Troubleshooting Virtual Private Networks (VPN's)
Document Version: 1b
ScreenOS Version: 5.0.0r4
NetScreen Products: Applies to all NetScreen Firewall Devices

Design:

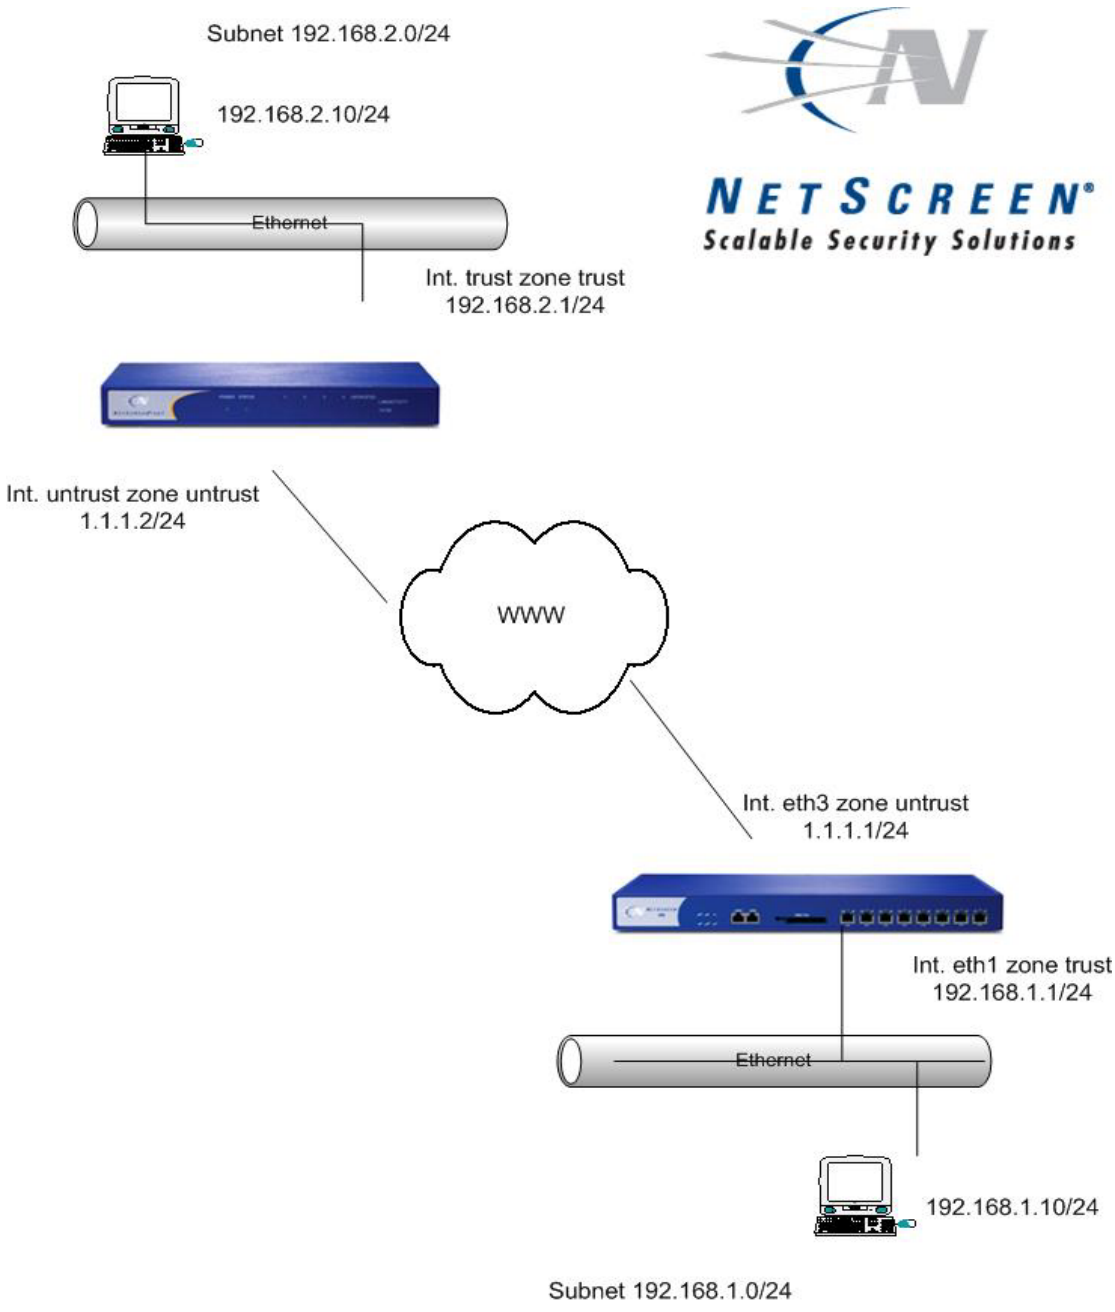

VPN Troubleshooting flow diagram

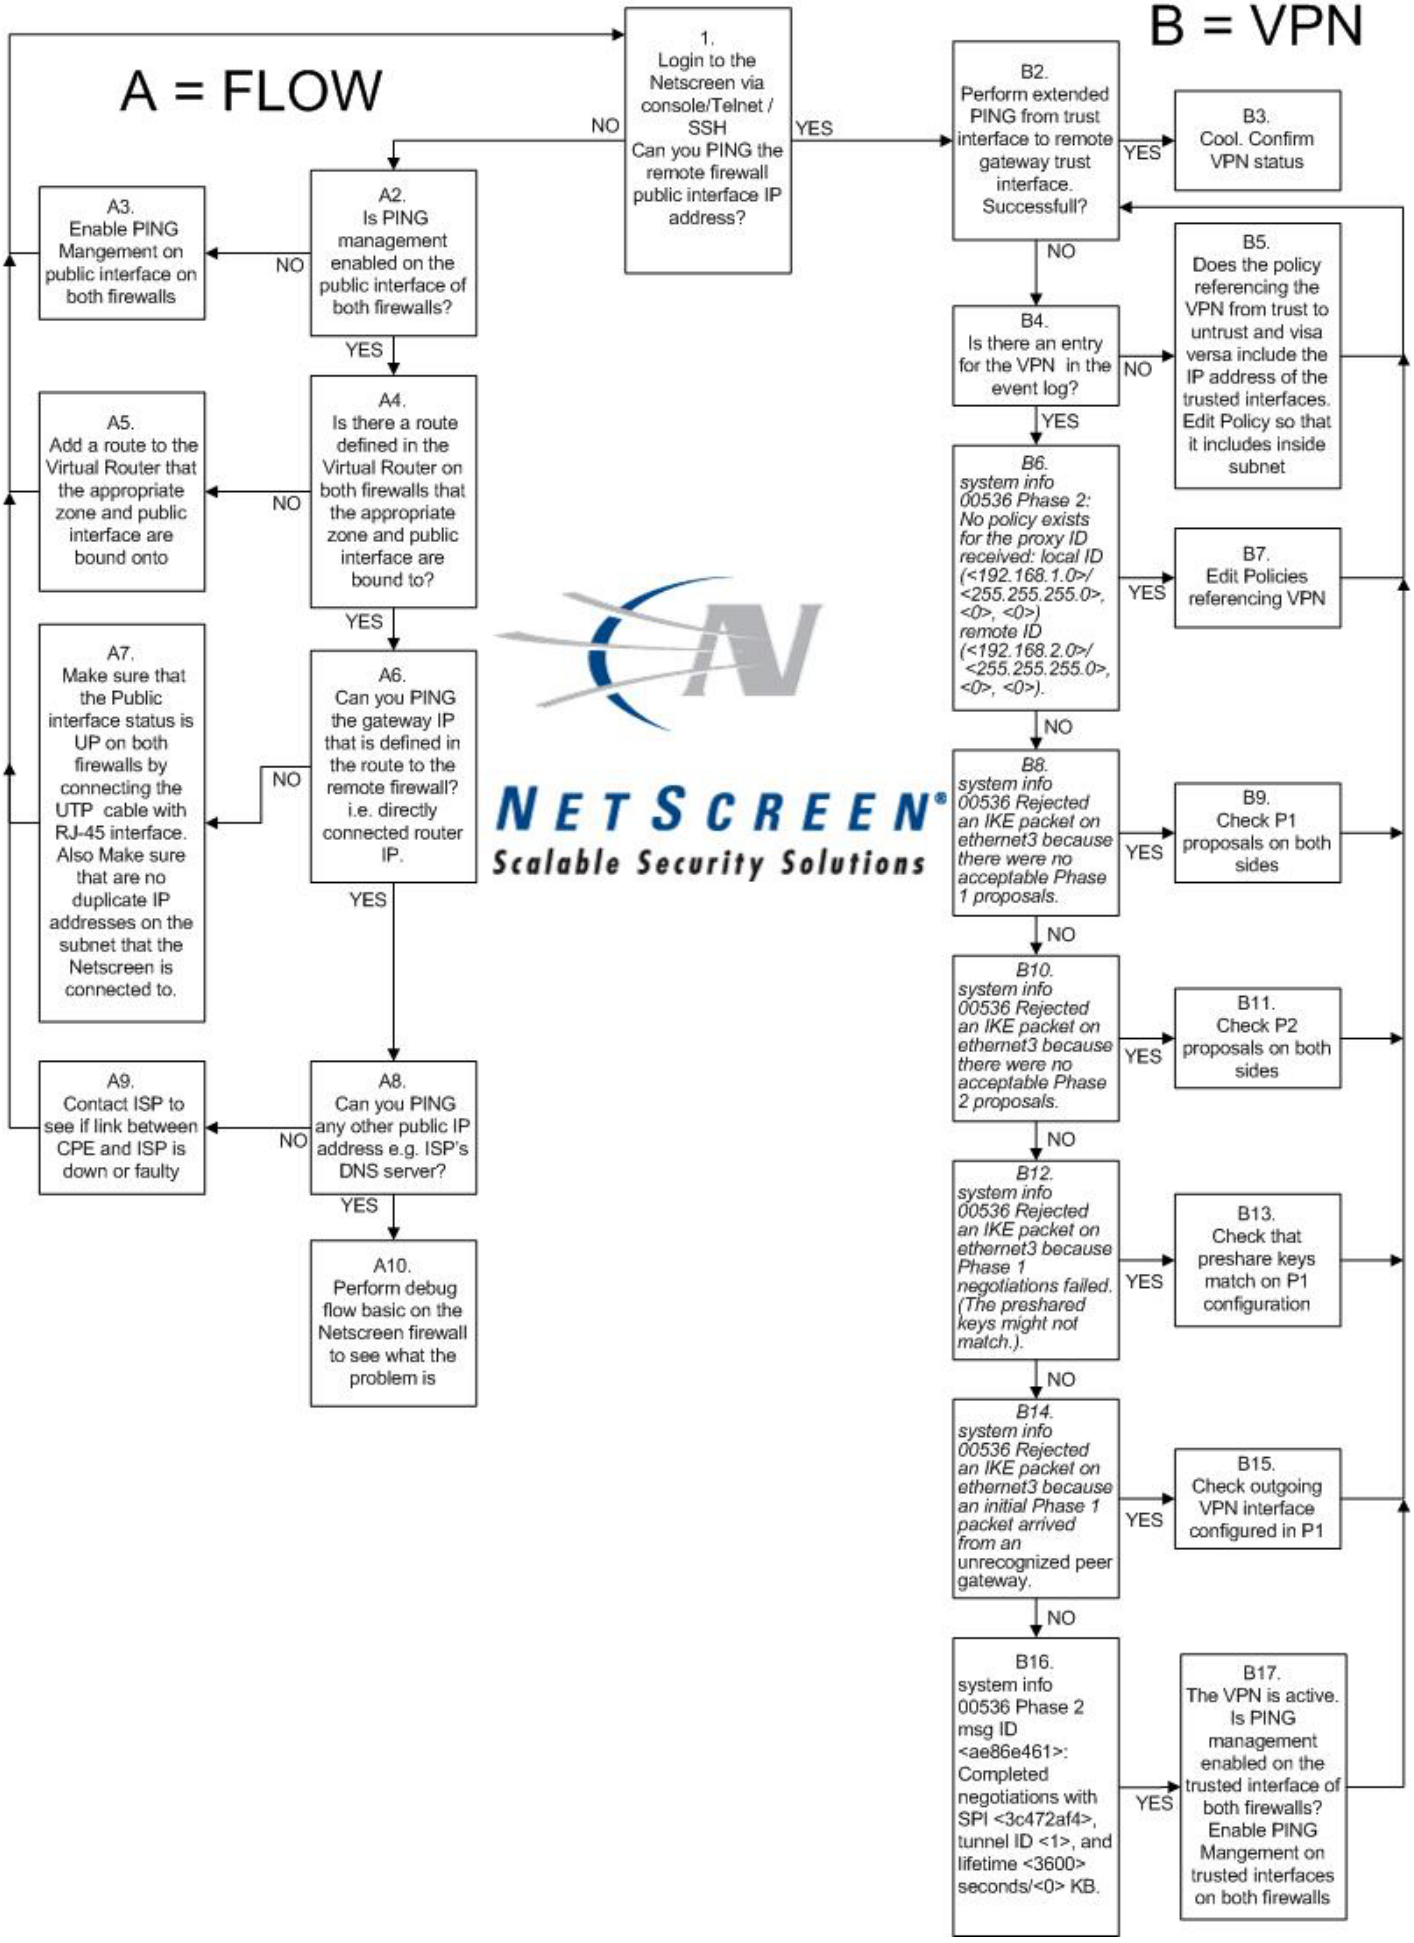

The idea of this document is to give an overview of the process involved in troubleshooting VPN's using the log entries in the event log. This can be done by turning on the debug for the packet flow, as well as for phase1 and phase2 internet key exchange (IKE) between the NetScreen firewalls. Using these methods it is possible to track the incoming and outgoing packet flows in detail.

Concepts:

- Ingress Interface
- Egress Interface
- Debug buffer
- Flow Filters
- Debug flow basic
- Session setup criteria and Packet flow process

Ingress Interface – 1st incoming interface a packet enters the NetScreen firewall

Egress Interface – exiting interface once a packet has been processed and sent to destination

Debug buffer - The NetScreen firewall by default will never send a debug output directly to the console or telnet/ssh session. The NetScreen firewall will send the debug output to a debug buffer. This greatly saves on CPU cycles. If possible always check the event log for log entries before using the debug. This will speed up the troubleshooting process and also does not place extra load on the CPU.
WARNING: If you send a debug output directly to the console, it is likely that you will lose connectivity to the NetScreen via SSH/Telnet during heavy traffic loads.

On the NetScreen Console, Telnet session or SSH session:

get console (make sure that the debug output will be sent to the debug buffer)

```
ns208-> get console
Console timeout: 0(minute), Page size: 0/0, debug: buffer
privilege 250, config was changed and not saved!
ID State  Duration Task Type   Host
 0 Logout       0 8480 Local
 1 Logout       0 3664 Local
 2 Login     6368 1968 Local
```

get db stream (allows you to view the debug buffer)
clear db (allows you to clear the debug buffer)

Flow Filters - Set a Flow filter to only capture the flow process on certain criteria on heavily loaded networks. This can be preformed from the console, telnet or SSH session. Flow filters are used in conjunction with the 'debug flow basic':

```
ns208-> set ffilter ?
<return>
dst-ip            flow filter dst ip
dst-port          flow filter dst port
ip-proto          flow filter ip proto
src-ip            flow filter src ip
src-port          flow filter src port
```

ns208-> set ffilter src-ip 192.168.1.10
filter added
ns208-> get ff
Flow filter based on:
id:0 src ip 192.168.1.10

The flow filter works on the 'AND' and 'OR' process. Multiple filter ID's will be matched to the 'OR' and if you have src-ip and dst-ip specified on the same ID it will be 'AND' matched.

ns208-> set ffilter src-ip 192.168.1.11
filter added
ns208-> get ff
Flow filter based on:
id:0 src ip 192.168.1.10
id:1 src ip 192.168.1.11

ns208-> get ff (Packet flow will be captured from src-ip 192.168.1.10 OR src-ip 192.168.1.11)
Flow filter based on:
id:0 src ip 192.168.1.10
id:1 src ip 192.168.1.11

ns208-> unset ff 1
filter 1 removed
ns208-> get ff
Flow filter based on:
id:0 src ip 192.168.1.10

ns208-> set ff src-ip 192.168.1.10 dst-ip 192.168.2.10 (Packet flow will be captured from only from src-ip 192.168.1.10 AND dst-ip 192.168.2.10)
ns208-> get ff
Flow filter based on:
id:0 src ip 192.168.1.10 dst ip 192.168.2.10

debug flow basic – debugs the flow process of a packet through the NetScreen i.e. route lookup, policy match, NAT etc. This can be preformed from the console, telnet or SSH session.

