TechNote:
Advantages and applications of interface/route based VPN tunnels

Introduction

With ScreenOS 3.1.0, many new features have been introduced to provide significant flexibility for NetScreen's security devices. Two of the most notable features introduced are definable and configurable security zones (beyond Trust, Untrust, and DMZ) and interfaces that now support near-generic features, including the ability to be bound to any security zone. One new feature that has received little attention is the functionality added for VPN deployments, in particular the ability to easily implement interface-based VPN tunnels in addition to policy based VPN tunnels.

This application note will define interface-based VPN tunnels (sometime referred to as route-based VPN tunnels) and describe some of the applications that interfaces-based VPN tunnels are best suited for and provide sample configurations for these.

This information is not designed to replace the Concepts & Examples guide, but rather to provide a comparison between the applications of similar features.

Definitions

Interface-based VPN tunnels are a concept that may be fairly new to many users. These VPN tunnels require a logical interface to be created called a tunnel interface. The tunnel interface can have its own IP address or share the IP address of the interface that it is associated with. This tunnel interface is then placed in a tunnel zone. To then route traffic to the VPN tunnel, there must be a policy that allows traffic to the zone, and then a route for the traffic to the tunnel interface. So the policy is a standard firewall policy, or group of many policies that allow or deny traffic. They do not specify the VPN tunnel to use. Instead, the VPN tunnel is selected based on the routing table that determines the interface to use. When using interface-based VPN tunnels, the administrator can define the proxy ID information to further restrict access on the VPN tunnel.

In contrast, *Policy-based VPN tunnels* are the more familiar of the two types of IPSec VPN tunnels. Policy-based VPN tunnels continue to be supported on NetScreen devices. Policy-based VPN tunnels are defined by creating a gateway or remote user, then the parameters used to identify the gateway, and finally the policy that defines the parameters (source, destination, and service) that will then invoke the VPN tunnel to encrypt the data. The final step is the most important in this process as the VPN tunnel is defined by the policy that uses it. The policy will be used as part of the IKE negotiations (to determine the proxy IDs unless policy checking is turned off). Policy-based VPN tunnels are very powerful when users or sites need to be identified based on the VPN parameters, such as information in the digital certificate. Because the VPN tunnel is part of the policy, the IPSec authentication and the firewall policy can be closely intertwined.

Interface-based VPN tunnels

To better understand the key components and interaction with interface-based VPN tunnels, the following diagram logically steps through the process involved (assuming that an existing session is not already in place).

1) Packet arrives at a interface ethernet2/1 in the Trusted zone
2) Virtual Router is queried to determine the destination security zone for policy look-up
3) Policy engine is queries to find a matching policy. If found, instructions are returned, such as permit.
4) Session is established and packet is sent to destination zone
5) Virtual router is queried to determine what interface to use
6) Packet is passed to the proper interface
7) And sent out to the next network hop

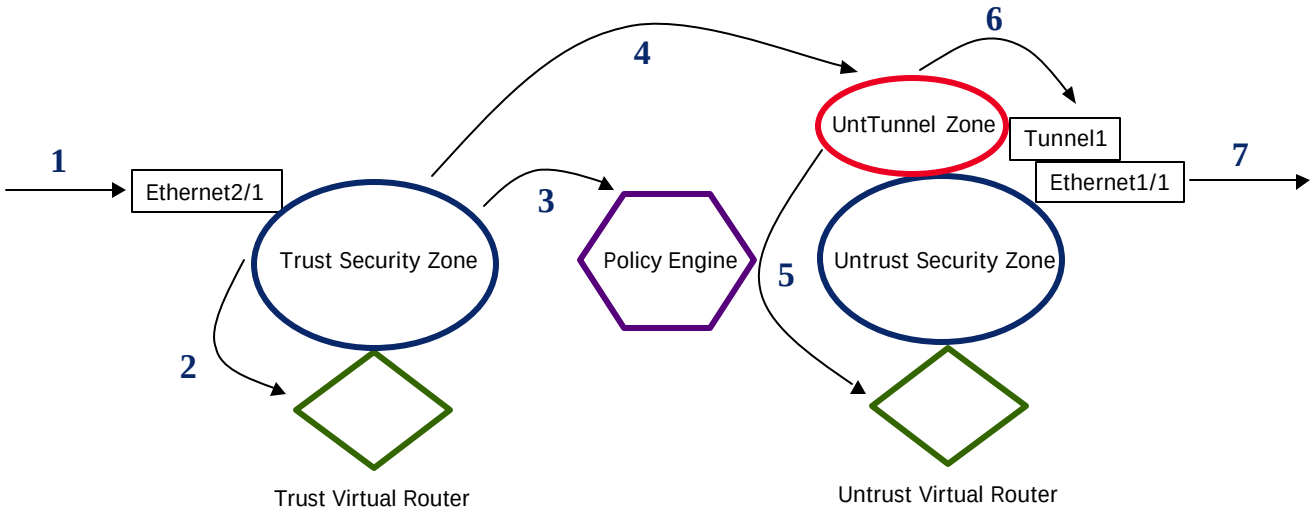

Notice that the policy engine is searched without regard to the VPN tunnel. The firewall policy will determine whether or not the packet will be allowed to progress to the destination security zone (in this case the UntTunnel security zone, a tunnel zone bound to the Untrust Zone). Based on the routing, the interface is selected, which in this case is Tunnel1, bound to the Untrust interface, and encrypted to the destination.

