Small Site Multi-Homing

By Ivan Pepelnjak, Dragan Spasojević and Damijan Markovič

1. 2. 2007

Unless your network operates under extreme security considerations or in places where the is no Internet access, your management has probably already asked you to lower the wide area network (WAN) costs by migrating from a traditional leased line or frame-relay-based network to an Internet- or MPLS VPN-based transport, while at the same time retaining or even increasing network reliability. This conflicting set of requirements might force you to make all your sites multi-homed (connected to more than one Internet Service Provider – ISP).

Multi-homing requirements aren't new; for example, every decent e-commerce solution should be multi-homed. However, most solutions you'll find with extensive help from Google require that you possess your own public IP prefix and an autonomous system number (both of them are scarce resources) and run Border Gateway Protocol (BGP) with all ISPs. Clearly, these requirements are completely unrealistic if you want to multi-home a small remote office.

In this article, you'll learn how to:

- connect a small remote site to more than one ISP;

- detect failures in the ISP networks and adjust your routing accordingly;
- increase overall availability of your sites with Service Level Agreement (SLA) monitoring;
- log all relevant changes in the remote site connectivity.

Basic Small Site Multi-Homing

Connecting a small site to multiple service providers can be extremely easy – you get two upstream links and two provider-assigned (PA) IP addresses (either static or dynamically assigned). Since each ISP will give you only a single IP address, you have to use private IP addresses on the LAN side of the router (Figure 1).

Figure 1: IP addressing in small multi-homed site

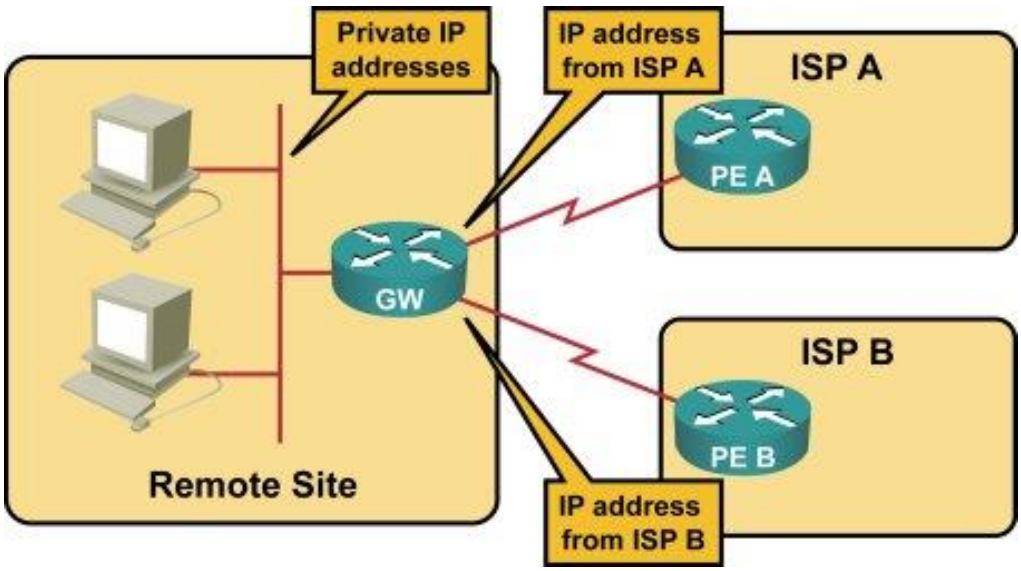

As most ISPs will not be willing to run a dynamic routing protocol with small sites, you have to configure static default routing on your end. You would almost always prefer one provider over the other, resulting in a primary and a backup default route (Figure 2).

Note: With careful configuration, it's also possible to achieve rudimentary load sharing with two equally-good default routes.

NIL LEARNING ©2014 All rights reserved.