NB – the debug is performed on packets arriving on the Ingress interface only. This means that if a packet arrives on the NS-208 eth1 interface, then you will see that flow process for the packet and in this process see what decision the NetScreen has made onto which destination interface to send the packet. You will not actually see the packet leave the egress interface with the debug flow basic. You will however see the return packet when it enters the ingress interface again.
Turn the debug flow basic off by pushing <ESC> on the keyboard

get debug (confirm if any existing debug is running on the NetScreen)
undebug all (turns off all debugs)

Session setup criteria and packet flow process– It is important to understand the criteria the NetScreen will check before it creates a session for the flow:

1. Is there an existing session? Yes, perform route lookup and send packet to upstream mac address of route entry(In ScreenOS 4.x, once a session is matched, the packet will be forwarded to the upstream mac address of the route. There will be no additional route lookup. This is a 5.x feature only, as the upstream route might change or become inactive, and the NetScreen can dynamically change the session entry to contain a new mac address for the new active route). If there is no session then proceed to step 2:
2. Is TCP-SYN-CHECK enabled? Yes, TCP Packet will be dropped if does not have SYN flag set. If TCP-SYN-CHECK is disabled, the NetScreen will process and start the session creation process for any packet e.g. SYN, SYN/ACK, PUSH/ACK except a packet with the RST or FIN bit set.
3. Is there a route to the destination IP address? No – drop packet. If yes, then perform route lookup and decide outgoing interface. Proceed step 4:
4. Once outgoing interface is found via a route lookup, do a policy search from ingress interface zone to egress interface zone. Policy match? No, drop packet. If yes, proceed to step 5:
5. Does policy or interface have NAT configured. No, don't perform NAT. Yes, Perform NAT.
6. Create Policy with ID x
7. If no MAC entry for IP address for the gateway, ARP out destination interface. If MAC address already in ARP table, packet sent out interface.

Successful session example:

ns208-> get session
alloc 176/max 128000, alloc failed 0, di alloc failed 0
*id 45/s**,vsys 0,flag 00000040/0000/00,policy 1,time 6, dip 0*
0(8801):192.168.1.10/4000->192.168.2.10/1024,1,0010db103040,2,vlan 0,tun 0,vsd 0,route 0
6(2800):192.168.1.10/4000<-192.168.2.10/1024,1,000000000000,3,vlan 0,tun 40000001,vsd 0,route 5

id = session ID (also seen in the debug)
Vsys = Virtual System (Available in the NS500, NS2000, and NS5000)
Flag = (Proprietary information to NetScreen)
Policy = Policy Matched
Time = Session timeout displayed in Ticks (1 Tick = 10 seconds)
Dip = If a DIP Pool is used for NAT
0(8801) = interface and session token (0 = eth1 and 6 = eth 3 can be seen with 'get system')
192.168.1.10/4000 = Source IP and Source port
192.168.2.10/1024,1 = Destination IP and destination port, followed by protocol number (ICMP = 1)
0010db103040 = upstream router MAC address
vlan = if a Vlan and tagged interface is used
tun = the VPN tunnel used
vsd = Virtual Security Device group in the case of using NSRP
route = route matched in the routing table

Section A: Packet Flow

A2 - Is PING management enabled on the public interface of both firewalls?

ns208-> get interface

A - Active, I - Inactive, U - Up, D - Down, R - Ready

Interfaces in vsys Root:

