part 2)** in the table of contents, scroll to the end of the page, and change **Authentication Timeout** to 60 (or more). The following image is from the AnyConnect 4.7 VPN Profile Editor; previous or subsequent versions might be different.
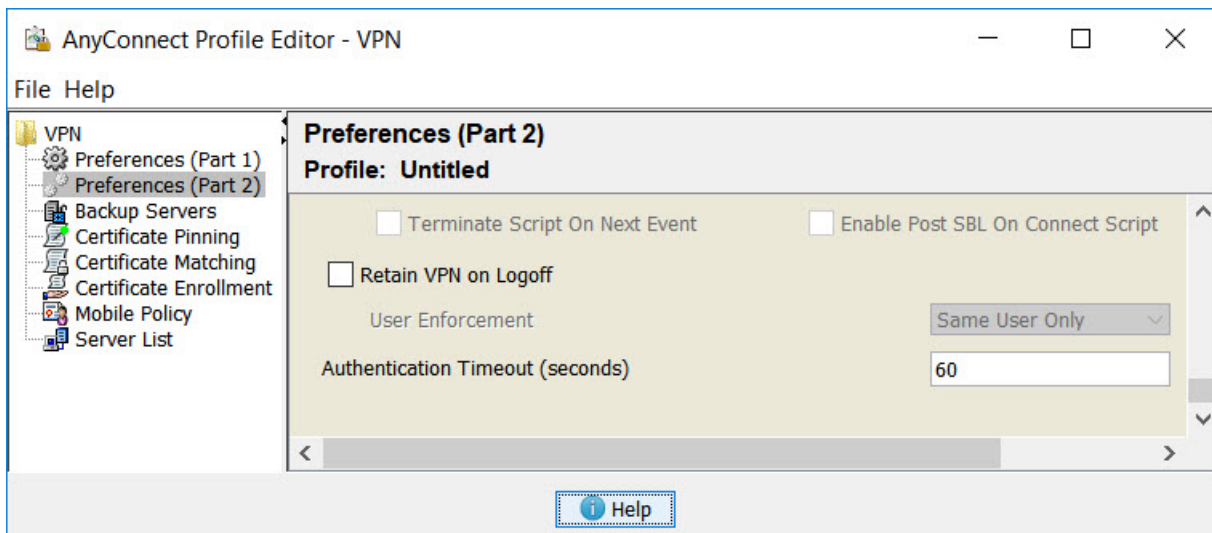
d) Choose **File** > **Save**, and save the profile XML file to your workstation with an appropriate name, for example, duo-ldap-profile.xml.

You can now close the VPN Profile Editor application.

e) In the device manager, choose **Objects** > **Secure Client Profiles**.

f) Click + to create a new profile object.

g) Enter a **Name** for the object. For example, Duo-LDAP-profile.

h) Click **Upload**, and select the XML file you created.

i) Click **OK**.

Step 5   Create a group policy and select the AnyConnect Client profile in the policy.

The group policy that you assign to a user controls many aspects of the connection. The following procedure explains how to assign the profile XML file to the group. For more details about what you can do with group policies, see Configure Group Policies for RA VPN, on page 22.

a) Click **View Configuration** in **Device** > **Remote Access VPN**.

b) Choose **Group Policies** in the table of contents.

c) Either edit DfltGrpPolicy, or click + and create a new group policy. For example, if you need a single remote access VPN connection profile for all users, editing the default group policy is appropriate.

d) On the General page, configure the following properties:

• **Name**—For a new profile, enter a name. For example, Duo-LDAP-group.

• **Secure Client Profiles**—Click + and select the AnyConnect Client profile object you created.

e) Click **OK** to save the group profile.

Step 6 Create or edit the remote access VPN connection profile to use for Duo-LDAP secondary authentication.

There are many steps to configuring a connection profile, which are explained in Configure an RA VPN Connection Profile, on page 15. The following procedure just mentions the key changes to make to enable Duo-LDAP as the secondary authentication source, and to apply the AnyConnect Client profile. For new connection profiles, you must configure the rest of the required fields. For this procedure, we assume you are editing an existing connection profile, and you simply need to change these two settings.

a) On the RA VPN page, choose **Connection Profiles** in the table of contents.
b) Either edit an existing connection profile, or create a new one.
c) Under Primary Identity Source, configure the following:

• **Authentication Type**—Choose either **AAA Only** or **AAA and Client Certificate**. You cannot configure two-factor authentication unless you use AAA.

• **Primary Identity Source for User Authentication**—Select your primary Active Directory or RADIUS server. Note that you can select a Duo-LDAP identity source as the primary source. However, Duo-LDAP provides authentication services only, not identity services, so if you use it as a primary authentication source, you will not see usernames associated with RA VPN connections in any dashboards, and you will not be able to write access control rules for these users. (You can configure fallback to the local identity source if you want to.)

• **Secondary Identity Source**—Select the Duo-LDAP identity source.

Primary Identity Source

Authentication Type

| AAA Only | Client Certificate Only | AAA and Client Certificate |

Primary Identity Source for User Authentication   Fallback Local Identity Source ⚠

AD ⌄   Please Select Local Identity Source ⌄

☐ Strip Identity Source server from username

☐ Strip Group from Username

Secondary Identity Source

Secondary Identity Source for User Authentication

Duo-LDAP-server ⌄

d) Click **Next**.
e) On the Remote User Experience page, select the **Group Policy** you created or edited.

