



Title: NSRP Clustering and Path Monitoring
Document Number: FW-400-004
Version: 1.0, August 31, 2002
OS Ver: Screen OS 4.0
HW Platforms this Paper Applies to: NS500
Audience (Internal or External): Internal
Version: DRAFT

NSRP Clustering and Path Monitoring

Introduction

This document is intended to provide generic guideline in using path monitoring to avoid any improper implementation that might cause unexpected behavior of the NSRP cluster pair. The discussion of configuring and tuning Netscreen Redundant Protocol (NSRP) is out of the scope of this paper.

This document does not intent to provide absolute use of path monitoring in your network infrastructures. The use of path monitoring is governed by the need of your unique network environment and requirement.

Concept

Path monitoring checks the network connection and link status of the Netscreen interface and interface of another devices. Path monitoring is one of the components in NSRP to provide the criteria that determine the operating state of the Netscreen firewalls. Interface link state monitoring and track IP are the two method used for path monitoring.

Track IP is a technique that uses either ARP request or ping to check the reachability of a particular IP address or a set of IP addresses.

Each track IP includes two parameters: weight and threshold. The total of the weighted sum of all track IP's on the firewall is compared to the threshold, and if this value exceeds the threshold, failover will occur.

Weighted sum is sum of all the weight of each tracked IP. The weighted sum is the value that the firewall used to compare against the failover threshold. Initially, the NetScreen will have a weighted sum of 0.

Once the fail count reaches the threshold, the weight will change from zero to the assign value.

When the weighted sum exceeds the failover threshold, the firewall will change the operating state from active or backup to inoperable. The device failover threshold is defaulted to 255 and it can be range from 1 to 255. Note that this value is global to the firewall.

Track IP configuration will not be synchronized between the Netscreen firewall in the cluster. Each tracked IP will be associated with,

1. Weight
2. Threshold
3. Interval
4. Outgoing interface
5. Method
6. Fail count
7. Success rate

Depending on the path reliability of the tracked IP, the track IP threshold may need to increase to avoid any false alarm.

The success rate is the cumulative success rate computed from either after the track IP was enabled or the firewall was reloaded. Thus, the longer the firewall is up, the more accurate this value is.

By default once the cluster identification is created all of the main interfaces are assign to Virtual System Device (VSD) zero. However, you can unset the VSD zero group and this will cause the configuration parameters on the main interfaces not to synchronized with the other Netscreen in the cluster.

Example: Track IP configuration on one of the firewall

Track IP	Interval	Threshold	Weight	Interface	Method	Fail count	Success Rate
2.2.2.1	1	15	10	E3	Ping	11	98%
172.16.1.1	1	15	5	E1	Ping	0	96%
172.16.1.2	1	15	5	E1	Ping	0	96%

The firewall failover threshold is 19. In the example above, if the IP address of 2.2.2.1 is not reachable by the firewall as you can see on the fail count column. The weighted sum will then changes from zero to 10 (10+0+0). But 10 is lower than the failover threshold of 19, thus state of the firewall will remain unchanged.

If all of the IP addresses listed above were not reachable by the Netscreen, the weighted sum would be $10+5+5 = 20$, which are greater than the failover threshold of 19, thus the firewall will change the operating state to inoperable.

Common Errors in Implementing Path Monitoring

The decision to implement track IP is not as simple as just tracking the default internal and external gateway of the Netscreen firewalls. Improperly implemented track IP will have adverse effect to the Netscreen high availability behavior and may lost all connectivity to and from the Netscreen firewalls.

From this section onward, path monitoring and track IP are synonymous.

Some of the common implementation errors include:

1. Both of the Netscreen firewalls configured in a cluster pair are tracking to the same set IP addresses.
2. Both of the Netscreen firewalls are using the same network path to reach the tracked IP addresses.
3. The triggered failover does not help in restoring any network connectivity.

In Screen 3.1 onward, assuming that all of the Netscreen are tracking the upstream default gateway only and using the same failover threshold. When the gateway was not reachable, all of the Netscreen firewalls in the cluster will change its state to inoperable. Hence all of the inter zones connectivity via the firewall are affected by this changes.