Name	IP Address	Zone	MAC	VLAN	State	VSD
eth1	192.168.1.1/24	Trust	0010.db19.a7d0	-	U	-
eth2	0.0.0.0/0	DMZ	0010.db19.a7d5	-	D	-
eth3	1.1.1.1/24	Untrust	0010.db19.a7d6	-	U	-
eth4	0.0.0.0/0	Null	0010.db19.a7d7	-	D	-
eth5	0.0.0.0/0	Null	0010.db19.a7d8	-	D	-
eth6	0.0.0.0/0	Null	0010.db19.a7d9	-	D	-
eth7	0.0.0.0/0	Null	0010.db19.a7da	-	D	-
eth8	0.0.0.0/0	HA	0010.db19.a7db	-	U	-
vlan1	0.0.0.0/0	VLAN	0010.db19.a7df	1	D	-

ns208-> get interface eth3
Interface ethernet3:
number 6, if_info 12336, if_index 0, mode route
link up, phy-link up/half-duplex
vsys Root, zone Untrust, vr trust-vr
dhcp client disabled
PPPoE disabled
**ip 1.1.1.1/24 mac 0010.db19.a7d6*
**manage ip 1.1.1.1, mac 0010.db19.a7d6*
route-deny disable
ping disabled, telnet disabled, SSH disabled, SNMP disabled
web disabled, ident-reset disabled, SSL disabled
webauth disabled, webauth-ip 0.0.0.0
OSPF disabled BGP disabled RIP disabled
bandwidth: physical 10000kbps, configured 0kbps, current 0kbps
total configured gbw 0kbps, total allocated gbw 0kbps
DHCP-Relay disabled
DHCP-server disabled

A3 - Enable PING Management on public interface on both firewalls

ns208-> set interface eth3 manage ping

ns208-> get interface eth3
Interface ethernet3:
number 6, if_info 12336, if_index 0, mode route
link up, phy-link up/half-duplex
vsys Root, zone Untrust, vr trust-vr
dhcp client disabled
PPPoE disabled
**ip 1.1.1.1/24 mac 0010.db19.a7d6*
**manage ip 1.1.1.1, mac 0010.db19.a7d6*
route-deny disable
ping enabled, telnet disabled, SSH disabled, SNMP disabled
web disabled, ident-reset disabled, SSL disabled
webauth disabled, webauth-ip 0.0.0.0
OSPF disabled BGP disabled RIP disabled
bandwidth: physical 10000kbps, configured 0kbps, current 0kbps
total configured gbw 0kbps, total allocated gbw 0kbps
DHCP-Relay disabled
DHCP-server disabled

A4 - Is there a route defined in the Virtual Router on both firewalls that the appropriate zone and public interface are bound to?

Look what zone the interface is Bound to:

ns208-> get interface

A - Active, I - Inactive, U - Up, D - Down, R - Ready

Interfaces in vsys Root:

Name	IP Address	Zone	MAC	VLAN	State	VSD
eth1	192.168.1.1/24	Trust	0010.db19.a7d0	-	U	-
eth2	0.0.0.0/0	DMZ	0010.db19.a7d5	-	D	-
eth3	1.1.1.1/24	Untrust	0010.db19.a7d6	-	U	-
eth4	0.0.0.0/0	Null	0010.db19.a7d7	-	D	-
eth5	0.0.0.0/0	Null	0010.db19.a7d8	-	D	-
eth6	0.0.0.0/0	Null	0010.db19.a7d9	-	D	-
eth7	0.0.0.0/0	Null	0010.db19.a7da	-	D	-
eth8	0.0.0.0/0	HA	0010.db19.a7db	-	U	-
vlan1	0.0.0.0/0	VLAN	0010.db19.a7df	1	D	-

Look what VR the zone is bound to:

```
ns208-> get zone
Total 13 zones created in vsys Root - 7 are policy configurable.
Total policy configurable zones for Root is 7.
--------------------------------------------------------------------
 ID Name          Type       Attr      VR          Default-IF   VSYS
  0 Null          Null       Shared    untrust-vr  hidden       Root
  1 Untrust       Sec(L3)    Shared    trust-vr    ethernet3    Root
  2 Trust         Sec(L3)              trust-vr    ethernet1    Root
  3 DMZ           Sec(L3)              trust-vr    ethernet2    Root
  4 Self          Func                 trust-vr    self         Root
  5 MGT           Func                 trust-vr    null         Root
  6 HA            Func                 trust-vr    ethernet8    Root
 10 Global        Sec(L3)              trust-vr    null         Root
 11 V1-Untrust    Sec(L2)              trust-vr    v1-untrust   Root
 12 V1-Trust      Sec(L2)              trust-vr    v1-trust     Root
 13 V1-DMZ        Sec(L2)              trust-vr    v1-dmz       Root
 14 VLAN          Func                 trust-vr    vlan1        Root
 16 Untrust-Tun   Tun                  trust-vr    hidden.1     Root
--------------------------------------------------------------------
```

Look what routes you have in the Trust-vr:

```
ns208-> get route
untrust-vr (0 entries)
------------------------------------------------------------------------------
C - Connected, S - Static, A - Auto-Exported, I - Imported, R - RIP
iB - IBGP, eB - EBGP, O - OSPF, E1 - OSPF external type 1
E2 - OSPF external type 2
trust-vr (2 entries)
------------------------------------------------------------------------------
   ID      IP-Prefix       Interface      Gateway   P Pref   Mtr    Vsys
------------------------------------------------------------------------------
*  2    192.168.1.0/24      eth1         0.0.0.0     C   0      0      Root
*  3      1.1.1.0/24        eth3         0.0.0.0     C   0      0      Root
```

Unless the peer VPN gateway was directly connected to same subnet as this device, the NetScreen will not know where to send an outgoing packet destined for a subnet that is not the same as its own. (Default gateway). Its good practice on a NetScreen device to always specify a default route in the routing table in case the NetScreen does not know where to route a packet. This can be true even for directly connected devices on the subnet.

NB – A quick way to see if you have a route to the destination IP:

```
ns208-> get route ip 1.1.1.2
Destination Routes for 1.1.1.2
--------------------
trust-vr      : => 1.1.1.1/24 (id=3) via 0.0.0.0 (vr: trust-vr)
                Interface ethernet3 , metric 0
```

This indicates that there is a route in the trust-vr via a directly connected interface eth3

A5 - Add a route to the Virtual Router that the appropriate zone and public interface are bound onto.

ns208-> set route 0.0.0.0/0 gateway 1.1.1.2

ns208-> get route
untrust-vr (0 entries)

C - Connected, S - Static, A - Auto-Exported, I - Imported, R - RIP
iB - IBGP, eB - EBGP, O - OSPF, E1 - OSPF external type 1
E2 - OSPF external type 2
trust-vr (3 entries)

ID	IP-Prefix	Interface	Gateway	P	Pref	Mtr	Vsys
* 5	0.0.0.0/0	eth3	1.1.1.2	S	20	1	Root
* 2	192.168.1.0/24	eth1	0.0.0.0	C	0	0	Root
* 3	1.1.1.0/24	eth3	0.0.0.0	C	0	0	Root

NB – The asterisk indicates that the route is active. If an interface is down(No cable connected) the route would not have an asterisk next to it, indicating the route is inactive.

A6 - Can you PING the gateway IP that is defined in the route to the remote firewall? i.e. directly connected router IP.

ns208-> ping 1.1.1.2
Type escape sequence to abort

Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds
ip 1.1.1.2 is unreachable in vr trust-vr

Success Rate is 0 percent.

ns208-> get route
untrust-vr (0 entries)

C - Connected, S - Static, A - Auto-Exported, I - Imported, R - RIP
iB - IBGP, eB - EBGP, O - OSPF, E1 - OSPF external type 1
E2 - OSPF external type 2
trust-vr (3 entries)

ID	IP-Prefix	Interface	Gateway	P	Pref	Mtr	Vsys
* 5	0.0.0.0/0	eth3	1.1.1.2	S	20	1	Root
* 2	192.168.1.0/24	eth1	0.0.0.0	C	0	0	Root
* 3	1.1.1.0/24	eth3	0.0.0.0	C	0	0	Root

ns208-> ping 1.1.1.2
Type escape sequence to abort

Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds
!!!!!
Success Rate is 100 percent (5/5), round-trip time min/avg/max=4/13/50 ms

A7. Make sure that the Public interface status is UP on both firewalls by connecting the UTP cable with RJ-45 interface.
Also Make sure that are no duplicate IP addresses on the subnet that the NetScreen is connected to.

Below you can see that the eth3 state is D = down

ns208-> get interface

A - Active, I - Inactive, U - Up, D - Down, R - Ready

Interfaces in vsys Root:

Name	IP Address	Zone	MAC	VLAN	State	VSD
eth1	192.168.1.1/24	Trust	0010.db19.a7d0	-	U	-
eth2	0.0.0.0/0	DMZ	0010.db19.a7d5	-	D	-
eth3	1.1.1.1/24	Untrust	0010.db19.a7d6	-	D	-
eth4	0.0.0.0/0	Null	0010.db19.a7d7	-	D	-
eth5	0.0.0.0/0	Null	0010.db19.a7d8	-	D	-
eth6	0.0.0.0/0	Null	0010.db19.a7d9	-	D	-
eth7	0.0.0.0/0	Null	0010.db19.a7da	-	D	-
eth8	0.0.0.0/0	HA	0010.db19.a7db	-	U	-
vlan1	0.0.0.0/0	VLAN	0010.db19.a7df	1	D	-

ns208-> get route
untrust-vr (0 entries)

C - Connected, S - Static, A - Auto-Exported, I - Imported, R - RIP
iB - IBGP, eB - EBGP, O - OSPF, E1 - OSPF external type 1
E2 - OSPF external type 2
trust-vr (3 entries)

	ID	IP-Prefix	Interface	Gateway	P	Pref	Mtr	Vsys
	6	0.0.0.0/0	eth3	1.1.1.2	S	20	1	Root
*	2	192.168.1.0/24	eth1	0.0.0.0	C	0	0	Root
	3	1.1.1.0/24	eth3	0.0.0.0	C	0	0	Root

Notice Above that the routes ID 6 and ID 3 that use interface eth3 are inactive (No Asterisk)
Make sure that all cables are connected.