Figure 2: Static default routing

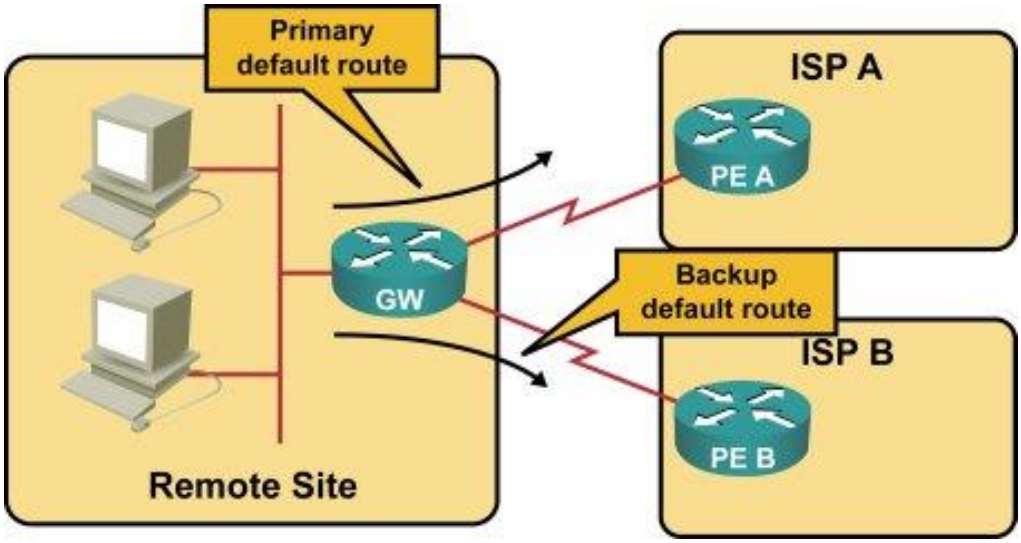

The router on the remote site would also have to perform two independent NAT translations, one for packets sent to ISP A (where local addresses get translated to the IP address assigned by ISP A) and another one for packets sent to ISP B (Figure 3).

Figure 3: NAT translation in small multi-homed site

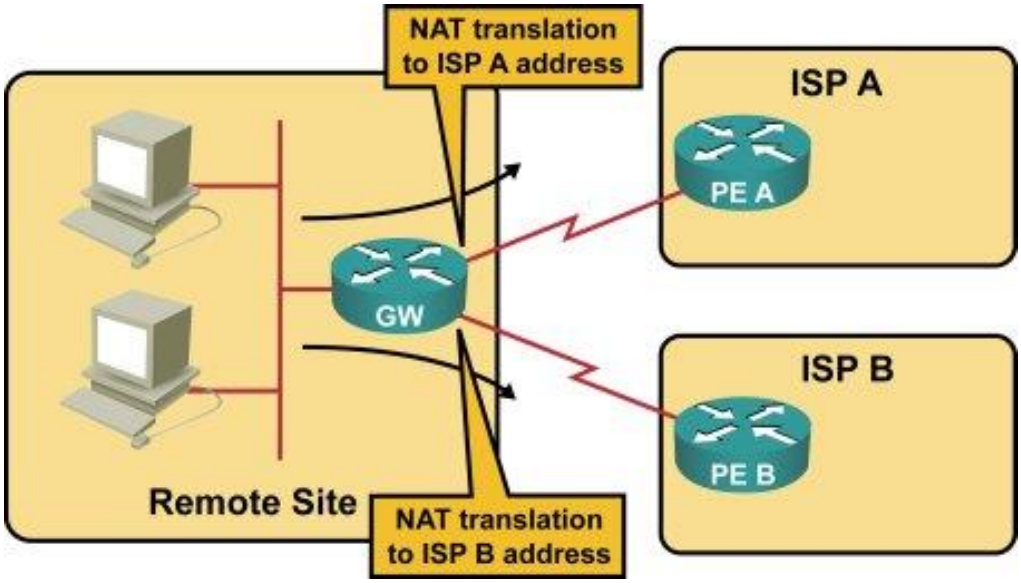

One of the major issues in multi-homed site design is the proper handling of the return traffic. It's not uncommon to experience performance problems if the outbound and return traffic flow over different links (also known as asymmetrical routing), while IP multicast and stateful packet inspection (part of IOS firewall feature set) almost always break under these conditions. Fortunately, asymmetrical routing is never a problem in a

NIL LEARNING ©2014 All rights reserved.

dual NAT design from Figure 3, as the source address of the outbound packet indicates the link that has been used to send it (see Figure 4).