This scenario must be avoided as in some scenario; the external connectivity may be broken due to Wide Area Network (WAN) link failure or hardware failure on the upstream devices but there are servers on the DMZ zone or custom zones that must still be able to access by internal users regardless of any external devices or link failures.

NSRP Clustering and Resilience Network Design

In a well architect resilience network design, track IP is generally not require because all possible single point of failure has been identified and eliminated, thus non of the interfaces, devices and links failure will cause failure to both incoming and outgoing network traffics. Routing redundancy and routes convergence are addressed by using dynamic routing protocol such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF) or Hot Standby Routing Protocol (HSRP).

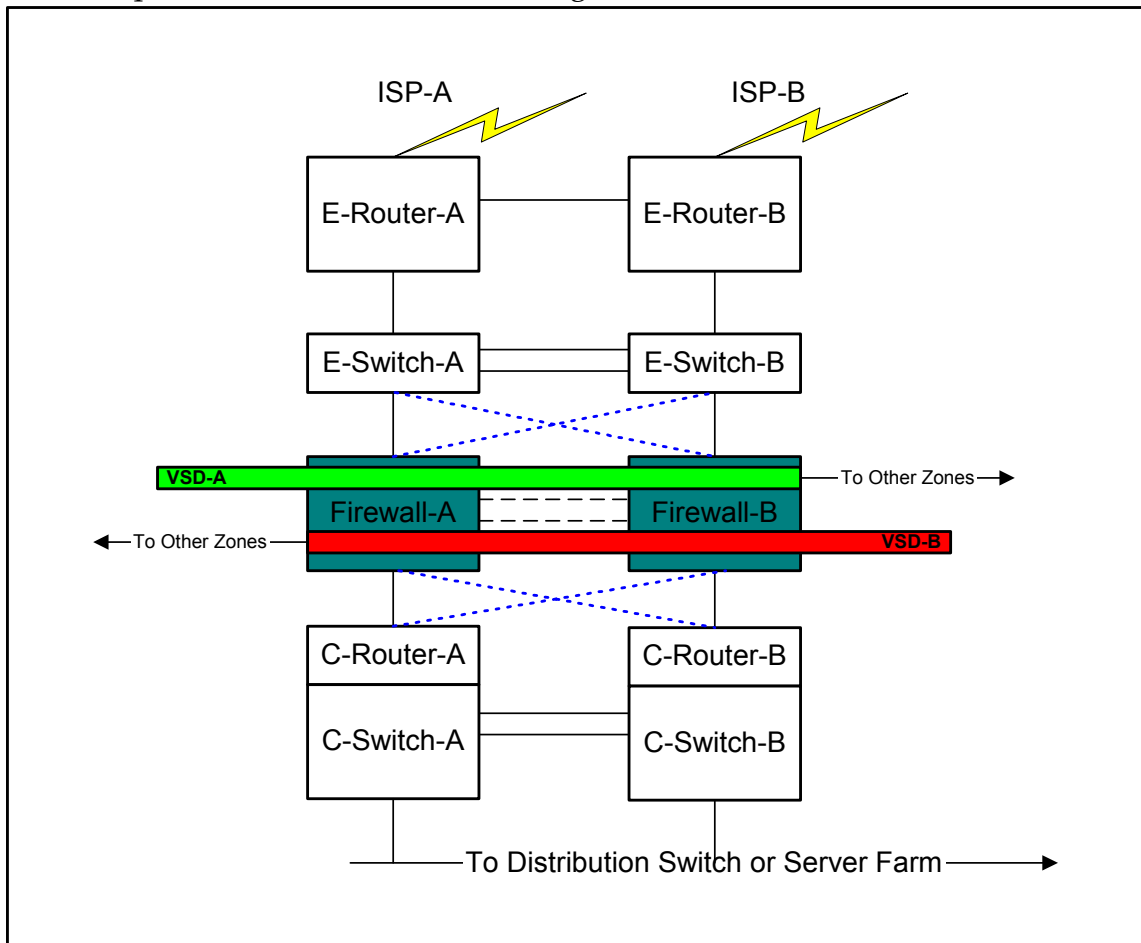
Monitoring the link status on the Netscreen firewall interfaces is sufficient to maintain network connectivity in most of the implementation.