Below shows that the eth3 state is U = Up

ns208-> get interface

A - Active, I - Inactive, U - Up, D - Down, R - Ready

Interfaces in vsys Root:

Name	*IP Address*	*Zone*	*MAC*	*VLAN*	*State*	*VSD*
eth1	*192.168.1.1/24*	*Trust*	*0010.db19.a7d0*	*-*	*U*	*-*
eth2	*0.0.0.0/0*	*DMZ*	*0010.db19.a7d5*	*-*	*D*	*-*
eth3	*1.1.1.1/24*	*Untrust*	*0010.db19.a7d6*	*-*	*U*	*-*
eth4	*0.0.0.0/0*	*Null*	*0010.db19.a7d7*	*-*	*D*	*-*
eth5	*0.0.0.0/0*	*Null*	*0010.db19.a7d8*	*-*	*D*	*-*
eth6	*0.0.0.0/0*	*Null*	*0010.db19.a7d9*	*-*	*D*	*-*
eth7	*0.0.0.0/0*	*Null*	*0010.db19.a7da*	*-*	*D*	*-*
eth8	*0.0.0.0/0*	*HA*	*0010.db19.a7db*	*-*	*U*	*-*
vlan1	*0.0.0.0/0*	*VLAN*	*0010.db19.a7df*	*1*	*D*	*-*

Below you can see that the routes associated with eth3 are now active:

ns208-> get route
untrust-vr (0 entries)

C - Connected, S - Static, A - Auto-Exported, I - Imported, R - RIP
iB - IBGP, eB - EBGP, O - OSPF, E1 - OSPF external type 1
E2 - OSPF external type 2
trust-vr (3 entries)

	ID	*IP-Prefix*	*Interface*	*Gateway*	*P*	*Pref*	*Mtr*	*Vsys*
*	*6*	*0.0.0.0/0*	*eth3*	*1.1.1.2*	*S*	*20*	*1*	*Root*
*	*2*	*192.168.1.0/24*	*eth1*	*0.0.0.0*	*C*	*0*	*0*	*Root*
*	*3*	*1.1.1.0/24*	*eth3*	*0.0.0.0*	*C*	*0*	*0*	*Root*

Duplicate IP on the Untrusted LAN - A symptom of a duplicate IP address on the Untrusted LAN is that the communication works sometimes, and then suddenly stops. As soon as you reboot the NetScreen and try and PING again it seems to work for a while, and then it stops suddenly again.

Lets take for example that there is a server on the Untrusted LAN that has a duplicate IP the same as the NetScreen Untrusted interface. The NetScreen will ARP out of the Untrusted interface for the MAC address for its default gateway IP found in the routing table if it does not have the MAC/IP entry in its ARP table . The gateway will respond with its MAC address. The NetScreen sends frames to the gateway and the gateway will route these packets out.
The problem will come into effect when the server with the duplicate IP also does an ARP request for its default gateway. The Gateway which did have the NetScreen MAC/IP in its ARP table, now will change it to the Server MAC/IP thinking that it has changed. So the NetScreen will send packets out merrily, but the return packets will be sent to the server, and never arrive at the NetScreen Interface.

So the reason why it works when the NetScreen is rebooted, is because when the NetScreen comes up again it will ARP out for the gateway MAC again, and the Gateway will update its ARP table with the NetScreen MAC/IP entry. This will work until the server is rebooted or the ARP entry for its default gateway times out.

A8 - Can you PING any other public IP address e.g. ISP's DNS server?

ns208-> ping 194.73.82.242
Type escape sequence to abort

Sending 5, 100-byte ICMP Echos to 194.73.82.242, timeout is 2 seconds
.....
Success Rate is 0 percent (0/5),

A9 - Contact ISP to see if link between CPE and ISP is down or faulty

A10 - Perform debug flow basic on the NetScreen firewall to see what the problem is

As stated under the concepts, the NetScreen can only catch the debug flow for packets incoming on the ingress interface. If you had to try and ping from the NetScreen to the ISP's DNS server, the outgoing packet would not be caught by the debug flow as it is considered an outgoing packet on the egress interface. Only the incoming return packet would be caught. Though if there was a problem with on layer 2 (ARP) the NetScreen would not be sending the packet out at all, and then would not receive a reply and there would be nothing caught in the debug buffer.

e.g. The best practice is to ping the DNS server or any IP on the Internet that you know is live from a host on the trusted side of the network. In this example we will run a PING from 192.168.1.10 to 194.73.82.242

ns208-> set ff dst-ip 194.73.82.242
filter added
ns208-> get ff
Flow filter based on:
id:0 dst ip 194.73.82.242
ns208-> debug flow basic
ns208-> clear db

Initiate Ping from 192.168.1.10 to destination 194.73.82.242
Below is a working Ping in respect to the flow and the packet actually being sent out of the interface eth3.
In this case we do not see a reply as the ISP is down.

ns208-> get db stream
****** 12553.0: <Trust/ethernet1> packet received [60]******* *Packet arrived on the eth1 interface*
ipid = 29503(733f), @d7806910 *IP id*
packet passed sanity check.
ethernet1:192.168.1.10/1280->194.73.82.242/512,1(8/0)<Root> *Src IP, Port, Dst IP, port incl Protocol 1*
chose interface ethernet1 as incoming nat if. *Int eth1 is placed in NAT mode*
search route to (192.168.1.10->194.73.82.242) in vr trust-vr for vsd-0/flag-0/ifp-null *Route lookup in trust-vr*
route 194.73.82.242->1.1.1.2, to ethernet3 *route found to gateway 1.1.1.2 exiting interface int eth3*
routed (194.73.82.242, 0.0.0.0) from ethernet1 (ethernet1 in 0) to ethernet3 *packet routed*
policy search from zone 2-> zone 1 *Policy lookup performed from Trust (2) to Untrust (1)*
Permitted by policy 3 *matched policy ID 3*
choose interface ethernet3 as outgoing phy if *choose physical interface eth3*
no loop on ifp ethernet3.
session application type 0, name None, timeout 60sec *session time created as 60 seconds for ICMP*
service lookup identified service 0. *service lookup performed*
existing vector list 1-559ef00.
Session (id:76) created for first pak 1 *Create session with ID 76*
route to 1.1.1.2 *Routed packet to 1.1.1.2*
arp entry found for 1.1.1.2 *Already had ARP entry for 1.1.1.2*
nsp2 wing prepared, ready
cache mac in the session *Cached MAC address in the session*
flow got session.
flow session id 76
post addr xlation: 1.1.1.1->194.73.82.242. *Translate src address to egress interface IP*
packet send out to 0010db103041 through ethernet3 *Packet sent out on the wire*

If the ISP was up and working correctly, you would receive the following reply caught in the debug buffer after the ICMP request is sent out.

****** 14466.0: <Untrust/ethernet3> packet received [60]******* *Reply caught on int eth3*
ipid = 1239(04d7), @d785d110
packet passed sanity check.
ethernet3:194.73.82.242/512->1.1.1.1/1036,1(0/0)<Root> *Src IP, port, Dst IP, port and protocol 1*
existing session found. sess token 3
flow got session.
flow session id 98 *Matched outgoing session already created*
post addr xlation: 194.73.82.242->192.168.1.10.
packet send out to 000bdb022203 (cached) through ethernet1 *Packet sent back to 192.168.1.10*

Other possibilities on not getting a reply is that the outgoing packet's IP address is not being source translated (NAT). This will mean that the NetScreen will send the packet out if it matches a policy, but the Private IP range that you are using will be the real source IP. If this packet is not dropped before it reaches the ISP, the ISP will route the reply back to you as the 192.168.0.0/16 range is a non routable IP range. (See RFC)

This is normally because of two factors.

1. If the Trusted interface, or in the case of the NS-208, the eth1 interface is in NAT mode, then you will not have to specify NAT in the advanced properties of the policy.
2. If the Trusted interface, or in the case of the NS-208, the eth1 interface is in route mode, then you must have a policy that has NAT (Source Address Translation) enabled.

```
ns208-> get interface eth1
Interface ethernet1:
  number 0, if_info 0, if_index 0, mode nat
  link up, phy-link up/half-duplex
  vsys Root, zone Trust, vr trust-vr
  dhcp client disabled
  PPPoE disabled
  *ip 192.168.1.1/24   mac 0010.db19.a7d0
  *manage ip 192.168.1.1, mac 0010.db19.a7d0
  route-deny disable
  ping enabled, telnet enabled, SSH enabled, SNMP enabled
  web enabled, ident-reset disabled, SSL enabled
  webauth disabled, webauth-ip 0.0.0.0
  OSPF disabled  BGP disabled  RIP disabled
  bandwidth: physical 100000kbps, configured 0kbps, current 0kbps
        total configured gbw 0kbps, total allocated gbw 0kbps
  DHCP-Relay disabled
  DHCP-server disabled
```

NB – Be very careful when using flow filters. If you have only set a flow filter defining the src-ip and dst- ip, then you would never catch the reply, as the NetScreen will only catch packets for the request going out from 192.168.1.10 to 194.73.82.242. It would not catch the return reply packet from 194.73.82.242 to 192.168.1.10.

So best practice on heavily utilised networks is to specify two flow filter entries. One for the src-ip and dst-ip for the outgoing packet, and one flow filter entry for the reply from the src-ip of the actual original destination only. Or you could specify the src-ip to the NAT ip that was used on the outgoing packet.
The reason for this is that if you are performing NAT on the outgoing packet, then the return packet will not actually be the original source IP (192.168.1.10), but in fact the Egress interface IP or an IP from the DIP pool.

Section B: VPN

Important concept for troubleshooting VPNs and debugging IKE Negotiation:

Remember when checking the event log and in debugging IKE negotiation is that you should always check the event log and run the debug on the NetScreen that will **receive** the P1 IKE request NOT the device that **initiates** the P1 IKE request.
The reason for this is because of the nature of IKE negotiation. If we run a debug on the NS-5GT and try an Extended PING from the inside interface of the NS-5GT (IP 192.168.2.1/24) to the inside interface of the NS-208 (IP address 192.168.1.1/24), and there is a fault with the VPN configuration, then the remote device (NS-208) will drop any IKE requests that are sent from the NS-5GT.

The event log and debug will only display the following type of messages:

Debug run on the NS-208 when itself has initiated P1 to the remote peer
16:03:41 : IKE<0.0.0.0 > IKE: phase-2 packet re-trans timer expired
16:03:41 : IKE<1.1.1.2 > phase-2 packet re-trans timer expired.

It will try and retransmit over and over, and the remote peer will continually drop the request packets. Though you will see whether the IKE negotiation has succeeded passed Phase1 or Phase2 but will never know what the actual problem in Phase2 was causing the fault.

B2 - Perform extended PING from trust interface to remote gateway trust interface.

Once you have the VPN configured, the best exercise is to perform an extended PING sourced from the trusted interface of one gateway destined to the trusted interface of the remote gateway. In the above scenario an extended PING would be from the trust interface on the NS-5GT (IP 192.168.2.1/24) to the eth1 interface on the NS-208 (IP 192.168.1.1/24). We would the then check the event log and debugs on the NS-208 as it is the receiver of the IKE P1 request. The NS-5GT is the initiator.