Figure 4: Symmetrical routing with dual NAT

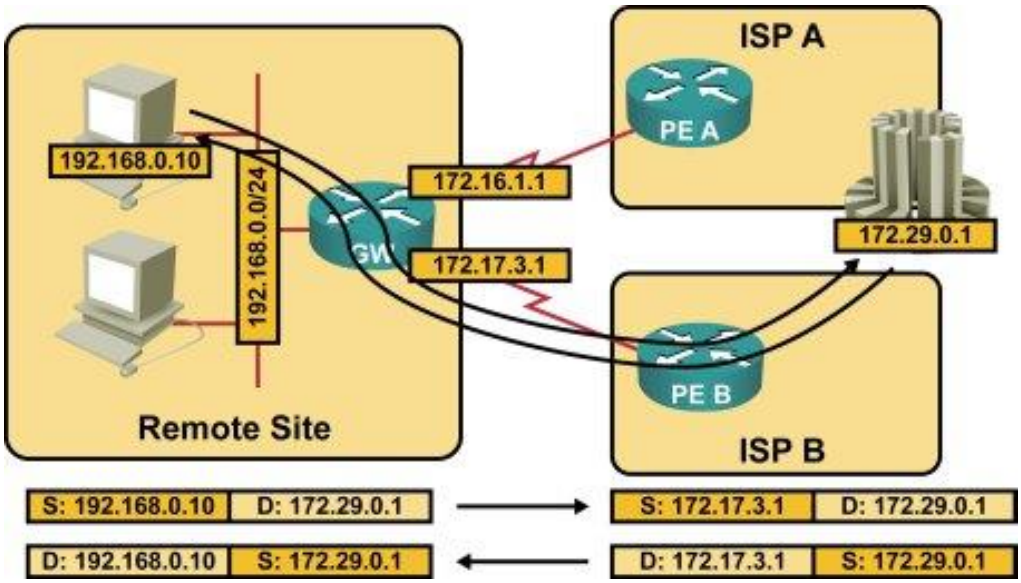

Configuring Small Multi-Homed Site

Configuring the gateway router in a small multi-homed site is very simple. You start by configuring the private and public IP addresses (Listing 1).

Listing 1: Initial router configuration

```
hostname GW
!
ip cef
!
ip dhcp pool LAN
   network 192.168.0.0 255.255.255.0
   default-router 192.168.0.1
!
interface FastEthernet0/0
 description *** Inside LAN interface ***
 ip address 192.168.0.1 255.255.255.0
!
interface Serial0/0/0
```

©2014 All rights reserved.

```
 description *** Link to ISP 1 ***
 ip address 172.16.1.1 255.255.255.252
!
interface Serial0/0/1 point-to-point
 description *** Link to ISP 2 ***
 ip address 172.17.3.1 255.255.255.252
```

NAT configuration is a bit more complex; you have to configure two NAT pools (one for each ISP), as displayed in Listing 2.

Listing 2: Network Address Translation configuration

```
interface FastEthernet0/0
 ip nat inside
!
interface Serial0/0/0
 ip nat outside
!
interface Serial0/0/1 point-to-point
 ip nat outside
!
ip nat inside source route-map ISP_A interface Serial0/0/0 overload
ip nat inside source route-map ISP B interface Serial0/0/1 overload
!
route-map ISP_A permit 10
 match interface Serial0/0/0
!
route-map ISP_B permit 10
 match interface Serial0/0/1
```

*Note: Having two **route-maps** matching outgoing interfaces (the **match interface** statement in a NAT **route-map** matches outgoing interface) is the only way to configure per-interface NAT pools in Cisco IOS.*

As most ISPs will not be willing to run a dynamic routing protocol with small sites, you have to configure static default routing on your end. You would almost always prefer one provider over the other (therefore one default route would have a lower administrative distance) as shown in

NIL LEARNING ©2014 All rights reserved.