Using track IP to trigger a fail over will not help in restoring any network connectivity, thus the fail over is not necessary. In the scenario of active-active

NSRP clustering, all traffic loads will be shifted to a single firewall and this defeat the purpose of active-active NSRP clustering.

If the implementation is desired, tracking of a set of internal and external IP addresses are recommended.

An Example of Resilience Network Design



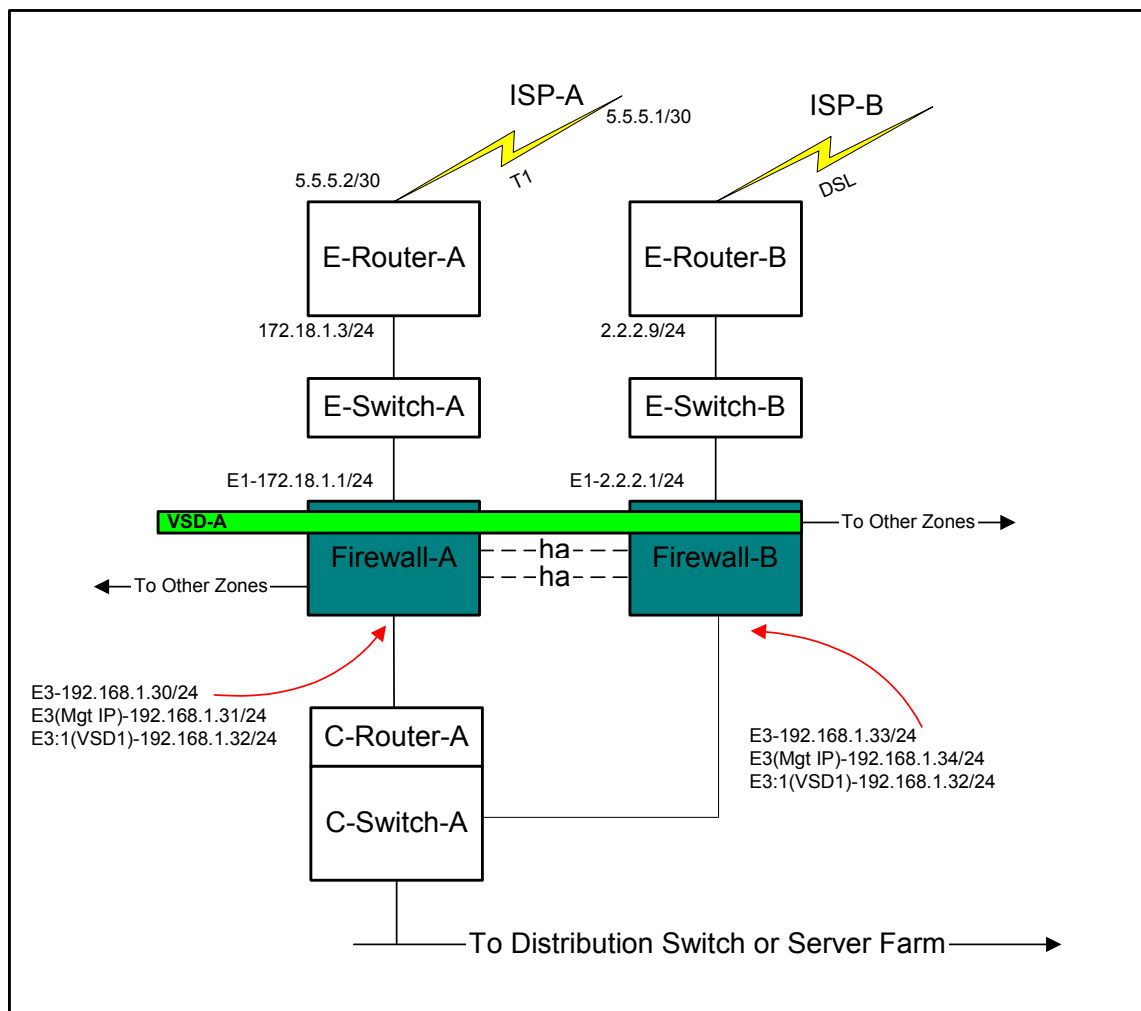
NSRP Clustering and Non Resilience Network Design