Group Policy

Duo-LDAP-group

    f) Click **Next** on this page and the next page, Global Settings.

    g) Click **Finish** to save your changes to the connection profile.

Step 7   Commit your changes.

   a) Click the **Deploy Changes** icon in the upper right of the web page.

    b) Click the **Deploy Now** button.

    You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

How to Provide Internet Access on the Outside Interface for Remote Access VPN Users (Hair Pinning)

In remote access VPN, you might want users on the remote networks to access the Internet through your device. However, because the remote users are entering your device on the same interface that faces the Internet (the outside interface), you need to bounce Internet traffic right back out of the outside interface. This technique is sometimes called hair pinning.

The following graphic shows an example. There is a remote access VPN configured on the outside interface, 198.51.100.1. You want to split the remote user's VPN tunnel, so that Internet-bound traffic goes back out the outside interface, while traffic to your internal networks continue through the device. Thus, when a remote user wants to go to a server on the Internet, such as www.example.com, the connection first goes through the VPN, then gets routed back out to the Internet from the 198.51.100.1 interface.

The following procedure explains how to configure this service.

Before you begin

This example assumes that you have already registered the device, applied a remote access VPN license, and uploaded the AnyConnect Client image. It also assumes that you have configured the identity realm, which is also used in Identity policies.

Procedure

Step 1   Configure the remote access VPN connection.

The configuration requires a customized group policy in addition to the connection profile. Because hair pinning is a common configuration, and the required settings in the group policy are generally applicable, in this example we will edit the default group policy instead of creating a new group policy. You can take either approach.

a) Click **View Configuration** in the **Device** > **Remote Access VPN** group.

b) Click **Group Policies** in the table of contents, then click the edit icon (🔵) for the DfltGrpPolicy object.

c) Make the following changes to the default group policy:

• On the **General** page, in **DNS Server**, select the DNS server group that defines the servers VPN endpoints should use to resolve domain names.

> DNS Server
>
> CustomDNSServerGroup

• On the **Split Tunneling** page, for both **IPv4** and **IPv6 Split Tunneling**, select the **Allow all traffic over tunnel option**. This is the default setting, so it might already be configured correctly.

> IPv4 Split Tunneling   IPv6 Split Tunneling
>
> Allow all traffic over tunnel   ⌄   Allow all traffic over tunnel

Note	**This is a critical setting to enable hair-pinning.** You want all traffic to go to the VPN gateway, whereas split tunneling is a way to allow remote clients to directly access local or Internet sites outside of the VPN.

d) Click **OK** to save the changes to the default group policy.

e) Click **Connection Profiles** and either edit an existing profile or create a new one.

f) In the connection profile, page through the wizard and configure all options as you would for any other RA VPN configuration. However, you must configure the following options correctly to enable hair-pinning:

• **Group Policy**, in step 2. Select the group policy you customized for hair-pinning.

> Group Policy
>
> DfltGrpPolicy

• **NAT Exempt**, in step 3. Enable this feature. Select the inside interface, then select a network object that defines the internal networks. In this example, the object should specify 192.168.1.0/24. RA VPN traffic going to the internal network will not get address translation. However, because hair-pinned traffic is going out the outside interface, it will still be NAT'ed because the NAT exemption applies to the inside interface only. Note that if you have other connection profiles defined, you need to add to the existing settings, as the configuration applies to all connection profiles.

NAT Exempt

Inside Interfaces
The interfaces through which remote access VPN users can connect to the internal networks

+

🖥 inside

Inside Networks
The internal networks remote access VPN users are allowed to use. The IP versions of the internal networks and address pools must match, either IPv4, IPv6, or both.

+

🖳 local-network

Note   The **NAT Exempt** option is the other critical setting for the hair pin configuration.

g) (Optional.) In the **Global Settings** step, select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option.

By selecting this option, you remove the need to configure access control rules to allow traffic from RA VPN pool addresses. This option provides improved security (external users cannot spoof addresses in the pool), but it means that RA VPN traffic is exempt from inspection, including URL filtering and intrusion protection. Consider the pros and cons before deciding on this option.

h) Review the RA VPN configuration, then click **Finish**.

Step 2   Configure the NAT rule to translate all connections going out the outside interface to ports on the outside IP address (interface PAT).

When you complete the initial device configuration, the system creates a NAT rule named InsideOutsideNatRule. This rule applies interface PAT to IPv4 traffic from any interface that exits the device through the outside interface. Because the outside interface is included in "Any" source interface, the rule you need already exists, unless you edited it or deleted it.

The following procedure explains how to create the rule you need.

a) Click **Policies** > **NAT**.
b) Do one of the following:

  • To edit the InsideOutsideNatRule, mouse over the **Action** column and click the edit icon (🖉).

  • To create a new rule, click +.

c) Configure a rule with the following properties:

  • **Title**—For a new rule, enter a meaningful name without spaces. For example, OutsideInterfacePAT.

  • **Create Rule For**—**Manual NAT**.

  • **Placement**—**Before Auto NAT Rules** (the default).

  • **Type**—**Dynamic**.

  • **Original Packet**—For **Source Address**, select either Any or any-ipv4. For **Source Interface**, ensure that you select Any (which is the default). For all other Original Packet options, keep the default, Any.

• **Translated Packet**—For **Destination Interface**, select outside. For **Translated Address**, select **Interface**. For all other Translated Packet options, keep the default, Any.

The following graphic shows the simple case where you select Any for the source address.

d) Click **OK**.

Step 3   (If you do not configure **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** in the connection profile.) Configure an access control rule to allow access from the remote access VPN address pool.