Non working Ping sourced from the NS-5GT

ns5gt-> ping
Target IP address:192.168.1.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds[2]:
Source interface:trust
Type escape sequence to abort

Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds from trust
.....
Success Rate is 0 percent (0/5),

NB – Make sure that the trusted interface on the NS-5GT and the NS-208 the eth1 is UP. You will not be able to source a packet from an interface that is down.

Working Ping sourced from the NS-5GT:

```
ns5gt-> ping
Target IP address:192.168.1.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds[2]:
Source interface:trust
Type escape sequence to abort

Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds from trust
.!!!!
Success Rate is 80 percent (4/5), round-trip time min/avg/max=5/5/6 ms
ns5gt->
```

NB - If the VPN was not active you will always lose the first few pings as IKE negotiation takes place. The amount of pings lost is dependant on the time latency for traffic between the two firewall.

B3 - Cool. Confirm VPN status

```
ns5gt-> get sa
total configured sa: 1
HEX ID      Gateway      Port Algorithm       SPI          Life:sec  kb     Sta  PID vsys
00000001< 1.1.1.1        500  esp:3des/sha1 e3270b99    3193      unlim  A/-  2   0
00000001> 1.1.1.1        500  esp:3des/sha1 3c472af5    3193      unlim  A/-  1   0
```

HEX ID = Unique VPN identifier
Gateway = Remote Peer gateway IP
Port = IKE Port
Algorithm = P2 Security Algorithm
SPI = Security Profile Identifier
Life:sec = P2 rekey timer in seconds
Kb = P2 rekey timer in kilobytes
Sta = SA status (A Active, D down, I inactive) Also used if VPN monitor is enabled in P2
PID = Policy ID that VPN is bound to)
Vsys = Virtual System where Policy and VPN reside in

```
ns5gt-> get sa id 1
index 0, name VPN to Core - Phase2, peer gateway ip 1.1.1.1. vsys<Root>
auto key. policy node, tunnel mode, policy id in:<2> out:<1> vpngrp:<-1>. sa_list_nxt:<-1>.
tunnel id 1, peer id 0, NSRP Local.     site-to-site. Local interface is untrust <1.1.1.2>.
  esp, group 0, 3des encryption, sha1 authentication
  autokey, IN active, OUT active
  monitor<0>, latency: 0, availability: 0
  DF bit: clear
  app_sa_flags: 0x2067
  proxy id: local 192.168.2.0/255.255.255.0, remote 192.168.1.0/255.255.255.0, proto 0, port 0
  ike activity timestamp: 62839
nat-traversal map not available
incoming: SPI e3270b99, flag 00004000, tunnel info 40000001, vector(d:65e4c0)
  life 3600 sec, 2248 remain, 0 kb, 0 bytes remain
  anti-replay on, last 0x0, window 0x0, idle timeout value <0>, idled 1352 seconds
  next pak sequence number: 0x0
outgoing: SPI 3c472af5, flag 00000000, tunnel info 40000001, vector(e:660f18)
  life 3600 sec, 2248 remain, 0 kb, 0 bytes remain
  anti-replay on, last 0x0, window 0x0, idle timeout value <0>, idled 1342 seconds
  next pak sequence number: 0x4

ns5gt-> get ike cookies

Active: 1, Dead: 0, Total 1

522f/0003, 1.1.1.2:500->1.1.1.1:500, PRESHR/grp2/3DES/SHA, xchg(2) (VPN to Core - Phase1/grp-1/usr-1)
resent-tmr 7166032 lifetime 28800 lt-recv 28800 nxt_rekey 28379 cert-expire 0     (next P1 rekey)
initiator, err cnt 0, send dir 0, cond 0x0
nat-traversal map not available                    (No NAT Traversal being used on the VPN)
ike heartbeat           : disabled
ike heartbeat last rcv time: 0
ike heartbeat last snd time: 0
XAUTH status: 0                (Xauth not enabled)
ns5gt-> get event
2000-04-12 15:31:33 system info  00536 IKE<1.1.1.1> Phase 2 msg ID <11f49203>:
                            Completed negotiations with SPI
                            <e3270b99>, tunnel ID <1>, and
                            lifetime <3600> seconds/<0> KB.
2000-04-12 15:31:33 system info  00536 IKE<1.1.1.1> Phase 2: Initiated
                            negotiations.
2000-04-12 15:31:33 system info  00536 IKE<1.1.1.1> Phase 1: Completed Main
                            mode negotiations with a
                            <28800>-second lifetime.
```

Working debug flow basic:

Packet1:

***** 12532.0: <Trust/ethernet1> packet received [128]******
ipid = 1250(04e2), @d780a910
packet passed sanity check.
ethernet1:192.168.1.10/3500->192.168.2.10/1024,1(8/0)<Root>
chose interface ethernet1 as incoming nat if.
search route to (192.168.1.10->192.168.2.10) in vr trust-vr for vsd-0/flag-0/ifp-null (Search Route)
route 192.168.2.10->1.1.1.2, to ethernet3 (Found Route)
routed (192.168.2.10, 0.0.0.0) from ethernet1 (ethernet1 in 0) to ethernet3
policy search from zone 2-> zone 1 (Policy Search)
Permitted by policy 1 (Found Policy match)
No src xlate choose interface ethernet3 as outgoing phy if (No NAT needed)
no loop on ifp ethernet3.
session application type 0, name None, timeout 60sec
service lookup identified service 0.
existing vector list 5-54a74c0.
Session (id:39) created for first pak 5 (Create session and ID no.)
cache mac in the session
flow got session.
flow session id 39
post addr xlation: 192.168.1.10->192.168.2.10.
going into tunnel 40000001. (Policy matches VPN Tunnel Configuration HEX id 40000001)
flow_encrypt: pipeline.
enqueue to IKE: timems 12527987, Q 1, saidx 0: spi:426ca4a0 done!
packet dropped, SA inactive (*Tunnel is inactive)
handle raw/no_session pakcet.
send no session packet (*Because the first packet arrived at the NetScreen and the SA for the IPSEC tunnel was inactive, the NetScreen must begin P1 and P2 IKE negotiations. Generally it will take 1-2 seconds, and you may drop 1 or 2 pings until the tunnel comes up. If the 2nd packet of the debug displays that the SA is still inactive, then we can move onto debugging the IKE negotiation, as we have confirmed that the packet is indeed arriving at the NetScreen interface, but some part of the IKE negotiation is the problem. Go to section Debugging IKE negotiation)

Packet2 – This is the 2nd packet that we can see matches the active SA once the IKE negotiations have successfully completed and the packet is sent encrypted to the peer gateway.

```
****** 12534.0: <Trust/ethernet1> packet received [128]******
 ipid = 1251(04e3), @d780b110
 packet passed sanity check.
 ethernet1:192.168.1.10/3600->192.168.2.10/1024,1(8/0)<Root>
 chose interface ethernet1 as incoming nat if.
 search route to (192.168.1.10->192.168.2.10) in vr trust-vr for vsd-0/flag-0/ifp-null
 route 192.168.2.10->1.1.1.2, to ethernet3
 routed (192.168.2.10, 0.0.0.0) from ethernet1 (ethernet1 in 0) to ethernet3
 policy search from zone 2-> zone 1
 Permitted by policy 1
 No src xlate   choose interface ethernet3 as outgoing phy if
 no loop on ifp ethernet3.
 session application type 0, name None, timeout 60sec
 service lookup identified service 0.
 existing vector list 5-54a74c0.
 Session (id:41) created for first pak 5
cache mac in the session
 flow got session.
 flow session id 41
 post addr xlation: 192.168.1.10->192.168.2.10.
 going into tunnel 40000001.
 flow_encrypt: pipeline.
chip info: DMA. Tunnel id 00000001                (VPN Tunnel ID)
(vn2)  doing ESP encryption and size =136         (Encrypting the packet)
ipsec encrypt prepare engine done
ipsec encrypt set engine done
ipsec encrypt engine released
ipsec encrypt done
      put packet(455f920) into flush queue.
      remove packet(455f920) out from flush queue.

**** jump to packet:1.1.1.1->1.1.1.2
 out encryption tunnel 40000001 gw:1.1.1.2
 no more encapping needed.
 packet send out to 0010db3a7171 through ethernet3    (ESP Packet sent out to upstream router)
 **** pak processing end.
```

Below is a session created by IKE negotiation (notice that the source port and destination port are 500, as well as the source and destination IP addresses are the public IP addresses of the internet facing interfaces):

```
id 46/s**,vsys 0,flag 00000040/0080/20,policy 320002,time 6, dip 0
 6(0601):1.1.1.2/500->1.1.1.1/500,17,000000000000,3,vlan 0,tun 0,vsd 0,route 0
 3(0010):1.1.1.2/500<-1.1.1.1/500,17,000000000000,4,vlan 0,tun 0,vsd 0,route 0
```

Successful VPN session creation from host 192.168.1.10 to host 192.168.2.10:

ns208-> get session
alloc 176/max 128000, alloc failed 0, di alloc failed 0
*id 45/s**,vsys 0,flag 00000040/0000/00,policy 1,time 6, dip 0*
0(8801):192.168.1.10/4000->192.168.2.10/1024,1,0010db103040,2,vlan 0,tun 0,vsd 0,route 0
6(2800):192.168.1.10/4000<-192.168.2.10/1024,1,000000000000,3,vlan 0,tun 40000001,vsd 0,route 5

B4 - Is there an entry for the VPN in the event log?

'Get event' will display the event log

B5 - Does the policy referencing the VPN from trust to untrust and visa versa include the IP address of the trusted interfaces. Edit Policy so that it includes inside subnet

This would be important for the flow to match the policy so to initiate IKE negotiations.
If the Policy for example was specific to allow a single host from the trust LAN to a single host on the remote LAN, then this would not match the VPN policy, and IKE negotiation would not occur.
If there is no such policy, it would be good practice to create a policy including the whole trusted LAN subnet to the remote trusted LAN so that you can be sure that the Policy will match the packet sourced from the trusted interface. Once all is working correctly, the policy can be defined to match the exact criteria that the live environment requires.

e.g. The following policy on the NS-5GT would be sufficient to match the PING sourced from the trust interface as the IP address is 192.168.2.1/24, and is in the 192.168.2.0/24 subnet.

set policy id 1 from "Trust" to "Untrust" "192.168.2.0/24" "192.168.1.0/24" "ANY" tunnel vpn "VPN to Core - Phase2" id 1 pair-policy 2

B6 - system info 00536 Phase 2: No policy exists for the proxy ID received: local ID (<192.168.1.0>/<255.255.255.0>, <0>, <0>) remote ID (<192.168.2.0>/<255.255.255.0>, <0>, <0>).