Using NSRP clustering together with track IP in a non-resilience network design may be useful in helping to restore either partial or complete network connectivity in the case of certain component failure.

A case study is presented for the discussion in regard to the use of path monitoring. The requirement is to provide partial network connectivity in the case of the primary Internet link failed.

Case study - Partial resilience network and NSRP Active-passive configuration

Network diagram



Note that Firewall-A is the active firewall for VSD-A and Firewall-B is the backup firewall for VSD-A

Implementation checklist,

Items	Requirement
Primary Internet link failure tracking	Yes
Backup Internet link failure tracking	No
NSRP clustering mode	Active-passive
NSRP path monitoring on primary link	Yes
NSRP path monitoring on backup link	No
NSRP preempt	Yes
Number of VSD group require	1
NSRP link status monitoring	Interfaces belong to trusted zone
Branch VPN redundancy	Yes, with redundant VPN configuration
Avoid both Netscreen in inoperable state	Yes
Restoring outgoing traffic	Yes
Restoring incoming traffic	Partial due to the routing at ISP
Routing strategy	Static
Number of virtual router	1

Configuration note,

1. Unset VSD group zero. This makes the main interface local and the configuration on the main interface will not synchronized with the master or vice versa.
2. Track IP is not configured on the backup Netscreen.
3. The default gateway must be configured using the set interface gateway command.
4. Manage IP must be configure in order for the track IP to work.
5. Select the outgoing interface for track IP. The default setting is auto.
6. The tracked IP is the remote point-to-point WAN IP address and the firewall default gateway IP.

Configuring active-passive NSRP

Configuration on the master Netscreen Firewall

----- Assigning interface into zones -----

```
set interface "ethernet1" zone "Untrust"  
set interface "ethernet2" zone "DMZ"  
set interface "ethernet3" zone "Trust"
```

----- Unset the default VLAN1 IP address -----

```
unset interface vlan1 ip
```

----- Generic NSRP configuration -----

```
set nsrp cluster id 2  
unset nsrp vsd-group id 0 -----> Make VSD id zero local  
set nsrp vsd-group id 1 priority 5  
set nsrp vsd-group id 1 preempt  
set nsrp monitor interface ethernet1  
set nsrp monitor interface ethernet3
```

----- Assigning interfaces IP address and subnet mask for VSD 0 -----

```
set interface ethernet1 ip 172.18.1.1/24  
set interface ethernet1 route  
set interface ethernet3 ip 192.168.1.30/24  
set interface ethernet3 route
```

----- Assigning interfaces IP address and subnet mask for VSD 1 -----

```
set interface ethernet3:1 ip 192.168.1.32/24  
set interface ethernet3:1 route
```

----- Default gateway -----

```
set interface ethernet1 gateway 172.18.1.3
```

----- Assigning manage IP to interfaces -----

```
set interface ethernet1 manage-ip 172.18.1.2  
set interface ethernet3 manage-ip 192.168.1.31
```

----- NSRP track IP configuration -----

```
set nsrp track-ip ip  
set nsrp track-ip threshold 9  
set nsrp track-ip ip 5.5.5.1 interface ethernet1  
set nsrp track-ip ip 5.5.5.1 weight 10  
set nsrp track-ip ip 172.18.1.3 weight 10
```

----- Policy configuration -----

```
set policy id 0 from "Trust" to "Untrust" "Any" "Any" "ANY" nat dip-id 2 Permit
```

