

## Re: Pix 515 VLAN NAT0 issues

---

*Source:* <http://newsgroups.derkeiler.com/Archive/Comp/comp.dcom.sys.cisco/2006-03/msg00959.html>

---

- *From:* [roberson@xxxxxxxxxxxxx](mailto:roberson@xxxxxxxxxxxxx) (Walter Roberson)
  - *Date:* Thu, 16 Mar 2006 20:26:02 GMT
- 

In article <1142535749.946683.143370@xx>, tartar813 <rtartar@xxxxxxxxxx> wrote:

Where is the conduit conversion tool?

<http://www.cisco.com/cgi-bin/tablebuild.pl/pix>

and log in to your account, then scroll down the list until you find occ-121 about 2/3 of the way down.

```
object-group network REGGIE_STATIC_HOSTS
network-object host 72.29.91.82
network-object host 72.29.91.83
network-object host 72.29.91.84
network-object host 72.29.91.85
network-object host 72.29.91.86
network-object host 72.29.91.87
network-object host 72.29.91.88
access-list reggie_out_acl permit ip object-group REGGIE_STATIC_HOSTS
any
nat (reggie) 0 access-list reggie_out_acl
```

Let me make sure I get it, This will not NAT all of the items going out from the REGGIE\_STATIC\_HOSTS network object group?

Right. Anything sourced "within" the reggie segment that matches that ACL will be exempt from NAT.

Does this automatically setup the inbound translations also?

Supressing some unimportant semantic quibbles, Yes, exactly. Any connection heading into a lower-security interface that matches the "reverse" of the ACL (i.e, exchange source and destination fields)

## Re: Pix 515 VLAN NAT0 issues

will be permitted inward, provided that the access-group on that lower interface permits that flow. It is a form of "static" for that purpose.

There is, though, the side effect that proxy arp will not be enabled for the IPs (not unless there is a regular static for that IP), so your WAN router will have to route those IPs to the outside IP of the PIX. This is usually not a problem unless you happen to have real hosts on the outside segment.

Thank you, I really appreciate this, I feel like an idiot since I've been using the conduits and stuff for so long.

Even the TAC ends up scratching their head over bidirectional policy NAT. Some stuff just isn't well documented.

Some ACL and translation fundamentals:

Each ACL should be written in terms of the IPs that would be in the packet at the time the PIX receives the packet. e.g., an ACL applied to an inside interface would have the internal IPs as the source and the outside IPs \*as known to the inside\* as the destinations.

Translation takes place after the interface controls have decided to accept the packet, based upon the ACL applied to the interface (or upon the default flow rules if there is no ACL.) But that's the rule for when the translation is actually performed: before the ACL is even looked at, the PIX checks to see that there a translation exists. Thus if a new connection attempt hits your outside interface and is addressed to a public IP that you do not have a "static" or "nat 0 access-list" for, then the packet will be dropped with a log entry about "no translation group" and only if there is a translation can you go on to "denied by access-list". {It wasn't that way before 6.2, and they might have modified this by now, as I griped about this.} The modification of packet content happens after the packet has been accepted as having a translation and satisfying the security policies.

The default rules, if you have no ACL applied to an interface, are that traffic to lower-security is allowed and to higher security is not allowed. If you do have an ACL, then that rule does not apply at all, and instead the rule becomes "anything which is not permitted by the ACL is not allowed."

An important difference you will hit is that "conduit" applies to all interfaces, but the access-group command applies an ACL

## Re: Pix 515 VLAN NAT0 issues

only to one interface. So before if you had a conduit that permitted traffic to something in your highest security zone, then you will need an ACL for each of the lower security zones if you want them to be able to reach that higher security zone.

Only one ACL is permitted "in" per interface. PIX 7.x adds ACLs "out" an interface, and modifies to "one per direction".

Never try to use the same ACL for two purposes. If you have two controls mention the same ACL name/number then you will likely have odd problems.

Translation to lower security interfaces normally changes the source IP, and translation to higher security interfaces normally changes the destination IP. [PIX 6.2 and later allow changing this.]

An ACL applied to an interface should refer to the private IP of a host on a lower security security interface, but to the public IP of a host on a higher security interface. Of course if you have used nat 0 access-list or static'd IPs to themselves between a pair of interfaces, then the public and private IP would be the same for that transaction.

Only one "nat 0 access-list" is permitted per interface, and it applies to traffic going to lower security interfaces. Indefinite numbers of "nat 0" (without access-list) are permitted per interface, and again apply to towards all lower security interfaces. "static" and all other "nat" commands work between pairs of interfaces, so the IP of an inside host as known to dmz1 could be different than the IP of the same host as known to dmz2.

Access-lists mentioned in crypto map (VPN) "match address" clauses should be written from the perspective of packets going out the interface that the crypto map is applied to. But unlike the other cases, the "match address" ACLs must be written in terms of what would be in the packet \*after\* translation (towards the outside). For incoming VPN packets, the "match address" ACL will automatically be read "in reverse" [like for the nat 0 access-list case], and the addresses used to check will be the ones after decapsulation but before any translation.

An incoming VPN packet will be decapsulated, and the inner packet first checked against the {implicitly reversed} appropriate "match address" ACL. After that, the inner packet will be checked against the ACL (or default policy) for the interface it was received on, -unless- "sysopt connection permit-ipsec" or similar has been turned on: If you use those commands, then all VPN packets that manage to make it to you will be permitted to go to any destination (except on the -same- interface) without any checking of access policies.

Similarly, an outgoing VPN packet will be checked first against the security policy of the interface it was received on, \*unless\* "sysopt connection permit-" is in effect and the packet would go out over the VPN -- those packets will go through even if the security policy says to block them. After the outgoing VPN packet is accepted by the interface, it undergoes translation, and the -translated- packet will be compared against the "match address" ACLs for dispatching.