



Redundant Small Site Multi-Homing

By Ivan Pepelnjak


1. 7. 2007

The IP Corner article *Small Site Multihoming* has generated lots of responses, most of them being questions about redundant implementation of the same principles. In this article, we'll thus extend the small site multi-homing design with a set of redundant routers. The final design will still retain the administrative simplicity of the original solution – with no need to own public IP address space, autonomous system number or to run Border Gateway Protocol (BGP).

Redundant Design

The basic design of a redundant multi-homed small site is very similar to the one described in IP Corner article *Small Site Multihoming* (I would strongly suggest you read that article before this one). The public IP addresses used by the site are still within the address space of the two service providers (Figure 1), and the private IP addresses are used within the site. Each gateway router performs independent Network Address Translation (NAT) from the private IP addresses to the public IP address pool (or a single IP address) assigned by the ISP.


Figure 1: IP addressing in small multi-homed site



Static default routes are configured on both gateway routers. The availability of these routes depends on tracked objects configured on each router (see the section *Not-so-Very-Static-Routes* in the Small Site Multihoming article). The static default routes are redistributed into a dynamic routing protocol to ensure that both gateway routers (as well as any other additional router within the site) have the same view of the Internet connectivity (Figure 2).


Note: In the primary/backup scenario, the backup static default route configured on the gateway routers is a floating static route to ensure the backup path will not be used as long as the primary path is available.

Figure 2: Static default routing



Due to NAT being performed by the gateway routers, the return traffic is always handled properly regardless of the path the outgoing packet is taking. For example, if the outgoing packet is forwarded by GW B, the NAT performed by the gateway router would replace the source IP address with the IP address assigned by ISP B; the return packet would thus automatically take the path through ISP B (Figure 3).

Figure 3: Symmetrical routing with NAT



Although NAT solves the return path problem, the solution is not perfect. For example, if the traffic is flowing over the backup link and the primary link is reestablished, all the traffic will be shifted to the primary link (regardless of the TCP session status), resulting in a different public source IP address, and the workstations will lose all TCP sessions established at the switchover moment.

Configuring Internet Routing

The configuration of the gateway routers follow the principles explained in the Small Site Multihoming article. IP addressing and NAT is configured on both gateway routers, as shown in Listing 1 (only GW-A configuration is included in most examples).

Listing 1: IP addressing, DHCP and NAT configuration

```
hostname GW-A

!

ip cef

!

ip dhcp pool LAN
    network 192.168.0.0 255.255.255.0
    default-router 192.168.0.1
!

ip dhcp excluded-addresses 192.168.0.1 192.168.0.10
ip dhcp excluded-addresses 192.168.0.128 192.168.0.255
!

interface FastEthernet0/0
    description *** Inside LAN interface ***
    ip address 192.168.0.1 255.255.255.0
    ip nat inside
!

interface Serial0/0/0
    description *** Link to ISP A ***
    ip address 172.16.1.1 255.255.255.252
```

```

ip nat outside
!
ip nat inside source route-map ISP_A interface Serial0/0/0 overload
!
route-map ISP_A permit 10
    match interface Serial0/0/0

```

*Note: To increase the overall reliability, the DHCP server is running on both gateway routers. Use the **ip dhcp excluded-addresses** configuration commands to ensure the routers always allocate addresses from non-overlapping pools.*

To implement reliable static routes on both gateway routers, you have to configure:

- An IP SLA object to track end-to-end connectivity to an IP address that is “far enough” (at least within the core of the ISP network, tracking a server of an upstream ISP is even better).
- A **track** object that monitors the state of the IP SLA object.
- A default route that is inserted in the IP routing table based on the state of the **track** object.
- Local policy routing to ensure that the IP SLA measurements always use the Internet interface (otherwise a gateway router with failed upstream link might use the default path provided by the other gateway router for its SLA measurements).

The relevant parts of GW-A configuration are included in Listing 2 (and the detailed description of the configuration and monitoring commands related to reliable static routing is available in the [Small Site Multihoming](#) article). The only major difference between GW-A and GW-B is the default route configuration, where you would use a high administrative on the backup router (GW-B in our example) to make the default route floating; if Internet connectivity on GW-A is operational, the default route received through the routing protocol should override the static default route.

Listing 2: Basic multi-homed default routing setup

```

hostname GW-A
!

```

```

ip sla 15

  icmp-echo 172.29.0.1 source-interface Serial0/0/0

  timeout 200

  frequency 10

!

ip sla schedule 15 life forever start-time now

!

track 17 rtr 15 reachability

  delay down 10 up 20

!

ip local policy route-map LocalPolicy

!

ip access-list extended PingISP_A

  permit icmp host 172.16.1.1 host 172.29.0.1

!

route-map LocalPolicy permit 10

  match ip address PingISP_A

  set interface Serial0/0/0

!

ip route 0.0.0.0 0.0.0.0 Serial0/0/0 10 name ISP_A track 17

```

Configuring Intra-Site Routing

The static default route configured on GW-A and GW-B has to be propagated between them to ensure that both routers have the same view of the Internet connectivity. The easiest way to implement this requirement is to redistribute the static default route into a dynamic routing protocol configured between the two routers, as shown in Listing 3.

Listing 3: Redistributing the static default route

```

router ospf 1

  redistribute static metric 10

  default-information originate

!