If you select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** in the connection profile, traffic from RA VPN pool addresses bypasses the access control policy. You cannot write access control rules that will apply to the traffic. You need to write rules only if you disable the option.

The following example allows traffic from the address pool to any destination. You can adjust this to meet your specific requirements. You can also precede the rule with block rules to filter out undesirable traffic.

a) Click **Policies** > **Access Control**.
b) Click + to create a new rule.
c) Configure a rule with the following properties:

• **Order**—Select a position in the policy before any other rule that might match these connections and block them. The default is to add the rule to the end of the policy. If you need to reposition the rule later, you can edit this option or simply drag and drop the rule to the right slot in the table.

• **Title**—Enter a meaningful name without spaces. For example, RAVPN-address-pool.

- **Action**—**Allow**. You can select Trust if you do not want this traffic to be inspected for protocol violations or intrusions.

- **Source/Destination** tab—For **Source** > **Network**, select the same object you used in the RA VPN connection profile for the address pool. Leave the default, Any, for all other Source and Destination options.

SOURCE			DESTINATION		
Zones	Networks	Ports	Zones	Networks	Ports/Protocols
ANY	⬚ ravpn-pool	ANY	ANY	ANY	ANY

- **Application**, **URL**, and **Users** tabs—Leave the default settings on these tabs, that is, nothing selected.

- **Intrusion**, **File** tabs—You can optionally select intrusion or file policies to inspect for threats or malware.

- **Logging** tab—You can optionally enable connection logging.

    d) Click **OK**.

Step 4   Commit your changes.

    a) Click the **Deploy Changes** icon in the upper right of the web page.

    b) Click the **Deploy Now** button.

    You can wait until deployment completes, or click **OK** and check the task list or deployment history later.

How to Use a Directory Server on an Outside Network with Remote Access VPN

You can configure a remote access VPN to allow mobile workers and telecommuters to securely connect to your internal networks. Security of the connection depends on your directory server, which authenticates the user connection to ensure that only authorized users can gain entry.

If your directory server is on an outside network rather than an inside network, you need to configure a site-to-site VPN connection from the outside interface to the network that includes the directory server. **There is one trick to the site-to-site VPN configuration:** you must include the outside interface address of the remote access VPN device within the "inside" networks of the site-to-site VPN connection, and also in the remote networks for the device behind which the directory server resides. This will be explained further in the following procedure.

✎

Note   If you use the data interfaces as a gateway for the virtual management interface, this configuration also enables usage of the directory for identity policies. If you do not use data-interfaces as the management gateway, ensure that there is a route from the management network to the inside network that participates in the site-to-site VPN connection.

This use case implements the following network scenario.

Figure Callout	Description
1	Remote access host that makes a VPN connection to 192.168.4.6. Clients will get an address in the 172.18.1.0/24 address pool.
2	Site A, which hosts the remote access VPN.
3	The site-to-site VPN tunnel between the outside interfaces of the Site A and Site B the threat defense devices.
4	Site B, which hosts the directory server.
5	The directory server, on the inside network of Site B.

Before you begin

This use case assumes that you followed the device setup wizard to establish a normal baseline configuration. Specifically:

- There is an Inside_Outside_Rule access control rule that allows (or trusts) traffic going from the inside_zone to the outside_zone.

- The inside_zone and outside_zone security zones contain the inside and outside interfaces (respectively).

- There is an InsideOutsideNATRule that performs interface PAT for all traffic coming from inside interfaces going to the outside interface. On devices that use an inside bridge group by default, there might be several rules for interface PAT.

- There is a static IPv4 route for 0.0.0.0/0 that points to the outside interface. This example assumes that you are using static IP addresses for the outside interfaces, but you could also use DHCP and obtain the static route dynamically. For this example, we are assuming the following static routes:

    - Site A: outside interface, gateway is 192.168.4.254.

    - Site B: outside interface, gateway is 192.168.2.254.

Procedure

Step 1 Configure the site-to-site VPN connection on **Site B**, which hosts the directory server.

a) Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

b) Click the + button.

c) Configure the following options for **Endpoint Settings**.

- **Connection Profile Name**—Enter a name, for example, SiteA (to indicate that the connection is to Site A).

- **Local Site**—These options define the local endpoint.

  - **Local VPN Access Interface**—Select the **outside** interface (the one with the 192.168.2.1 address in the diagram).

  - **Local Network**—Click + and select the network object that identifies the local network that should participate in the VPN connection. Because the directory server is on this network, it can participate in the site-to-site VPN. Assuming that the object does not already exist, click **Create New Network** and configure an object for the 192.168.1.0/24 network. After saving the object, select it in the drop-down list and click **OK**.

Add Network Object

Name

Network192.168.1.0

Description

Type

● Network   ○ Host

Network

192.168.1.0/24

- **Remote Site**—These options define the remote endpoint.

  - **Remote IP Address**—Enter 192.168.4.6, which is the IP address of the remote VPN peer's interface that will host the VPN connection.

  - **Remote Network**—Click + and select the network objects that identify the remote networks that should participate in the VPN connection. Click **Create New Network**, configure the following objects, then select them in the list.