In the Event log:

2004-04-06 16:24:25 system info 00536 IKE<1.1.1.2> Phase 2 msg ID <4717e7ea>: Negotiations have failed.
2004-04-06 16:24:25 system info 00536 Rejected an IKE packet on ethernet3 from 1.1.1.2:500 to 1.1.1.1:500 with cookies cfaf76fe7f73ae52 and 57436be50cbe5372 because the peer sent a proxy ID that did not match the one in the SA config.
2004-04-06 16:24:25 system info 00536 IKE<1.1.1.2> Phase 2: No policy exists for the proxy ID received: local ID (<192.168.1.0>/<255.255.255.0>, <0>, <0>) remote ID (<192.168.2.0>/<255.255.255.0>, <0>, <0>).
2004-04-06 16:24:25 system info 00536 IKE<1.1.1.2> Phase 2 msg ID <4717e7ea>: Responded to the peer's first message.

In the debug ike detail:

```
## 16:23:57 : IKE<1.1.1.2      > Recv*: [HASH] [SA] [NONCE] [ID] [ID]
## 16:23:57 : IKE<1.1.1.2      >   extract payload (136):
## 16:23:57 : IKE<1.1.1.2      >   QM in state OAK_QM_SA_ACCEPT.
## 16:23:57 : IKE<1.1.1.2      >   Start by finding matching member SA (verify -1/-1)
## 16:23:57 : IKE<1.1.1.2      >   IKE<1.1.1.2> Matching policy: gw ip <1.1.1.2> peer entry id<0>
## 16:23:57 : IKE<1.1.1.2      >   Local  ID: 192.168.1.0/24 prot<0> port<0> type<4>
## 16:23:57 : IKE<1.1.1.2      >   Remote ID: 192.168.2.0/24 prot<0> port<0> type<4>
## 16:23:57 : IKE<0.0.0.0      >   protocol matched expected<0>.
## 16:23:57 : IKE<0.0.0.0      >   port matched expect<0>.
## 16:23:57 : IKE<1.1.1.2      > Proxy ID match: No policy exists for the proxy ID received
## 16:23:57 : IKE<1.1.1.2      > proxy-id do not match ipsec sa config
## 16:23:57 : IKE<1.1.1.2      >   oakley_process_quick_mode():exit
## 16:23:57 : IKE<1.1.1.2      > Phase 2 msg-id <4717e7ea>: Negotiations have failed.
```

The most common problem made in VPN configurations is that the proxy ID's are incorrect. In Phase2 of IKE negotiation, as part of the security process, each side of the tunnel will send its proxy ID's across as a security measure.

The mistake many administrators make is that, they will have different address objects on each side, so that the proxy id's sent from one side does not match what is expected at the other side.

An example that will not work:

On the NS-208 a policy is configured from zone trust to zone untrust , source address ANY to destination address 192.168.2.0/24 service ANY. On the NS-5GT there is a policy from zone untrust to zone trust, source address 192.168.1.0/24 service ANY to destination address 192.168.2.0/24 service ANY.

The polices will look like this:

On the NS-208:

set policy id 1 from "Trust" to "Untrust" "Any" "192.168.2.0/24" "ANY" tunnel vpn "VPN to Remote Site1 - Phase2" id 2 pair-policy 2
set policy id 2 from "Untrust" to "Trust" "192.168.2.0/24" "Any" "ANY" tunnel vpn "VPN to Remote Site1 - Phase2" id 2 pair-policy 1

On the NS-5GT:

set policy id 2 from "Untrust" to "Trust" "192.168.1.0/24" "192.168.2.0/24" "ANY" tunnel vpn "VPN to Core - Phase2" id 1 pair-policy 1
set policy id 1 from "Trust" to "Untrust" "192.168.2.0/24" "192.168.1.0/24" "ANY" tunnel vpn "VPN to Core - Phase2" id 1 pair-policy 2

This above will cause the VPN to fail on Phase2. To view the proxy ID's that will be sent in Phase2 by each device and what each device will expect to receive in P2 is performed by looking at the policy in more detail.

To view any policy in more detail, get policy id x:
On the NS-208:

ns208-> get policy id 1
name:"none" (id 1), zone Trust -> Untrust,action Tunnel, status "enabled", pair policy 2
src "Any", dst "192.168.2.0/24", serv "ANY"
Policies on this vpn tunnel: 0
nat off, url filtering OFF
vpn VPN to Remote Site1 - Phase2, nsp tunnel 40000002, sa index 1, sa tunnel id 2
policy flag 0000, session backup: on
traffic shapping off, scheduler n/a, serv flag 00
log no, log count 0, alert no, counter no(0) byte rate(sec/min) 0/0
total octets 285846, counter(session/packet/octet) 0/0/0
priority 7, diffserv marking Off
tadapter: state off, gbw/mbw 0/-1
proxy id:
local 0.0.0.0/0.0.0.0, remote 192.168.2.0/255.255.255.0, proto 0, port 0
No Authentication
No User, User Group or Group expression set

On the NS-5GT:

ns5gt-> get policy id 2
name:"none" (id 2), zone Untrust -> Trust,action Tunnel, status "enabled", pair policy 1
src "192.168.1.0/24", dst "192.168.2.0/24", serv "ANY"
Policies on this vpn tunnel: 1
[192.168.1.0/24, 192.168.2.0/24, 0-65535, 0-65535, 0]
nat off, url filtering OFF
vpn VPN to Core - Phase2, nsp tunnel 40000001, sa index 0, sa tunnel id 1
policy flag 0000, session backup: on
traffic shapping off, scheduler n/a, serv flag 00
log no, log count 0, alert no, counter no(0) byte rate(sec/min) 0/0
total octets 284852, counter(session/packet/octet) 0/0/0
priority 7, diffserv marking Off
tadapter: state off, gbw/mbw 0/-1
proxy id:
local 192.168.2.0/255.255.255.0, remote 192.168.1.0/255.255.255.0, proto 0, port 0
No Authentication
No User, User Group or Group expression set
Working Configuration:

B7 - Edit Policies referencing VPN

set policy id 2 from "Untrust" to "Trust" "192.168.2.0/24" "192.168.1.0/24" "ANY" tunnel vpn "VPN to Remote Site1 - Phase2" id 3 pair-policy 1
set policy id 1 from "Trust" to "Untrust" "192.168.1.0/24" "192.168.2.0/24" "ANY" tunnel vpn "VPN to Remote Site1 - Phase2" id 3 pair-policy 2

ns208-> get policy id 1
name:"none" (id 1), zone Trust -> Untrust,action Tunnel, status "enabled", pair policy 2
src "192.168.1.0/24", dst "192.168.2.0/24", serv "ANY"
Policies on this vpn tunnel: 0
nat off, url filtering OFF
vpn VPN to Remote Site1 - Phase2, nsp tunnel 40000003, sa index 0, sa tunnel id 3
policy flag 0000, session backup: on
traffic shapping off, scheduler n/a, serv flag 00
log no, log count 0, alert no, counter no(0) byte rate(sec/min) 0/0
total octets 285846, counter(session/packet/octet) 0/0/0
priority 7, diffserv marking Off
tadapter: state off, gbw/mbw 0/-1
proxy id:
local 192.168.1.0/255.255.255.0, remote 192.168.2.0/255.255.255.0, proto 0, port 0
No Authentication
No User, User Group or Group expression set

B8 - system info 00536 Rejected an IKE packet on ethernet3 because there were no acceptable Phase 1 proposals.

Another common problem made is that the P1 proposals that are chosen do not match on both sides of the IPSEC gateway. There must at least be one set of Phase1 matching proposals on both sides of the IPSEC tunnel to complete Phase1 negotiations.

1. Check the event log by 'get event' on the NetScreen Console, SSH or Telnet session
2. Run a 'debug ike detail'

ns208-> get event
2004-04-06 16:52:58 system info 00536 IKE<1.1.1.2> Phase 1: Discarded a second initial packet, which arrived within 5 seconds after the first.
2004-04-06 16:52:54 system info 00536 Rejected an IKE packet on ethernet3 from 1.1.1.2:500 to 1.1.1.1:500 with cookies 1b61a768b134a44a and 99cd98ba0e182b07 because there were no acceptable Phase 1 proposals.

Ns208->get db stream
16:52:15 : IKE<1.1.1.2 > Process [SA]:
16:52:15 : IKE<1.1.1.2 > Proposal received:
16:52:15 : IKE<1.1.1.2 > auth(1)<PRESHRD>, encr(5)<3DES>, hash(2)<SHA>, group(2)
16:52:15 : IKE<1.1.1.2 > xauth: disabled
16:52:15 : IKE<1.1.1.2 >
16:52:15 : IKE<1.1.1.2 > xauth flag: 0, 0
16:52:15 : IKE<1.1.1.2 > auth value: 1, 1
16:52:15 : IKE<1.1.1.2 > enc value: 5, 5
16:52:15 : IKE<1.1.1.2 > [0] expect:
16:52:15 : IKE<1.1.1.2 > auth(1)<PRESHRD>, encr(5)<3DES>, hash(1)<MD5>, group(2)
16:52:15 : IKE<1.1.1.2 > xauth: disabled
16:52:15 : IKE<1.1.1.2 > Phase 1: Rejected proposals from peer. Negotiations failed.
16:52:15 : IKE<1.1.1.2> (0,r): ERROR:
IKE error, when processing oak_no_sate

Above you can see in the debug on the NS-208, that the NS-5GT sent the P1 proposal:

auth(1)<PRESHRD>, encr(5)<3DES>, hash(2)<SHA>, group(2)

But the NS-208 was expecting below:

auth(1)<PRESHRD>, encr(5)<3DES>, hash(1)<MD5>, group(2)

Config on the NS-208:

set ike gateway "VPN to Remote Site1 - Phase1" address 1.1.1.2 Main outgoing-interface "ethernet3" preshare "JSNnvXl4NzpQSfsMJiCIhVnT5WnkdLOZyw==" proposal "pre-g2-3des-md5"

Config on the NS-5GT:

set ike gateway "VPN to Core - Phase1" address 1.1.1.1 Main outgoing-interface "untrust" preshare "M4E4rIY7NuxSrlso6HCrT2oJ4yn2VYWEbA==" proposal "pre-g2-3des-sha"

B9 - Check P1 proposals on both sides
FIX: Change the proposals on one side to match the other:

Config on the NS-208:

set ike gateway "VPN to Remote Site1 - Phase1" address 1.1.1.2 Main outgoing-interface "ethernet3" preshare "JSNnvXl4NzpQSfsMJiCIhVnT5WnkdLOZyw==" proposal "pre-g2-3des-sha"

Config on the NS-5GT:

set ike gateway "VPN to Core - Phase1" address 1.1.1.1 Main outgoing-interface "untrust" preshare "M4E4rIY7NuxSrlso6HCrT2oJ4yn2VYWEbA==" proposal "pre-g2-3des-sha"

B10 - system info 00536 Rejected an IKE packet on ethernet3 because there were no acceptable Phase 2 proposals.