Listing 3, although it's possible (with CEF switching using per-destination load sharing) to use two default routes in 1-to-1 load-balancing setup.

Listing 3: Basic multihomed default routing setup

```
ip route 0.0.0.0 0.0.0.0 Serial0/0/0 10

ip route 0.0.0.0 0.0.0.0 Serial0/0/1 251
```

The simplistic static routing in Listing 3 represents a major availability issue – if you cannot detect the link failure on the link toward ISP A reliably, the default static route toward ISP B will never be used. While you can almost always detect leased-line or cable failure (due to loss of carrier signal) and usually detect Frame-Relay failures through Local Management Interface (LMI) messages or end-to-end keepalives, it's almost impossible to detect layer-2 failures in PPPoE (ADSL) or Metro Ethernet access layers. In these scenarios, the primary default route will never disappear (even though the next-hop router is no longer reachable), making static multi-homing impossible. This problem is solved, however, in Cisco IOS release 12.3(8)T (integrated in release 12.4) with static routes tied to IP SLA measurements.

Not-so-Very-Static Routes

Cisco IOS release 12.3(4)T introduced *Enhanced Object Tracking*, which together with *Reliable Static Routing Using Object Tracking* introduced in IOS release 12.3(8)T solves the problem. *Enhanced Object Tracking* introduces a generic **track** object that can track a state of an interface (layer-2 or layer-3 state), presence or metric of an IP route, state of an SLA measurement or even availability of Mobile IP home agent or GPRS nodes. Even more, you can combine various track objects (including weighing them) into a compound object.

The *Reliable Static Routing Using Object Tracking* feature ties a **track** object to a static route – whenever the **track** object's state is *down*, the static route is removed from the routing table; exactly what you would need to support reliable multi-homing. To configure a static route based on the state of the next-hop router, you need to:

Configure an **ip sla** (previously known as Response Time Reporter – **rtr**) object pinging the next-hop router on primary Internet link (Listing 4). The

NIL LEARNING ©2014 All rights reserved.

polling **frequency** you specify (in seconds) depends on the reliability requirements, but anything below a few seconds would place unnecessary burden on the next-hop router (as you might not be the only one tracking its availability).

Listing 4: Pinging next-hop router

```
ip sla 100
 icmp-echo 172.16.1.2 source-interface Serial0/0/0
 timeout 500
 frequency 3
ip sla schedule 100 life forever start-time now
```

Note: You cannot change the parameters of an SLA object once you've scheduled it. To change the target IP address, timeouts or polling frequency, you need to delete the SLA object and recreate it.

Create a **track** object monitoring the reachability of the SLA target (Listing 5). As you probably don't want to respond to a single lost ICMP packet, you should use the **delay** option of the **track** object to specify how long the next-hop router should remain unreachable before it's declared to be lost (the **down** delay should be approximately three times the SLA polling frequency and the **up** delay should be even longer).

*Note: When calculating the **up** delay, remember that a router can temporarily respond to pings during the bootstrap process.*

Listing 5: Tracking the state of the next-hop router

```
track 100 rtr 100 reachability
 delay down 10 up 20
```

After configuring the **track** object, attach it to the primary static default route to ensure that the default route is removed if the next-hop router is not reachable (Listing 6).

Listing 6: Conditional static default route

```
ip route 0.0.0.0 0.0.0.0 Serial0/0/0 10 track 100
ip route 0.0.0.0 0.0.0.0 Serial0/0/1 251
```

You can check the proper operation of the reliable static routing with the **show ip route** command. Listing 7 displays the IP routing table on the GW

NIL LEARNING ©2014 All rights reserved.

router when the primary next-hop router is available, Listing 8 shows the routing table after primary next-hop router failure.