    1. SiteAInside, Network, 192.168.3.0/24.

Add Network Object

Name

SiteAInside

Description

Type

◉ Network   ○ Host

Network

192.168.3.0/24

2. SiteAInterface, Host, 192.168.4.6. **This is key: you must include the remote access VPN connection point address as part of the remote network for the site-to-site VPN connection so that the RA VPN hosted on that interface can use the directory server.**

Add Network Object

Name

SiteAInterface

Description

Type

○ Network   ◉ Host

Host

192.168.4.6

When you are finished, the endpoint settings should look like the following:

Connection Profile Name

SiteA

LOCAL SITE

Local VPN Access Interface

outside

Local Network

+

🖥 Network192.168.1.0

REMOTE SITE

◉ Static   ◯ Dynamic

Remote IP Address

192.168.4.6

Remote Network

+

🖥 SiteAInside

🖥 SiteAInterface

d) Click **Next**.

e) Define the privacy configuration for the VPN.

For this use case, we assume you qualify for export controlled features, which allows the use of strong encryption. Adjust these example settings to meet your needs and your license compliance.

- **IKE Version 2**, **IKE Version 1**—Keep the defaults, **IKE Version 2** enabled, **IKE Version 1** disabled.

- **IKE Policy**—Click **Edit** and enable **AES-GCM-NULL-SHA** and **AES-SHA-SHA**, and disable **DES-SHA-SHA**.

- **IPsec Proposal**—Click **Edit**. In the Select IPSec Proposals dialog box, click +, then click **Set Default** to choose the default AES-GCM proposals.

- **Local Preshared Key**, **Remote Peer Preshared Key**—Enter the keys defined on this device and on the remote device for the VPN connection. These keys can be different in IKEv2. The key can be 1-127 alphanumeric characters. **Remember these keys, because you must configure the same strings when creating the site-to-site VPN connection on the Site A device.**

The IKE policy should look like the following:

IKE Version 2

IKE Version 1

IKE Policy

Globally applied EDIT...

IPSec Proposal

Default set selected EDIT...

Authentication Type

◉ Pre-shared Manual Key   ○ Certificate

Local Pre-shared Key

••••••••••

Remote Peer Pre-shared Key

••••••••••

f) Configure the **Additional Options**.

- **NAT Exempt**—Select the interface that hosts the inside network, in this example, the **inside** interface. Typically, you do not want traffic within a site-to-site VPN tunnel to have their IP addresses translated. This option works only if the local network resides behind a single routed interface (not a bridge group member). If the local network is behind more than one routed interface, or one or more bridge group members, you must manually create the NAT exempt rules. For information on manually creating the required rules, see Exempting Site-to-Site VPN Traffic from NAT.

- **Diffie-Helman Group for Perfect Forward Secrecy**—Select **Group 19**. This option determines whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. For an explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use.

The options should look like the following.

Additional Options

NAT Exempt

inside

Diffie-Hellman Group for Perfect Forward Secrecy

19

g) Click **Next**.

h) Review the summary and click **Finish**.

The summary information is copied to the clipboard. You can paste the information in a document and use it to help you configure the remote peer, or to send it to the party responsible for configuring the peer.

i) Click the **Deploy Changes** icon in the upper right of the web page.

j) Click the **Deploy Now** button and wait for deployment to complete successfully.

Now the Site B device is ready to host one end of the site-to-site VPN connection.

Step 2 Log out of the **Site B** device and log into the **Site A** device.

Step 3 Configure the site-to-site VPN connection on **Site A**, which will host the remote access VPN.

a) Click **Device**, then click **View Configuration** in the Site-to-Site VPN group.

b) Click the + button.

c) Configure the following options for **Endpoint Settings**.

- **Connection Profile Name**—Enter a name, for example, SiteB (to indicate that the connection is to Site B).

- **Local Site**—These options define the local endpoint.

  - **Local VPN Access Interface**—Select the **outside** interface (the one with the 192.168.4.6 address in the diagram).

  - **Local Network**—Click + and select the network objects that identify the local networks that should participate in the VPN connection. Click **Create New Network**, configure the following objects, then select them in the list. **Note that you created the same objects in the Site B device, but you have to create them again in the Site A device.**

    1. SiteAInside, Network, 192.168.3.0/24.

    2. SiteAInterface, Host, 192.168.4.6. **This is key: you must include the remote access VPN connection point address as part of the inside network for the site-to-site VPN connection so that the RA VPN hosted on that interface can use the directory server on the remote network.**

Add Network Object

Name

SiteAInterface

Description

Type

○ Network   ⦿ Host

Host

192.168.4.6

- **Remote Site**—These options define the remote endpoint.
    - **Remote IP Address**—Enter 192.168.2.1, which is the IP address of the remote VPN peer's interface that will host the VPN connection.
    - **Remote Network**—Click + and select the network object that identifies the remote network that should participate in the VPN connection, the one that includes the directory server. Click **Create New Network** and configure an object for the 192.168.1.0/24 network. After saving the object, select it in the drop-down list and click **OK**. **Note that you created the same object in the Site B device, but you have to create it again in the Site A device.**

Add Network Object

Name

Network192.168.1.0

Description

Type

⦿ Network   ○ Host

Network

192.168.1.0/24

When you are finished, the endpoint settings should look like the following. Notice that the local/remote networks are flipped compared to the Site B settings. This is how the two ends of a point-to-point connection should always look.

Connection Profile Name

SiteB

LOCAL SITE

Local VPN Access Interface

outside

Local Network

+

SiteAInside

SiteAInterface

REMOTE SITE

● Static   ○ Dynamic

Remote IP Address

192.168.2.1

Remote Network

+

Network192.168.1.0

d) Click **Next**.

e) Define the privacy configuration for the VPN.