Another common problem made is that the P2 proposals that are chosen do not match on both sides of the IPSEC gateway. There must at least be one set of Phase2 matching proposals on both sides of the IPSEC tunnel to complete Phase2 negotiations.

ns208-> get event
Date Time Module Level Type Description
2004-04-06 17:15:01 system info 00536 IKE<1.1.1.2> Phase 2 msg ID <231af9e6>:
Negotiations have failed.
2004-04-06 17:15:01 system info 00536 Rejected an IKE packet on ethernet3
from 1.1.1.2:500 to 1.1.1.1:500 with
cookies 83e1f5209ce4640e and
fc666af880db4a5c because there were no
acceptable Phase 2 proposals.
2004-04-06 17:15:01 system info 00536 IKE<1.1.1.2> Phase 2 msg ID <231af9e6>:
Responded to the peer's first message.
2004-04-06 17:15:01 system info 00536 IKE<1.1.1.2> Phase 1: Completed Main
mode negotiations with a
<28800>-second lifetime.
2004-04-06 17:15:01 system info 00536 IKE<1.1.1.2> Phase 1: Responder starts
MAIN mode negotiations.
2004-04-06 17:14:54 system notif 00767 All logged events or alarms were
cleared by admin NetScreen
Total entries matched = 6

Above and on the last page you can see that Phase 1 completes successfully, but Phase 2 fails:
Below you can see from the debug ike detail that Phase2 fails because the proposals were rejected:

17:15:01 : IKE<1.1.1.2 > Phase 2 received:
17:15:01 : IKE<1.1.1.2 > atts<00000003 00000000 00000003 00000002 00000001 00000000>
17:15:01 : IKE<1.1.1.2 > proto(3)<ESP>, esp(3)<ESP_3DES>, auth(2)<SHA>, encap(1)<TUNNEL>, group(0)
17:15:01 : IKE<1.1.1.2 > expect [0]:
17:15:01 : IKE<1.1.1.2 > atts<00000003 00000000 00000003 00000001 00000001 00000000>
17:15:01 : IKE<1.1.1.2 > proto(3)<ESP>, esp(3)<ESP_3DES>, auth(1)<MD5>, encap(1)<TUNNEL>, group(0)
17:15:01 : IKE<1.1.1.2 > proposal not acceptable, but no more proposal in payload.
17:15:01 : IKE<1.1.1.2 > Phase 2: Rejected proposals from peer. Negotiations failed.
17:15:01 : IKE<1.1.1.2> (3,r): ERROR:

Config on the NS-208:

set vpn "VPN to Remote Site1 - Phase2" gateway "VPN to Remote Site1 - Phase1" replay tunnel idletime 0 proposal "nopfs-esp-3des-md5"

Config on the NS-5GT:

set vpn "VPN to Core - Phase2" gateway "VPN to Core - Phase1" replay tunnel idletime 0 proposal "nopfs-esp-3des-sha"

B11 - Check P2 proposals on both sides

FIX: Change the proposals on one side to match the other:

Config on the NS-208:

set vpn "VPN to Remote Site1 - Phase2" gateway "VPN to Remote Site1 - Phase1" replay tunnel idletime 0 proposal "nopfs-esp-3des-sha"

Config on the NS-5GT:

set vpn "VPN to Core - Phase2" gateway "VPN to Core - Phase1" replay tunnel idletime 0 proposal "nopfs-esp-3des-sha"

B12 - system info 00536 Rejected an IKE packet on ethernet3 because Phase 1 negotiations failed. (The preshared keys might not match.)

The preshared key is different on both gateways.

ns208-> get event
Date Time Module Level Type Description
2004-04-06 17:22:26 system info 00536 Rejected an IKE packet on ethernet3 from 1.1.1.2:500 to 1.1.1.1:500 with cookies 61561626aed0e308 and b78f53c619774596 because Phase 1 negotiations failed. (The preshared keys might not match.).
2004-04-06 17:22:26 system info 00536 IKE<1.1.1.2> Phase 1: Responder starts MAIN mode negotiations.

On the debug ike detail:

```
## 17:22:26 : IKE<1.1.1.2      > Construct ISAKMP header.
## 17:22:26 : IKE<1.1.1.2      >   Msg header built (next payload #4)
## 17:22:26 : IKE<1.1.1.2      > Construct [KE] for ISAKMP
## 17:22:26 : IKE<1.1.1.2      > Construct [NONCE]
## 17:22:26 : IKE<1.1.1.2      >   throw packet to the peer, paket_len=184
## 17:22:26 : IKE<1.1.1.2      > Xmit : [KE] [NONCE]
## 17:22:26 : IKE<1.1.1.2      >   send_request to peer
## 17:22:26 : IKE<1.1.1.2      >   Send Phase 1 packet (len=184)
## 17:22:26 : IKE<1.1.1.2      >   ike packet, len 96, action 0
## 17:22:26 : IKE<0.0.0.0      >   coach. sock 1024
## 17:22:26 : IKE<1.1.1.2      > ****** Recv packet if <ethernet3> of vsys <Root> ******
## 17:22:26 : IKE<1.1.1.2      >   Catcher: get 68 bytes. src port 500
## 17:22:26 : IKE<1.1.1.2      >   SA: (Root, local 1.1.1.1, state 2/620f, r):
## 17:22:26 : IKE<1.1.1.2      >   ISAKMP msg: len 68, nxp 5[ID], exch 2[MM], flag 01  E
## 17:22:26 : IKE<1.1.1.2      >   gen_skeyid()
## 17:22:26 : IKE<1.1.1.2      >   Decrypting payload (length 40)
## 17:22:26 : IKE<1.1.1.2      > Recv*: [ID]  +++ Corrupted MSG
## 17:22:26 : IKE<1.1.1.2      >     Validate (40): bad 30
## 17:22:26 : IKE<1.1.1.2      > Packet is invalid!
## 17:22:26 : IKE<1.1.1.2      > Pre-shared key might not match.
```

On the last page you will see that after the NS-208 responder constructs message 3 containing the nonce generated with the preshare key and sends it across to the NS-5GT. In Message 4, the NS-5GT sends its nonce generated with its preshare key and when decrypted by the NS-208 the packet is invalid

Config on the NS-208

set ike gateway "VPN to Remote Site1 - Phase1" address 1.1.1.2 Main outgoing-interface "ethernet3" preshare "hvaP3DoZN6LF8RsWmgCZCJzkzynt6rGlhQ==" proposal "pre-g2-3des-sha"

Config on the NS-5GT:

set ike gateway "VPN to Core - Phase1" address 1.1.1.1 Main outgoing-interface "untrust" preshare "M4E4rIY7NuxSrlso6HCrT2oJ4yn2VYWEbA==" proposal "pre-g2-3des-sha"

B13 - Check that preshare keys match on P1 configuration

FIX: Change the Preshare Keys to the same values on both sides of the gateway.

set ike gateway "VPN to Remote Site1 - Phase1" address 1.1.1.2 Main outgoing-interface "ethernet3" preshare "FtEollXnN0+hsbsSmACbhDeohunzf6rFtQ==" proposal "pre-g2-3des-sha"

As seen above the output values in the NetScreen config is encrypted with a random value generator. If the preshare key is unknown, then the best bet is to change the keys on both sides to a known value.

B14 - system info 00536 Rejected an IKE packet on ethernet3 because an initial Phase 1 packet arrived from an unrecognized peer gateway.

Incorrect outgoing interface specified under Autokey Advanced>gateway configuration

ns208-> get event
Date Time Module Level Type Description
2004-04-06 17:37:36 system info 00536 Rejected an IKE packet on ethernet3 from 1.1.1.2:500 to 1.1.1.1:500 with cookies 2744c0e8d008c6e3 and 0000000000000000 because an initial Phase 1 packet arrived from an unrecognized peer gateway.

Get db stream

17:38:00 : IKE<1.1.1.2 > Recv : [SA] [VID] [VID]
17:38:00 : IKE<1.1.1.2 > Find gateway by peer IP and local ifp.
17:38:00 : IKE<1.1.1.2 > Getting the 1st peer_ent that is used, with no peer IP, and right local IP.
17:38:00 : IKE<1.1.1.2 > Rejected an initial Phase 1 packet from an unrecognized peer gateway.
17:38:04 : IKE<1.1.1.2 > ike packet, len 164, action 1
17:38:04 : IKE<0.0.0.0 > coach. sock 1025
*## 17:38:04 : IKE<1.1.1.2 > ****** Recv packet if <ethernet3> of vsys <Root> ******
17:38:04 : IKE<1.1.1.2 > Catcher: get 136 bytes. src port 500*
17:38:04 : IKE<1.1.1.2 > New Phase 1 SA
17:38:04 : IKE<1.1.1.2 > ISAKMP msg: len 136, nxp 1[SA], exch 2[MM], flag 00
17:38:04 : IKE<1.1.1.2 > Recv : [SA] [VID] [VID]
17:38:04 : IKE<1.1.1.2 > Find gateway by peer IP and local ifp.
17:38:04 : IKE<1.1.1.2 > Getting the 1st peer_ent that is used, with no peer IP, and right local IP.
17:38:04 : IKE<1.1.1.2 > Rejected an initial Phase 1 packet from an unrecognized peer gateway.

