
# Setting General VPN Parameters

The security appliance implementation of virtual private networking includes useful features that do not fit neatly into categories. This chapter describes some of these features. It includes the following sections:



## Configuring VPNs in Single, Routed Mode

VPNs work only in single, routed mode. VPN functionality is unavailable in configurations that include either security contexts, also referred to as multi-mode firewall, or Active/Active stateful failover.

The exception to this caveat is that you can configure and use one connection for administrative purposes to (not through) the security appliance in transparent mode.

## Configuring IPSec/SSL to Bypass ACLs

**Note** Unless clienltess (browser-mode) SSL VPN is specified, the term SSL VPN in this chapter refers to the SSL VPN client (AnyConnect 2.x or previous SVC 1.x).

To permit any packets that come from an IPSec/SSL tunnel without checking ACLs for the source and destination interfaces, enter the **sysopt connection permit-ipsec** command in global configuration mode.

You might want to bypass interface ACLs for IPSec/SSL traffic if you use a separate VPN concentrator behind the security appliance and want to maximize the security appliance performance. Typically, you create an ACL that permits IPSec/SSL packets using the **access-list** command and apply it to the source interface. Using an ACL is more secure because you can specify the exact traffic you want to allow through the security appliance.

The syntax is **sysopt connection permit-vpn**. The command has no keywords or arguments.

The following example enables IPSec/SSL traffic through the security appliance without checking ACLs:

```
hostname(config)# sysopt connection permit-vpn
```

# Permitting Intra-Interface Traffic (Hairpinning)

The security appliance includes a feature that lets a VPN client send IPSec-protected traffic to another VPN user by allowing such traffic in and out of the same interface. Also called “hairpinning”, this feature can be thought of as VPN spokes (clients) connecting through a VPN hub (security appliance).

In another application, this feature can redirect incoming VPN traffic back out through the same interface as unencrypted traffic. This would be useful, for example, to a VPN client that does not have split tunneling but needs to both access a VPN and browse the Web.

Figure 24-1 shows VPN Client 1 sending secure IPSec/SSL traffic to VPN Client 2 while also sending unencrypted traffic to a public Web server.

***Figure 24-1 VPN Client Using Intra-Interface feature for Hairpinning***

To configure this feature, use the **same-security-traffic** command in global configuration mode with its **intra-interface** argument.

The command syntax is **same-security-traffic permit** {**inter-interface** | **intra-interface**}.

The following example shows how to enable intra-interface traffic:

```
hostname(config)# same-security-traffic permit intra-interface
hostname(config)#
```

**Note** You use the **same-security-traffic** command, but with the **inter-interface** argument, to permit communication between interfaces that have the same security level. This feature is not specific to IPSec/SSL connections. For more information, see the “Configuring Interface Parameters” chapter of this guide.

To use hairpinning, you must apply the proper NAT rules to the security appliance interface, as discussed in the following section.

## NAT Considerations for Intra-Interface Traffic

For the security appliance to send unencrypted traffic back out through the interface, you must enable NAT for the interface so that publicly routable addresses replace your private IP addresses (unless you already use public IP addresses in your local IP address pool). The following example applies an interface PAT rule to traffic sourced from the client IP pool:

```
hostname(config)# ip local pool clientpool 192.168.0.10-192.168.0.100
hostname(config)# global (outside) 1 interface
hostname(config)# nat (outside) 1 192.168.0.0 255.255.255.0
```

When the security appliance sends encrypted VPN traffic back out this same interface, however, NAT is optional. The VPN-to-VPN hairpinning works with or without NAT. To apply NAT to all outgoing traffic, implement only the commands above. To exempt the VPN-to-VPN traffic from NAT, add commands (to the example above) that implement NAT exemption for VPN-to-VPN traffic, such as:

```
hostname(config)# access-list nonat permit ip 192.168.0.0 255.255.255.0 192.168.0.0
255.255.255.0
hostname(config)# nat (outside) 0 access-list nonat
```

For more information on NAT rules, see the "Applying NAT" chapter of this guide.

# Setting Maximum Active IPSec/SSL VPN Sessions

To limit VPN sessions to a lower value than the security appliance allows, enter the **vpn-sessiondb max-session-limit** command in global configuration mode.

- This command applies to all types of VPN sessions, including WebVPN.
- This limit affects the calculated load percentage for VPN Load Balancing.

The syntax is **vpn-sessiondb max-session-limit {***session-limit***}.**

The following example shows how to set a maximum VPN session limit of 450:

```
hostname (config)# vpn-sessiondb max-session-limit 450
hostname (config)#
```

To set both SSL VPN client and clientless max sessions enter the **vpn-sessiondb max-webvpn-session-limit {session-limit}** command in global configuration mode.

# Configuring Client Update

The client update feature lets administrators at a central location automatically notify VPN client users when it is time to update the VPN client software and the VPN 3002 hardware client image.

To configure client update, enter the **client-update** command in tunnel-group ipsec-attributes configuration mode. If the client is already running a software version on the list of revision numbers, it does not need to update its software. If the client is not running a software version on the list, it should update. You can specify up to 4 client update entries.