Configure the same IKE version, policy, and IPsec proposal, and the same preshared keys, as you did for the Site B connection, **but make sure that you reverse the Local and Remote preshared keys**.

The IKE policy should look like the following:

IKE Version 2

IKE Version 1

IKE Policy

Globally applied   EDIT...

IPSec Proposal

Default set selected   EDIT...

Authentication Type

● Pre-shared Manual Key   ○ Certificate

Local Pre-shared Key

••••••••••

Remote Peer Pre-shared Key

••••••••••

f) Configure the **Additional Options**.

- **NAT Exempt**—Select the interface that hosts the inside network, in this example, the **inside** interface. Typically, you do not want traffic within a site-to-site VPN tunnel to have their IP addresses translated. This option works only if the local network resides behind a single routed interface (not a bridge group member). If the local network is behind more than one routed interface, or one or more bridge group members, you must manually create the NAT exempt rules. For information on manually creating the required rules, see Exempting Site-to-Site VPN Traffic from NAT.

- **Diffie-Helman Group for Perfect Forward Secrecy**—Select **Group 19**.

The options should look like the following.

Additional Options

NAT Exempt

| inside | ∨ |

Diffie-Hellman Group for Perfect Forward Secrecy

| 19 | ∨ |

g) Click **Next**.

h) Review the summary and click **Finish**.

i) Click the **Deploy Changes** icon in the upper right of the web page.

j) Click the **Deploy Now** button and wait for deployment to complete successfully.

Now the Site A device is ready to host the other end of the site-to-site VPN connection. Because Site B is already configured with compatible settings, the two devices should negotiate a VPN connection.

You can confirm the connection by logging into the device CLI and pinging the directory server. You can also use the **show ipsec sa** command to view the session information.

Step 4 Configure the directory server on **Site A**. Click **Test** to verify that there is a connection.

a) Select **Objects**, then select **Identity Sources** from the table of contents.

b) Click + > **AD**.

c) Configure the basic realm properties.

- **Name**—A name for the directory realm. For example, AD.

- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.

- **Directory Username**, **Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

Note The system generates ldap-login-dn and ldap-login-password from this information. For example, Administrator@example.com is translated as cn=adminisntrator,cn=users,dc=example,dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name "users" folder.

- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, cn=users,dc=example,dc=com. For information on finding the base DN, see Determining the Directory Base DN.

- **AD Primary Domain**— The fully qualified Active Directory domain name that the device should join. For example, example.com.

Name

AD

Type

Active Directory (AD)

Directory Username

Administrator@example.com

e.g. user@example.com

Directory Password

••••••••

Base DN

cn=users,dc=example,dc=com

e.g. ou=user, dc=example, dc=com

AD Primary Domain

example.com

e.g. example.com

d) Configure the directory server properties.

- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address. For this example, enter 192.168.1.175.

- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method. For this example, keep 389.

- **Encryption**—To use an encrypted connection for downloading user and group information. The default is **None**, which means that user and group information is downloaded in clear text. For RA VPN, you can use **LDAPS**, which is LDAP over SSL. Use port 636 if you select this option. RA VPN does not support STARTTLS. For this example, select **None**.

- **Trusted CA Certificate**—If you select an encryption method, upload a Certificate Authority (CA) certificate to enable a trusted connection between the system and the directory server. If you are using a certificate to authenticate, the name of the server in the certificate must match the server Hostname / IP Address. For example, if you use 192.168.1.175 as the IP address but ad.example.com in the certificate, the connection fails.

Directory Server Configuration

Hostname / IP Address

192.168.1.175

e.g. ad.example.com

Port

389

Encryption

NONE

Trusted CA certificate

Please select a certificate

e) Click the **Test** button to verify the system can contact the server.

The system uses separate processes to access the server, so you might get errors indicating that the connection works for one type of use but not another, for example, available for Identity policies but not for remote access VPN. If the server cannot be reached, verify that you have the right IP address and host name, that the DNS server has an entry for the hostname, and so forth. Also, verify that the site-to-site VPN connection is working and that you included Site A's outside interface address in the VPN, and that NAT is not translating traffic for the directory server. You might also need to configure a static route for the server.

f) Click **OK**.

Step 5 Click **Device** > **Smart License** > **View Configuration**, and enable the RA VPN license.

When enabling the RA VPN license, select the type of license you purchased: Plus, Apex (or both), or VPN Only. For more information, see Licensing Requirements for Remote Access VPN, on page 7.

RA VPN License

Type   PLUS ˅   DISABLE

✔ Enabled

Please select the license type that you purchased to enable remote access VPN. Note that Firepower Device Manager does not support any of the advanced features covered by the Apex license.

Includes: RA-VPN

Step 6 Configure the remote access VPN on Site A.

a) Click **View Configuration** in the **Device** > **Remote Access VPN** group. Ensure that you are on the **Connection Profiles** page.

b) Create or edit a connection profile.

c) In the first step of the wizard, configure the profile name and then select the AD realm as the primary authentication source. You can optionally select the local database as fallback identity source.

Primary Identity Source

Authentication Type

| AAA Only | Client Certificate Only | AAA and Client Certificate |

Primary Identity Source for User Authentication

AD   ˅

Fallback Local Identity Source ⚠

LocalIdentitySource   ˅

d) Configure the address pool.

For this example, click +, then select **Create New Network** in the IPv4 address pool and create an object for the 172.18.1.0/24 network, then select the object. Clients are assigned an address from this pool. Leave the IPv6 pool blank. The address pool cannot be on the same subnet as the IP address for the outside interface.