Listing 7: IP routing table with operational primary next-hop router

```
GW#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
     172.17.0.0 255.255.255.252 is subnetted, 1 subnets
C       172.17.3.0 is directly connected, Serial0/0/1
     172.16.0.0 255.255.255.252 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, Serial0/0/0
C    192.168.0.0 255.255.255.0 is directly connected, FastEthernet0/0
S*   0.0.0.0 0.0.0.0 is directly connected, Serial0/0/0
```

Listing 8: IP routing table after the next-hop router failure

```
GW#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
     172.17.0.0 255.255.255.252 is subnetted, 1 subnets
C       172.17.3.0 is directly connected, Serial0/0/1
     172.16.0.0 255.255.255.252 is subnetted, 1 subnets
C       172.16.1.0 is directly connected, Serial0/0/0
C    192.168.0.0 255.255.255.0 is directly connected, FastEthernet0/0
S*   0.0.0.0 0.0.0.0 is directly connected, Serial0/0/1
```

Monitoring Reliable Static Routing

The reliable static routes silently appear or disappear from the IP routing table based on the state of the attached **track** object; the only means of monitoring their state is with the **show ip route track-table** command (Listing 9) or with the **debug track** command (Listing 10).

Listing 9: Show tracked routes

```
GW#show ip route track-table
 ip route 0.0.0.0 0.0.0.0 Serial0/0/0 10 name ISP_A track 100 state is [down
]
```

©2014 All rights reserved.

Listing 10: Debug tracking

```
GW#debug track

06:49:44: Track: 100 Down change delayed for 10 secs

06:49:54: Track: 100 Down change delay expired

06:49:54: Track: 100 Change #26 rtr 100, reachability Up->Down

06:50:24: Track: 100 Up change delayed for 20 secs

06:50:34: Track: 100 Up change delay cancelled

06:58:59: Track: 100 Up change delayed for 20 secs

06:59:19: Track: 100 Up change delay expired

06:59:19: Track: 100 Change #25 rtr 100, reachability Down->Up
```

Note: The debugging printout in Listing 10 illustrates a real-life scenario where the next-hop router became temporarily reachable during the bootstrap process and disappeared a few seconds later (the change delay cancelled printout).

While the silent modification of the IP routing table might be acceptable in most situations (after all, you don't get notified when a regular IP route disappears from the routing table either), you might want to know if your primary ISP is unreachable (similar to the interface up/down events you would get with traditional access methods like leased lines or Frame Relay access). The *Embedded Event Manager 2.2* (introduced in IOS release 12.4(2)T) is the ideal solution, as you can trigger EEM applets (or TCL scripts) whenever a **track** object's state changes with the **event track** configuration command.

To display the changes in a tracked object state, you can define two EEM applets, one triggered on the **down** change, another one triggered on the **up** change. If you only want to be notified that the state has changed, the only **action** you need to specify is the **syslog msg** action, but you can perform any number of actions you want (for example, send an e-mail to the network manager or even reconfigure the router). A sample EEM configuration is shown in Listing 11 and the printouts generated by it are included in Listing 12.

NIL LEARNING ©2014 All rights reserved.

Listing 11: IOS EEM generates syslog messages on tracked object state change

```
event manager applet ISP_A_down
 event track 100 state down
 action 1.0 syslog msg "ping to 172.16.1.2 from Serial 0/0/0 failed"
event manager applet ISP_A_up
 event track 100 state up
 action 1.0 syslog msg "172.16.1.2 is reachable"
```

Listing 12: Sample EEM printouts

```
07:02:19: %HA_EM-6-
LOG: ISP_A_down: ping to 172.16.1.2 from Serial 0/0/0 failed
07:03:19: %HA_EM-6-LOG: ISP_A_up: 172.16.1.1 is reachable
```

End-to-End Connectivity Test

After you've successfully implemented the tracking of the primary next-hop router's availability, you might be tempted to improve the solution to track end-to-end connectivity through ISP A and switch to the backup ISP whenever your central site is not reachable through the primary ISP. In theory, the required configuration change should be minimal – you only have to change the destination IP address in the IP SLA definition (Listing 13).