Configuration on the backup Netscreen Firewall

----- Assigning interface into zones -----

```
set interface "ethernet1" zone "Untrust"  
set interface "ethernet2" zone "DMZ"  
set interface "ethernet3" zone "Trust"
```

----- Unset the default VLAN1 IP address -----

```
unset interface vlan1 ip
```

----- Generic NSRP configuration -----

```
set nsrp cluster id 2  
unset nsrp vsd-group id 0 -----> Make VSD id zero local  
set nsrp vsd-group id 1 priority 100  
set nsrp monitor interface ethernet1  
set nsrp monitor interface ethernet3
```

----- Assigning interfaces IP address and subnet mask for VSD 0 -----

```
set interface ethernet1 ip 2.2.2.1/24  
set interface ethernet1 route  
set interface ethernet3 ip 192.168.1.33/24  
set interface ethernet3 route
```

----- Assigning interfaces IP address and subnet mask for VSD 1 -----

```
set interface ethernet3:1 ip 192.168.1.32/24  
set interface ethernet3:1 route
```

----- Default gateway -----

```
set interface ethernet1 gateway 2.2.2.9
```

----- Assigning manage IP to interfaces -----

```
set interface ethernet1 manage-ip 2.2.2.2  
set interface ethernet3 manage-ip 192.168.1.34
```

----- NSRP track IP configuration -----

Not configured

----- Synchronizing the configuration with the master Netscreen -----

```
exec nsrp syn global save
```

A reset is required after running the command above.

Verification

Verify the interface binding to VSD, note that interface eth3:1 is the active interface for VSD 1.

VSD-1-MASTER(M)-> get interf

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

Interfaces in vsys Root:

Name	IP Address	Zone	MAC	VLAN	State	VSD
eth1	172.18.1.1/24	Untrust	0010.db19.ac14	-	U	-
eth2	0.0.0.0/0	DMZ	0010.db19.ac15	-	U	-
eth3	192.168.1.30/24	Trust	0010.db19.ac16	-	U	-
eth3:1	192.168.1.32/24	Trust	0010.dbff.4061	-	A	1
eth4	0.0.0.0/0	Null	0010.db19.ac17	-	D	-
eth5	0.0.0.0/0	Null	0010.db19.ac18	-	D	-
eth6	0.0.0.0/0	Null	0010.db19.ac19	-	D	-
eth7	0.0.0.0/0	Null	0010.db19.ac1a	-	D	-
eth8	0.0.0.0/0	HA	0010.db19.ac1b	-	U	-
vlan1	0.0.0.0/0	MGT	0010.db19.ac1f	1	I	-
v1-trust	0.0.0.0/0	V1-Trust	0010.db19.ac1f	-	D	-
v1-untrust	0.0.0.0/0	V1-Untrust	0010.db19.ac1f	-	D	-
v1-dmz	0.0.0.0/0	V1-DMZ	0010.db19.ac1f	-	D	-

VSD-1-MASTER(M)-> get nsrp

nsrp version: 2.0

cluster info:

cluster id: 2, no name

local unit id: 1682448

active units discovered:

index: 0, unit id: 1682448, ctrl mac: 0010db19ac1b, index: 1, unit id: 19085

28, ctrl mac: 0010db1d1f3b, data mac: ffffffff

total number of units: 2

VSD group info:

init hold time: 5

heartbeat lost threshold: 3

heartbeat interval: 1000(ms)

group	priority	preempt	holddown	inelig	master	PB	other members
1	5	yes	3	no	myself	1908528	

total number of vsd groups: 1

Total iteration=16858,time=120266128,max=18147,min=268,average=7134

RTO mirror info:

run time object sync: disabled

ping session sync: enabled

coldstart sync done

nsrp link info:

control channel: ethernet8 (ifnum: 11) mac: 0010db19ac1b state: up

ha data link not available

ha secondary path link not available

NSRP encryption: disabled

NSRP authentication: disabled

NSRP monitor interface: Ethernet1 ethernet3

number of gratuitous arps: 4 (default)

track ip: enabled

VSD-1-MASTER(M)-> get nsrp track

ip address	interval	threshold	wei	interface	method	fail-count	success-rate
5.5.5.1	1	3	10	ethernet1	ping	0	98%

0 ip(s) failed, weighted sum = 0. fail over threshold is set to 9.