The object should look like the following:

Name

ra-vpn-pool

Description

Type

⦿ Network

Network

172.18.1.0/24

The pool specification should look like the following:

Client Address Pool Assignment

IPv4 Address Pool
Endpoints are provided an address from this pool

+

🖥 ra-vpn-pool

IPv6 Address Pool
Endpoints are provided an address from this pool

+

DHCP Servers

+

e) Click **Next**, then select an appropriate group policy.

Check the summary information about the policy you select. Ensure that the DNS servers are configured. If they are not, edit the policy now and configure DNS.

f) Click **Next**, and in global settings, select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option, and configure the **NAT Exempt** options.

For **NAT Exempt**, you need to configure the following options. Note that if you have other connection profiles defined, you need to add to the existing settings, as the configuration applies to all connection profiles.

- **Inside Interfaces**—Select the **inside** interface. These are the interfaces for the internal networks remote users will be accessing. NAT rules are created for these interfaces.

- **Inside Networks**—Select the SiteAInside network object. These are the network objects that represent internal networks remote users will be accessing.

g) Upload the AnyConnect Client packages for the platforms you support.

h) Click **Next** and verify the settings.

First, verify that the summary is correct.

Then, click **Instructions** to see what end users need to do to initially install the AnyConnect Client software and test that they can complete a VPN connection. Click **Copy** to copy these instructions to the clipboard, and paste them in a text file or email.

i) Click **Finish**.

Step 7   Click the **Deploy Changes** icon in the upper right of the web page.

Step 8   Click the **Deploy Now** button and wait for deployment to complete successfully.

Now the Site A device is ready to accept RA VPN connections. Have an external user install the AnyConnect Client client and complete a VPN connection.

You can confirm the connection by logging into the device CLI and using the **show vpn-sessiondb anyconnect** command to view the session information.

How to Control RA VPN Access By Group

You can configure remote access VPN connection profiles to provide differential access to internal resources based on group policy. For example, if you want to provide unrestricted access to employees, but for contractors provide access to a single internal network and nothing else, you can use group policies to define different ACLs to restrict access appropriately.

The following example shows how to set up an RA VPN connection for contractors who should get access to the 192.168.2.0/24 internal subnet only. For regular employees, you can use the default group policy, which

does not have a traffic filter defined for the VPN. You can edit the default group policy if you want to apply restrictions to these users, and apply an ACL constructed as described below.

Before you begin

This procedure assumes that you have already created the identity source to use for the contractors. This might be a different source from the one you use for regular employees. Because the identity source is not strictly relevant to restricting access, we omit it from this example.

This example also assumes that the "inside2" interface is configured to host the 192.168.2.0/24 subnet, with the IP address 192.168.2.1 (any other address on the subnet is also acceptable).

Procedure

Step 1 Configure the extended access control list (ACL) for restricting RA VPN traffic.

You need to first configure the network object that defines the target 192.168.2.0/24, then create the Smart CLI object that defines the access list. Because the ACL has an implicit deny at the end, you need only permit access to the subnet, and traffic directed to any IP address outside the subnet will be denied. This example applies to IPv4 only; you can also configure objects for restricting IPv6 access to particular subnets. Simply create the network object and add an IPv6-based ACE to the same ACL.

a) Choose **Objects** > **Networks**, and create the required object.

For example, name the object ContractNetwork. The object should look similar to the following:

Name

ContractNetwork

Description

Type

⦿ Network   ◯ Host

Network

192.168.2.0/24

e.g. 192.168.2.0/24

b) Choose **Device** > **Advanced Configuration** > **Smart CLI** > **Objects**.
c) Click + to create a new object.
d) Enter a name for the ACL. For example, **ContractACL**.
e) For **CLI Template**, select **Extended Access List**.
f) Configure the following in the **Template** body:

   • configure access-list-entry action = permit

   • source-network = any-ipv4

   • destination-network = ContractNetwork object

- configure permit port = any
- configure logging = default

The ACE should look like the following:

g) Click **OK**.

This ACL will be configured the next time you deploy changes. You do not need to use the object in any other policy to force deployment.

Step 2 Create a group policy that uses the ACL.

At minimum, you should also configure DNS servers for the group policy. You can configure other options as needed. The following procedure focuses on the one setting that is relevant for this use case.

a) Choose **Device** > **RA VPN** > **Group Policies**.
b) Click + to create a new group policy.
c) On the **General** page, enter a name for the policy, such as **ContractGroup**.
d) Click **Traffic Filters** in the table of contents.
e) For **Access List Filter**, select the ContractACL object.

For this example, leave the VLAN option empty. Note that you could alternatively set up a VLAN for filtering purposes, and configure a subinterface for the VLAN.

f) Click **OK** to save the group policy.

Step 3 Configure the connection profile for contractors.

a) On the RA VPN page, click **Connection Profiles** in the table of contents.

b) Click + to create a new connection profile.

c) Complete step 1 of the wizard and click **Next**.

Enter a name for the profile, for example, Contractors.

Configure the rest of the options as normal. This includes selecting the appropriate authentication source for the contractors, and defining an address pool.

d) Select the group policy you configured for contractors and click **Next**.

Group Policy

ContractGroup

e) In the global settings, select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option, and configure the **NAT Exempt** options.

For **NAT Exempt**, you need to configure the following options. Note that if you have other connection profiles defined, you need to add to the existing settings, as the configuration applies to all connection profiles.

• **Inside Interfaces**—Select the **inside2** interface. These are the interfaces for the internal networks remote users will be accessing. NAT rules are created for these interfaces.