```




```
interface FastEthernet0/0
  ip ospf 1 area 0
```

*Note: OSPF will not announce the redistributed default route until you configure **default-information originate** within the OSPF process.*

If there are no workstations attached to the LAN between GW-A and GW-B, we're finished; all routers attached to that LAN will get the default route pointing to the currently-active gateway router through a dynamic routing protocol (Figure 4).

Figure 4: Intra-site routing in a large site



Intra-site routing with workstations attached to the same LAN as GW-A and GW-B is a bit more complex. You can usually configure only a single default gateway on the workstations, so you have to provide dynamic switchover of the default gateway with Hot Standby Router Protocol (HSRP) or an equivalent (for example, Gateway Load Balancing Protocol; GLBP, or Virtual Router Redundancy Protocol; VRRP). The configuration is very simple, since the **track** objects that you can use to adjust the router's HSRP priority based on the state of the upstream link have already been configured (see Listing 4 and Listing 5; the only difference is the default HSRP priority, which is higher on GW A).

Note: HSRP and VRRP address only the needs of primary/backup router scenarios, whereas GLBP supports outbound load balancing. As our design

does not address load balancing toward the Internet, HSRP nicely fits the requirements.

Listing 4: HSRP configuration on GW A

```
interface FastEthernet0/0

ip address 10.0.0.5 255.255.255.0

standby 1 priority 95

standby 1 ip 10.0.0.1

standby 1 preempt

standby 1 track 17 decrement 20
```

Listing 5: HSRP configuration on GW B

```
interface FastEthernet0/0

ip address 10.0.0.6 255.255.255.0

standby 1 ip 10.0.0.1

standby 1 priority 90

standby 1 preempt

standby 1 track 17 decrement 20
```

Summary

The design described in this article gives you the ability to implement fully redundant Internet connectivity without having an allocated public IP address space or autonomous system number. Even better, it's completely static on the Internet side, thus alleviating the need to configure BGP on the gateway routers. However, the simplicity of the design brings a few drawbacks as well; you should use this design only in a stable environment where the switchover from primary to backup ISP is unlikely (but you still need the secondary connection to ensure reliability), as every switchover will cause all established TCP sessions to be terminated.

The article focused solely on the primary/backup scenario. It's possible to extend it to support rudimentary load sharing, but you have to be careful to make certain that all the IP packets between a pair of inside/outside hosts will always flow across the same gateway router (otherwise the NAT

configured on the gateway router will destroy the TCP session). Similarly, it's possible (although not trivial) to implement publicly accessible inside servers; this topic will be covered in a later IP Corner article.

NIL – More Than Just a Training Company

NIL Learning delivers the leading-edge Cisco training to IT professionals and companies around the globe. Through field-proven experts — each both active engineer and instructor — NIL Learning enhances the standard learning curriculum with real-life experience and helps clients to maximize their training investment.

NIL Learning is part of NIL, a leading global IT solutions provider. Since 1992, NIL has been at the forefront of advanced contributors to strategic partner Cisco's technologies, learning curriculum and value-added solutions deployed to clients around the globe. Today, NIL has earned the highest certifications offered by Cisco, VMware, EMC, HP, IBM, Microsoft, F5, Jive, MobileIron, RSA, VCE and others. Their portfolio of solutions consists of managed services, professional services and learning services.

NIL is headquartered in Slovenia, with regional offices in Croatia, Serbia, Saudi Arabia, the U.S., Turkey, South Africa, Morocco, Nigeria, Kenya and Botswana.

Why learn at NIL LEARNING?

- All NIL LEARNING instructors are **field-proven experts** - each both active engineer and instructor.
- **75% of NIL LEARNING engineers hold CCSI certifications, and 18 have already achieved the respected CCIE rank.**
- NIL LEARNING **enhances the standard learning curriculum** with real-life experience and helps clients to maximize their training investment.
- NIL has been a Cisco Training Partner for many years; it became a **Cisco Learning Partner in 1993**, and has been a **Cisco Gold Partner since 1995**.

- NIL was awarded the Cisco Most Business Relevant Learning Partner in MEA in 2010 and the most innovative learning partner in MEA.
- NIL received the Innovation Award for its Technology Led Training and its extensive contribution to Cisco learning solutions at the Cisco EMEAR Learning Partner Summit in 2012.
- NIL received the Innovation Award for its Technology Led Training and Advanced Engineer Program at the Cisco Global Learning Partner Summit in 2013.
- NIL LEARNING runs a centralized training schedule across the whole EMEAR region.

More Info

Slovenia

T: +386 1 4746 500

E: sales-support@nil.com

Saudi Arabia

T: +966 1 465 4641

E: info.nilme@nil.com

Botswana

T: +267 318 1684

E: training@it-iq.bw

Serbia

T: +381 11 2282 818

E: info-nilserbia@nil.co.rs

Croatia

T: +385 (0)51 583 255

E: info-nilcroatia@nil.com

South Africa

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Kenya

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Turkey

T: +902 123 81 8639

E: info-nilturkey@nil.com

Morocco

T: +212(0) 660 808 394

E: info-nilmorocco@nil.com

USA

T: +1 612 886 3900

E: info-nilusa@nil.com

Nigeria

T: +27 (0)11 575 4637

E: mea_sales@nil.com

www.learning.nil.com