When the VPN is configured, the NetScreen needs to know which interface it should expect to receive P1 IKE negotiations on. If it receives P1 IKE packet destined to its own interface IP address, but that Autokey Advanced>gateway configuration has not specified that interface, then the packet will be dropped:

Config on NS-208:

set interface ethernet1 ip 192.168.1.1/24
set interface ethernet1 nat
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 route
set ike gateway "VPN to Remote Site1 - Phase1" address 1.1.1.2 Main outgoing-interface "ethernet1" preshare "FtEollXnN0+hsbsSmACbhDeohunzf6rFtQ==" proposal "pre-g2-3des-sha"

Config on the NS-5GT:

set ike gateway "VPN to Core - Phase1" address 1.1.1.1 Main outgoing-interface "untrust" preshare "M4E4rIY7NuxSrlso6HCrT2oJ4yn2VYWEbA==" proposal "pre-g2-3des-sha"

Above you can see that the NS-208 is expecting the P1 from the initiator on interface eth1, but has received it on eth3. In actual fact, the NS-208 is configured incorrectly and should have the eth3 specified in the configuration, as this is internet facing interface, and the remote IP set on the NS-5GT for the NS-208 gateway.

B15 - Check outgoing VPN interface configured in P1
FIX: Make sure the outgoing interface in P1 is the correct interface that the NetScreen will receive the P1 request from the initiator on.

set ike gateway "VPN to Remote Site1 - Phase1" address 1.1.1.2 Main outgoing-interface "ethernet3" preshare "FtEollXnN0+hsbsSmACbhDeohunzf6rFtQ==" proposal "pre-g2-3des-sha"

B16 - system info 00536 Phase 2 msg ID <ae86e461>: Completed negotiations with SPI <3c472af4>, tunnel ID <1>, and lifetime <3600> seconds/<0> KB.

The tunnel IKE negotiation has completed successfully.

B17 - The VPN is active. Is PING management enabled on the trusted interface of both firewalls? Enable PING Management on trusted interfaces on both firewalls

ns208-> get interface eth1
Interface ethernet1:
number 0, if_info 0, if_index 0, mode nat
link up, phy-link up/half-duplex
vsys Root, zone Trust, vr trust-vr
dhcp client disabled
PPPoE disabled
**ip 192.168.1.1/24 mac 0010.db19.a7d0*
**manage ip 192.168.1.1, mac 0010.db19.a7d0*
route-deny disable
ping disabled, telnet enabled, SSH enabled, SNMP enabled
web enabled, ident-reset disabled, SSL enabled
webauth disabled, webauth-ip 0.0.0.0
OSPF disabled BGP disabled RIP disabled
bandwidth: physical 100000kbps, configured 0kbps, current 0kbps
total configured gbw 0kbps, total allocated gbw 0kbps
DHCP-Relay disabled
DHCP-server disabled

ns208-> set interface eth1 manage ping

ns208-> get interface eth1
Interface ethernet1:
number 0, if_info 0, if_index 0, mode nat
link up, phy-link up/half-duplex
vsys Root, zone Trust, vr trust-vr
dhcp client disabled
PPPoE disabled
**ip 192.168.1.1/24 mac 0010.db19.a7d0*
**manage ip 192.168.1.1, mac 0010.db19.a7d0*
route-deny disable
ping enabled, telnet enabled, SSH enabled, SNMP enabled
web enabled, ident-reset disabled, SSL enabled
webauth disabled, webauth-ip 0.0.0.0
OSPF disabled BGP disabled RIP disabled
bandwidth: physical 100000kbps, configured 0kbps, current 0kbps
total configured gbw 0kbps, total allocated gbw 0kbps
DHCP-Relay disabled
DHCP-server disabled

Working NS-208 Config:

```
ns208-> get conf
Total Config size 2930:
unset hardware wdt-reset
set clock timezone 0
set vrouter trust-vr sharable
unset vrouter "trust-vr" auto-route-export
set auth-server "Local" id 0
set auth-server "Local" server-name "Local"
set auth default auth server "Local"
set admin name "NetScreen"
set admin password "nKVUM2rwMUzPcrkG5sWIHdCtqkAibn"
set admin auth timeout 10
set admin auth server "Local"
set admin format dos
set zone "Trust" vrouter "trust-vr"
set zone "Untrust" vrouter "trust-vr"
set zone "DMZ" vrouter "trust-vr"
set zone "VLAN" vrouter "trust-vr"
set zone "Trust" tcp-rst
set zone "Untrust" block
unset zone "Untrust" tcp-rst
set zone "MGT" block
set zone "DMZ" tcp-rst
set zone "VLAN" block
set zone "VLAN" tcp-rst
set zone "Untrust" screen tear-drop
set zone "Untrust" screen syn-flood
set zone "Untrust" screen ping-death
set zone "Untrust" screen ip-filter-src
set zone "Untrust" screen land
set zone "V1-Untrust" screen tear-drop
set zone "V1-Untrust" screen syn-flood
set zone "V1-Untrust" screen ping-death
set zone "V1-Untrust" screen ip-filter-src
set zone "V1-Untrust" screen land
set interface "ethernet1" zone "Trust"
set interface "ethernet2" zone "DMZ"
set interface "ethernet3" zone "Untrust"
unset interface vlan1 ip
set interface ethernet1 ip 192.168.1.1/24
set interface ethernet1 nat
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 route
unset interface vlan1 bypass-others-ipsec
unset interface vlan1 bypass-non-ip
set interface ethernet1 ip manageable
set interface ethernet3 ip manageable
set interface ethernet3 manage ping
set interface ethernet3 manage ssh
set interface ethernet3 manage ssl
unset flow no-tcp-seq-check
set flow tcp-syn-check
set console timeout 0
set console page 0
set hostname ns208
set address "Trust" "192.168.1.0/24" 192.168.1.0 255.255.255.0 "Core Trusted LAN"
set address "Untrust" "192.168.2.0/24" 192.168.2.0 255.255.255.0 "Remote Site1 Trusted LAN"
set ike gateway "VPN to Remote Site1 - Phase1" address 1.1.1.2 Main outgoing-interface "ethernet3" preshare "JSNnvXl4NzpQSfsMJiCIhVnT5WnkdLOZyw==" proposal "pre-g2-3des-sha"
set ike respond-bad-spi 1
set vpn "VPN to Remote Site1 - Phase2" gateway "VPN to Remote Site1 - Phase1" replay tunnel idletime 0 proposal "nopfs-esp-3des-sha"
set policy id 2 from "Untrust" to "Trust"  "192.168.2.0/24" "192.168.1.0/24" "ANY" tunnel vpn "VPN to Remote Site1 - Phase2" id 1 pair-policy 1
set policy id 1 from "Trust" to "Untrust"  "192.168.1.0/24" "192.168.2.0/24" "ANY" tunnel vpn "VPN to Remote Site1 - Phase2" id 1 pair-policy 2
set pki authority default scep mode "auto"
set pki x509 default cert-path partial
set ssh version v2
set ssh enable
set config lock timeout 5
set snmp port listen 161
set snmp port trap 162
set vrouter "untrust-vr"
exit
set vrouter "trust-vr"
unset add-default-route
set route  0.0.0.0/0 interface ethernet3 gateway 1.1.1.2
exit
```

Working NS-5GT config:

```
ns5gt-> get conf
Total Config size 3293:
unset hardware wdt-reset
set clock timezone 0
set vrouter trust-vr sharable
unset vrouter "trust-vr" auto-route-export
set auth-server "Local" id 0
set auth-server "Local" server-name "Local"
set auth default auth server "Local"
set admin name "NetScreen"
set admin password "nKVUM2rwMUzPcrkG5sWIHdCtqkAibn"
set admin auth timeout 10
set admin auth server "Local"
set admin format dos
set zone "Trust" vrouter "trust-vr"
set zone "Untrust" vrouter "trust-vr"
set zone "VLAN" vrouter "trust-vr"
set zone "Trust" tcp-rst
set zone "Untrust" block
unset zone "Untrust" tcp-rst
set zone "MGT" block
set zone "VLAN" block
set zone "VLAN" tcp-rst
set zone "Untrust" screen tear-drop
set zone "Untrust" screen syn-flood
set zone "Untrust" screen ping-death
set zone "Untrust" screen ip-filter-src
set zone "Untrust" screen land
set interface "trust" zone "Trust"
set interface "untrust" zone "Untrust"
unset interface vlan1 ip
set interface trust ip 192.168.2.1/24
set interface trust nat
set interface untrust ip 1.1.1.2/24
set interface untrust route
unset interface vlan1 bypass-others-ipsec
unset interface vlan1 bypass-non-ip
set interface trust ip manageable
set interface untrust ip manageable
set interface untrust manage ping
set interface untrust manage ssh
set interface untrust manage ssl
set interface trust dhcp server service
set interface trust dhcp server auto
set interface trust dhcp server option gateway 192.168.1.1
set interface trust dhcp server option netmask 255.255.255.0
set interface trust dhcp server ip 192.168.1.33 to 192.168.1.126
set flow tcp-mss
unset flow no-tcp-seq-check
set flow tcp-syn-check
set console timeout 0
set console page 0
set hostname ns5gt
set address "Trust" "192.168.2.0/24" 192.168.2.0 255.255.255.0 "Remote Site1 Trusted LAN"
set address "Untrust" "192.168.1.0/24" 192.168.1.0 255.255.255.0 "Core Trusted LAN"
set ike gateway "VPN to Core - Phase1" address 1.1.1.1 Main outgoing-interface "untrust" preshare "M4E4rIY7NuxSrlso6HCrT2oJ4yn2VYWEbA==" proposal "pre-g2-
3des-sha"
set ike respond-bad-spi 1
set vpn "VPN to Core - Phase2" gateway "VPN to Core - Phase1" replay tunnel idletime 0 proposal "nopfs-esp-3des-sha"
set pki authority default scep mode "auto"
set pki x509 default cert-path partial
set policy id 2 from "Untrust" to "Trust"  "192.168.1.0/24" "192.168.2.0/24" "ANY" tunnel vpn "VPN to Core - Phase2" id 1 pair-policy 1
set policy id 1 from "Trust" to "Untrust"  "192.168.2.0/24" "192.168.1.0/24" "ANY" tunnel vpn "VPN to Core - Phase2" id 1 pair-policy 2
set global-pro policy-manager primary outgoing-interface untrust
set global-pro policy-manager secondary outgoing-interface untrust
set ssh version v2
set ssh enable
set config lock timeout 5
set modem speed 115200
set modem retry 3
set modem interval 10
set modem idle-time 10
set snmp name "ns5gt"
set snmp port listen 161
set snmp port trap 162
set vrouter "untrust-vr"
exit
set vrouter "trust-vr"
unset add-default-route
set route  0.0.0.0/0 interface untrust gateway 1.1.1.1
exit
```