• **Inside Networks**—Select the ContractNetwork network object. These are the network objects that represent internal networks remote users will be accessing.

Access Control for VPN Traffic
Decrypted VPN traffic is subjected to access control policy inspection by default. Enabling the Bypass Access Control policy for decrypted traffic option bypasses the access control policy, but for remote access VPN, the VPN Filter ACL and the authorization ACL downloaded from the AAA server are still applied to VPN traffic

☑ Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)

NAT Exempt

Inside Interfaces
The interfaces through which remote access VPN users can connect to the internal networks

+

🖥 inside2

Inside Networks
The internal networks remote access VPN users are allowed to use. The IP versions of the internal networks and address pools must match, either IPv4, IPv6, or both.

+

🖵 ContractNetwork

f) Upload the AnyConnect Client packages for the platforms you support.

g) Click **Next** and verify the settings.

First, verify that the summary is correct.

Then, click **Instructions** to see what end users need to do to initially install the AnyConnect Client software and test that they can complete a VPN connection. Click **Copy** to copy these instructions to the clipboard, and paste them in a text file or email.

h) Click **Finish**.

How to Allow RA VPN Access to Internal Networks in Different Virtual Routers

If you configure multiple virtual routers on a device, you must configure the RA VPN in the global virtual router. You cannot configure RA VPN on an interface that is assigned to a custom virtual router.

Because the routing tables for virtual routers are separate, you must create static routes if your RA VPN users need to have access to networks that are part of a different virtual router.

Consider the following example. In this case, the RA VPN user connects to the outside interface at 172.16.3.1, and is given an IP address within the pool of 192.168.80.0/24. This user can now access the inside network that is attached to the global virtual router. However, the user cannot reach the 192.168.1.0/24 network that is part of virtual router VR1. To allow traffic flow between the VR1 network and the RA VPN user, you must configure static routes going both ways.

Before you begin

This example assumes that you have already configured the RA VPN, defined the virtual routers, and configured and assigned the interfaces to the appropriate virtual routers.

Procedure

Step 1 Configure the route leak from the Global virtual router to VR1.

This route allows the AnyConnect Clients assigned IP addresses in the VPN pool to access the 192.168.1.0/24 network in the VR1 virtual router.

a) Choose **Device** > **Routing** > **View Configuration**.

b) Click the view icon ( ) for the Global virtual router.

c) On the **Static Routing** tab for the Global router, click + and configure the route:

- **Name**—Any name will do, such as **ravpn-leak-vr1**.

- **Interface**—Select **vr1-inside**.

- **Protocol**—Select **IPv4**.

- **Networks**—Select an object that defines the 192.168.1.0/24 network. Click **Create New Network** to create the object now if necessary.

Name

nw-192-168.1.0

Description

Type
● Network   ○ Host

Network

192.168.1.0/24

e.g. 192.168.2.0/24 or 2001:DB8:0:C

- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:

Name

ravpn-leak-vr1

Description

⚠ The selected interface belongs to a different virtual router. If you create this static route, the route will cross virtual router boundaries, with the risk that traffic from this virtual router will leak into another virtual router. Proceed with caution.

Interface

vr1-inside (GigabitEthernet0/2) ⌄   ⬚   Belongs to different Router

⇥ VR1

Protocol

⦿ IPv4   ○ IPv6

Networks

+

🖥 nw-192-168-1.0

Gateway   **Metric**

Please select a gateway ⌄   1

SLA Monitor Applicable only for IPv4 Protocol type

Please select an SLA Monitor ⌄

    d) Click **OK**.

Step 2 Configure the route leak from VR1 to the Global virtual router.

This route allows endpoints on the 192.168.1.0/24 network to initiate connections to the AnyConnect Clients assigned IP addresses in the VPN pool.

    a) Choose **VR1** from the virtual routers drop-down list to switch to the VR1 configuration.

    b) On the **Static Routing** tab for the VR1 virtual router, click + and configure the route:

- **Name**—Any name will do, such as **ravpn-traffic**.

- **Interface**—Select **outside**.

- **Protocol**—Select **IPv4**.

- **Networks**—Select the object you created for the VPN pool, for example, **vpn-pool**.

- **Gateway**—Leave this item blank. When leaking a route into another virtual router, you do not select the gateway address.

The dialog box should look similar to the following:

Name

ravpn-traffic

Description

⚠ The selected interface belongs to a different virtual router. If you create this static
route, the route will cross virtual router boundaries, with the risk that traffic from
this virtual router will leak into another virtual router. Proceed with caution.

Interface   Belongs to different Router

outside (GigabitEthernet0/0)   ⌄   ▭   ⇄ Global

Protocol
● IPv4   ○ IPv6

Networks
+

🖵 vpn-pool

Gateway   Metric

Please select a gateway   ⌄   1

SLA Monitor Applicable only for IPv4 Protocol type

Please select an SLA Monitor   ⌄

c) Click **OK**.

What to do next

If there is overlap between the RA VPN address pool and the IP addresses in the custom virtual router, you
must also use static NAT rules on the IP addresses to enable proper routing. However, it is far easier to simply
change your RA VPN address pool so that there is no overlap.

How to Customize the AnyConnect Client Icon and Logo

You can customize the icon and logo for the AnyConnect Client app on Windows and Linux client machines.
The names of the icons are pre-defined, and there are specific limits to the file type and size for the images
you upload.