Key applications of interface-based VPN tunnels
Policy-based VPN tunnels are potentially the easier of the two types of VPN tunnels to configure (primarily because they are more familiar to most people). However, interface-based VPN tunnels can be used for most situations that a policy-based tunnel can be and they can also be used to simplify many VPN networks or to solve new application environments.

Application 1) Hub & Spoke VPN networks often have been complicated to set up since networks that are located at the remote sites were actually trusted networks when setting policies. In addition, all VPN tunnels require policies to direct traffic back out to the other spokes in the network, creating large and potentially confusing policy tables that shouldn't be necessary.

With interface-based VPN tunnels, the route associated with the tunnel zone will determine where to send the traffic. If the routing table determines that the traffic is destined for another remote site, it simply routes it to the correct tunnel interface to send to the remote site. This process skips the policy lookup, which often isn't needed in hub & spoke connections, and greatly simplifies the policy tables. Very important to note is that since the policy table is not queried, the VPN hub just acts as a central point, but not an enforcement point.

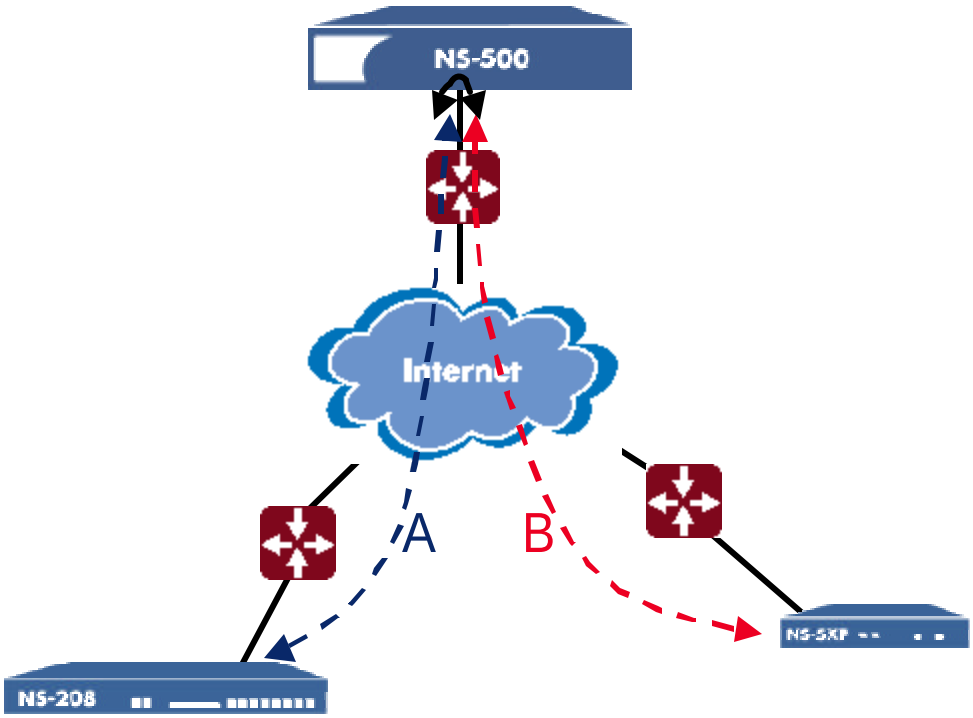

Application 2) Site-to-site tunnels, with a need for many policies. Many times, when VPNs are used to secure communications between two larger offices, one policy is not enough to appropriately filter the traffic through the VPN tunnel. With policy-based VPN tunnels, an administrator creates many policies using the same VPN configuration to provide the level of policy detail that is required. This in turn creates many SAs, which can lead to a capacity problem as each SA is regenerating keys and is essentially treated as a separate VPN tunnel with the associated overhead.

With interface-based VPN tunnels, a single VPN tunnel is set up between two tunnel interfaces. Firewall policies are instead used to determine what traffic is able to get to the tunnel zone being used. This allows for many policies to filter traffic before the VPN tunnel, and then the VPN tunnel just encrypts traffic that is able to reach it. Basically, the VPN tunnel becomes a means of reaching the remote sites, and firewall policies control access separately.

Application 3) Redundant VPN tunnels are desired in VPN networks where business-critical traffic relies on the connection, such as applications, e-mail servers, etc. Many networks have redundant ISPs to ensure the network's availability. This infrastructure can be leveraged with interface-based VPN tunnels by using two interfaces in the same untrust zone, each with a different VPN tunnel associated with it for redundancy. If one ISP fails, there could be a backup that uses a different ISP connection. This typically would be accomplished with two interfaces being assigned to the same external (Untrust) security zone. Routes would then determine the appropriate interface to use. In addition, remote offices would set up redundant VPN gateways so that they can fail over as well in the case of an ISP failure.

Note: Large numbers of dial-up users accessing the network with similar parameters and access rights are a good case for policy-based VPN tunnels. In this case, a large number of users can be grouped together, then grouped together, and finally a policy is created that is the same for all users as they come into the network. This is an example where interface-based VPNs do not provide an advantage, and policy-based VPN tunnels are easily deployed to resolve these topology requirements.

Summary

Interface-based VPN tunnels are a new concept that provide a significant advantage in certain VPN applications, as outlined above. For more information on specific configuration differences, please refer to NetScreen's Concepts & Examples Guide.