VSD-1-BACKUP(B)-> get interf

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

Interfaces in vsys Root:

Name	IP Address	Zone	MAC	VLAN	State	VSD
eth1	2.2.2.1/24	Untrust	0010.db1d.1f34	-	U	-
eth2	0.0.0.0/0	DMZ	0010.db1d.1f35	-	U	-
eth3	192.168.1.33/24	Trust	0010.db1d.1f36	-	U	-
eth3:1	192.168.1.32/24	Trust	0010.dbff.4061	-	I	1
eth4	0.0.0.0/0	Null	0010.db1d.1f37	-	D	-
eth5	0.0.0.0/0	Null	0010.db1d.1f38	-	D	-
eth6	0.0.0.0/0	Null	0010.db1d.1f39	-	D	-
eth7	0.0.0.0/0	Null	0010.db1d.1f3a	-	D	-
eth8	0.0.0.0/0	HA	0010.db1d.1f3b	-	U	-
vlan1	0.0.0.0/0	MGT	0010.db1d.5f3f	1	I	-
v1-trust	0.0.0.0/0	V1-Trust	0010.db1d.1f3f	-	D	-
v1-untrust	0.0.0.0/0	V1-Untrust	0010.db1d.1f3f	-	D	-
v1-dmz	0.0.0.0/0	V1-DMZ	0010.db1d.1f3f	-	D	-

VSD-1-BACKUP(B)->

VSD-1-BACKUP(B)-> get nsrp
nsrp version: 2.0

cluster info:

cluster id: 2, no name

local unit id: 1908528

active units discovered:

index: 0, unit id: 1908528, ctrl mac: 0010db1d1f3b, index: 1, unit id: 16824
48, ctrl mac: 0010db19ac1b, data mac: ffffffff

total number of units: 2

VSD group info:

init hold time: 5

heartbeat lost threshold: 3

heartbeat interval: 1000(ms)

group	priority	preempt	holddown	inelig	master	PB	other members
1	100	no	3	no	1682448	myself	

total number of vsd groups: 1

Total iteration=16040,time=110602505,max=17423,min=3170,average=6895

RTO mirror info:

run time object sync: disabled

ping session sync: enabled

coldstart sync done

nsrp link info:

control channel: ethernet8 (ifnum: 11) mac: 0010db1d1f3b state: up

ha data link not available

ha secondary path link not available

NSRP encryption: disabled

NSRP authentication: disabled

NSRP monitor interface: ethernet1 ethernet3

number of gratuitous arps: 4 (default)

track ip: disabled

VSD-1-BACKUP(B)->

Appendix

Known Issues

There are several known issues to address,

1. Static routes configuration which use local interfaces are synchronizing within the Netscreen cluster.
2. If the outgoing interface selected is Auto, which is the default, how would the Netscreen decide which interface to use?
3. Implementing track IP will become extremely difficult when the NSRP clustering is configured in active-active mode and custom zone are being used. The main challenge is to avoid the Netscreen in the cluster to become inoperable when the tracked IP failed. For example, if the Engineering Group are in the Engineering security zone, Finance Group is in the Finance security zone. These security zones will not be able to communicate with each other when all of the Netscreen in the cluster become inoperable even though all network connectivity between these zones are functional.

Multiple tracked IP can be configured on the Netscreen to avoid all Netscreen within the cluster to become inoperable when a single tracked IP failed. But, in most cases this will make the desire fail over not possible to trigger.

A detection mechanism should be implemented to avoid the above side effect to happen. For example, for a user configurable period of time, if all of the Netscreen within the cluster recognized that all of Netscreen are in inoperable state, at least one of the Netscreen in the cluster should be promoted to be the master for all of the VSD.

References

Netscreen Concept and Example NSRP Guide for Screen OS 4.0