Although you can use any filename if you deploy your own executable to customize the GUI, this example assumes you are simply swapping icons and logos without deploying a fully-customized framework.

There are a number of images you can replace, and their file names differ based on platform. For complete information on customization options, file names, types, and sizes, please see the chapter on customizing and localizing the AnyConnect Client and installer in the *Cisco Secure Client Administrator Guide*. For example, the chapter for the 4.8 client is available at:

https://www.cisco.com/c/en/us/td/docs/security/vpn_client/anyconnect/anyconnect48/administration/guide/b_AnyConnect_Administrator_Guide_4-8/customize-localize-anyconnect.html

Before you begin

For the purposes of this example, we will replace the following images for Windows clients. Note that if your image is a different size than the maximum, the system will automatically resize it to the maximum, and stretch the image if necessary.

- app_logo.png

  This application logo image is the application icon, and it can have a maximum size of 128 x 128 pixels.

- company_logo.png

  This company logo image appears in the top-left corner of the tray flyout and Advanced dialogs. The maximum size is 97 x 58 pixels.

- company_logo_alt.png

  The alternative company logo image appears in the bottom-right corner of the About dialog box. The maximum size is 97 x 58 pixels.

To upload these files, you must place them on a server that the threat defense device can access. You can use a TFTP, FTP, HTTP, HTTPS, or SCP server. The URLs to get images from these files can include paths and uesrname/password, as required by your server setup. This example will use TFTP.

Procedure

Step 1 Upload the image files to each threat defense device that is acting as an RA VPN headend that should use the customized icons and logos.

a) Log into the device CLI using an SSH client.

b) In the CLI, enter the **system support diagnostic-cli** command to enter diagnostic CLI mode.

```
> system support diagnostic-cli
Attaching to Diagnostic CLI ... Press 'Ctrl+a then d' to detach.
Type help or '?' for a list of available commands.

ftdv1>
```

Note Read the message! You must press **Ctrl+a, then d**, to get out of the diagnostic CLI and back into the normal threat defense CLI mode.

c) Note the command prompt. The normal CLI uses > only, whereas the diagnostic CLI's user EXEC mode uses the hostname plus >. In this example, ftdv1>. You need to get into privileged EXEC mode, which uses # as the ending character, for example, ftdv1#. If your prompt already has #, skip this step. Otherwise, enter the enable command, and simply press Enter at the password prompt without entering a password.

```
ftdv1> enable
Password:
ftdv1#
```

d) Use the **copy** command to copy each file from the hosting server to the threat defense device's disk0. You can place them in a subdirectory, such as disk0:/anyconnect-images/. You can create a new folder using the **mkdir** command.

For example, if the TFTP server's IP address is 10.7.0.80, and you want to create a new directory, the commands would be similar to the following. Note that responses to the **copy** command are omitted after the first example.

```
ftdv1# mkdir disk0:anyconnect-images

Create directory filename [anyconnect-images]? yes

Created dir disk0:/anyconnect-images

ftdv1# copy /noconfirm tftp://10.7.0.80/app_logo.png
disk0:/anyconnect-images/app_logo.png

Accessing tftp://10.7.0.80/app_logo.png...!!!!!!
Writing file disk0:/anyconnect-images/app_logo.png...
!!!!!!
12288 bytes copied in 1.000 secs (12288 bytes/sec)

ftdv1# copy /noconfirm tftp://10.7.0.80/company_logo.png
disk0:/anyconnect-images/company_logo.png
ftdv1# copy /noconfirm tftp://10.7.0.80/company_logo_alt.png
disk0:/anyconnect-images/company_logo_alt.png
```

Step 2   Use the **import webvpn** command in the diagnostic CLI to instruct the AnyConnect Client to download these images when installing itself on client machines.

import webvpn AnyConnect-customization type resource platform win name *filename* **disk0:/***directoryname/filename*

This command is for Windows. For Linux, replace the **win** keyword with **linux** or **linux-64**, as appropriate for your clients.

For example, to import the files uploaded in the previous step, and assuming we are still in the diagnostic CLI:

```
ftdv1# import webvpn AnyConnect-customization type resource platform win
name app_logo.png disk0:/anyconnect-images/app_logo.png

ftdv1# import webvpn AnyConnect-customization type resource platform win
name company_logo.png disk0:/anyconnect-images/company_logo.png

ftdv1# import webvpn AnyConnect-customization type resource platform win
name company_logo_alt.png disk0:/anyconnect-images/company_logo_alt.png
```

Step 3   Verify the configuration:

- To verify the imported files, use the **show import webvpn AnyConnect-customization** command in the diagnostic CLI privileged EXEC mode.

- To verify that the images were downloaded to a client, they should appear when the user runs the client. You can also check the following folder on Windows clients, where %PROGRAMFILES% typically resolves to c:\Program Files.

  %PROGRAMFILES%\Cisco\Cisco AnyConnect Secure Mobility Client\res

What to do next

If you want to return to the default images, use the **revert webvpn** command (in the diagnostic CLI privileged EXEC mode) for each image you customized. The command is:

revert webvpn AnyConnect-customization type resource platform win name *filename*

As with **import webvpn**, replace **win** with **linux** or **linux-64** if you customized those client platforms, and issue the command separately for each image filename you imported. For example:

```
ftdv1# revert webvpn AnyConnect-customization type resource platform win
name app_logo.png

ftdv1# revert webvpn AnyConnect-customization type resource platform win
name company_logo.png

ftdv1# revert webvpn AnyConnect-customization type resource platform win
name company_logo_alt.png
```