The command syntax follows:

**client-update type** *type* {**url** *url-string*} {**rev-nums** *rev-nums*}

**no client-update** [*type*]

**Syntax Description**

| | |
|---|---|
| **rev-nums** *rev-nums* | Specifies the software or firmware images for this client. Enter up to 4, separated by commas. |
| *type* | Specifies the operating systems to notify of a client update. The list of operating systems comprises the following:<br>• Windows: all windows-based platforms<br>• WIN9X: Windows 95, Windows 98, and Windows ME platforms<br>• WinNT: Windows NT 4.0, Windows 2000, and Windows XP platforms<br>• vpn3002: VPN 3002 hardware client |
| **url** *url-string* | Specifies the URL for the software/firmware image. This URL must point to a file appropriate for the client. |

The following example configures client update parameters for the remote-access tunnel-group called remotegrp. It designates the revision number 4.6.1 and the URL for retrieving the update, which is https://support/updates.

```
hostname(config)# tunnel-group remotegrp type ipsec_ra
hostname(config)# tunnel-group remotegrp ipsec-attributes
hostname(config-ipsec)# client-update type windows url https://support/updates/ rev-nums
4.6.1
hostname(config-ipsec)#
```

# General Considerations

The following section are questions that you should consider as you set up VPN load balancing. This section is formatted in a question and answer format.

**Q.** Does the ASA consider IP Pool exhaustion as part of its VPN load balancing mechanism?

**A.** No. If the VPN remote access session is directed to the least_loaded unit, which has exhausted its IP pools, then the session will fail to establish. The algorithm is based on Load, and is computed as an integer percentage (# of active/max sessions) supplied by each secondary cluster member.

**Q.** There are four ASAs in a cluster using a VIP via the ASA's own internal load balancing. Can we use the same group-url on all four members of the cluster w/o issues? And from a DNS perspective, can we just create an A record pointing at the VIP; or do we have to do something else?

**A.** It appears that on each cluster member that we cannot use **group-url https://vpn.rob.com/eng enable**. Instead we have to use the real IP address (not the VIP) of the ASA. If we use the URL and/or the VIP IP, Anyconnect is unable to connect.

For example: I have a 2 ASA cluster setup and it turns out I have both the FQDN and IP address for group-url. When trying to access the cluster the ASA will use the IP address of the machines in the cluster. I removed the FQDN group-url and it stopped working.

ASA1 with **`group-url group-url`** **https://10.94.147.93/BasicGroup**

and

ASA2 with **group-url group-url https://10.94.147.92/BasicGroup**

I can then access the cluster and BasicGroup using the cluster name and group-url: **cvc-asa.cisco.com/BasicGroup**.

**Q.** When we implement VPN load balancing, shouldn't the address pools for AnyConnect clients (or IPSec/SSL clients) on different ASA's participating in cluster be different?

**A.** Correct. If using address pools, they must be unique per device

**Q.** Can load load balancing and failover be combined?

**A.** Yes.

You can also have a configuration that combines both load balancing and failover. For example, the client connects to the IP address of the cluster and is redirected to the least-loaded ASA in the cluster. If that ASA goes down, the standby unit takes over immediately, and there is no impact to the client's tunnel.

**Note** Only the Active units participate in load balancing. Should the Active unit of a failover pair go down, then its Standby mate would become active and then join the Load Balancing cluster mechanism to distribute the VPN session load.

**Q.** If we have WebVPN enabled on multiple interfaces, is it possible to have VPN load balancing implemented for both of them?

**A.** You can only define one interface to participate in the cluster as the 'public' interface. The idea is to balance the CPU loads. Multiple interfaces still converge on the same cpu, so the concept of load-balancing on interfaces doesn't have any value. At this time there is no plans to support this.

**Q.** By default, when a cluster master redirects an incoming connection, it redirects it by IP address so it would show up at the ASA with an IP address rather than FQDN.

**A.** The options are to add a group-url for the local ASA https://ip_address/group-url or add the following command to the ASA to allow them to forward by FQDN rather than IP address:

```
(config)# vpn load-balancing
(config-load-balancing)# redirect-fqdn enable
```

**Q.** When trying to implement SSL licensing and failover, consider the following deployment:

Two ASA5520's , each with 100-user SSL VPN licenses, in a load balancing cluster.

Does the maximum total number of users allow 200 simultaneous users or only a maximum of 100? If you add a third device later with 100 users, can you now support 300 simultaneous users?

**A.** With VPN load balancing, all devices are active. This allows you to take the licensed amount per device, and add them together to determine the maximum number of users that your cluster can support. For this example, 200 sessions for twoASAs and 300 sessions for three ASAs, respectively.

**Q.** Is there a limit on the number of appliances that can participate in load balancing clustering?

**A.** There is no hard limit. Engineering tests up to ten nodes in a cluster. Additional nodes may work, but we do not officially support that topology.

**Q.** How does load balancing work for the adapative security appliance?

**A.** Basically, load balancing works like this:

- The phase 1 negotiation is done on the virtual master.
- An IKE redirect packet with the IP of a slave device was sent by the virtual master to the client.
- The client will start a new phase 1 and 2 negotiation on the slave device just like a standalone vpn connection.

For remote access, there is no need to setup any route manually. The situation is the same for a standalone as well as a load balancing redirected tunnel. Basically, a host route of the assigned IP address pointing to the public ip of the client device is installed on the inside interface of the ASA. "show route" will display the host route. Because of this reverse route, the inside interface of the ASA will respond to the ARP request of the client's assigned IP and hence, can return traffic from a server on the inside network to the client through the tunnel.

Load balancing works for IPsec Hardware Clients (VPN3002, PIX501, ASA5505)client/PAT mode and Network Extension Mode(NEM) as well.