Listing 13: Pinging a remote host

```
hostname GW
!
ip sla 100
 icmp-echo 172.29.0.1 source-interface Serial0/0/0
 timeout 200
 frequency 10
ip sla schedule 100 life forever start-time now
```

In most cases, that's all you have to do. As the ICMP echoes sent to the central site come from an IP address belonging to ISP A (the IP address configured on Serial 0/0/0 in the example), it's highly unlikely that you would get a return packet if the ISP A has problems. However, the return

NIL LEARNING ©2014 All rights reserved.

packet might still reach your router under rare circumstances (misconfigured access lists or one-way connectivity in ISP A). The results are astonishing:

As the pings through ISP A (primary default route) fail, the router removes the primary default route and the backup default route through ISP B is installed.

Pings are now sent from an IP address belonging to ISP A on a path going through ISP B.

If there is a return path from the central site to the IP address sending the ICMP packets, the central site will yet again appear reachable and the primary default route will be reinstalled (resulting in connectivity loss).

Due to renewed connectivity loss, the router will oscillate between the two default routes (Listing 14).

Listing 14: Oscillating routing

```
GW#debug track

07:15:09: Track: 100 Change #32 rtr 100, reachability Up->Down

07:15:09: %HA_EM-6-
LOG: ISP_1_down: ping to 172.29.0.1 from Serial 0/0/0 failed

07:15:19: Track: 100 Up change delayed for 20 secs

07:15:39: Track: 100 Up change delay expired

07:15:39: Track: 100 Change #33 rtr 100, reachability Down->Up

07:15:39: %HA_EM-6-LOG: ISP_1_up: 172.29.0.1 is reachable

07:15:49: Track: 100 Change #34 rtr 100, reachability Up->Down

07:15:49: %HA_EM-6-
LOG: ISP_1_down: ping to 172.29.0.1 from Serial 0/0/0 failed

07:15:59: Track: 100 Up change delayed for 20 secs
```

To fix this (admittedly rare) problem you have to configure a local policy routing (as the **ip sla** packets originate within the router, they are only affected by the **ip local policy**) that matches ICMP packets being sent from the Serial0/0/0 interface (based on their IP address; the *PingISP_A* access list) and forces them to be sent out through the same interface with the **set interface** configuration command (Listing 15).

NIL LEARNING ©2014 All rights reserved.

Listing 15: Fix the oscillating routing with local policy

```
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
```

If you want to, you can extend the concepts presented in this section even further. For example, if the central site is not reachable through either ISP (it might be down), it could make more sense to retain ISP A as the primary ISP. You would thus need to track the central site's availability through both ISPs and configure a reliable static default route for both of them (the backup one with a higher administrative distance, of course) with a third (last-resort) default route pointing to ISP A. The complete configuration is included in Listing 16 and its interpretation is left as an exercise for the reader.

Listing 16: GW router tracking central site availability through both ISPs

```
hostname GW
!
ip cef
!
ip dhcp pool LAN
   network 192.168.0.0 255.255.255.0
   default-router 192.168.0.1
!
ip sla 100
 icmp-echo 172.29.0.1 source-interface Serial0/0/0
 timeout 200
 frequency 3
ip sla schedule 100 life forever start-time now
!
```

NIL LEARNING ©2014 All rights reserved.

```
ip sla 101
 icmp-echo 172.29.0.1 source-interface Serial0/0/1
 timeout 500
 frequency 3
ip sla schedule 101 life forever start-time now
!
track 100 rtr 100 reachability
 delay down 10 up 20
!
track 101 rtr 101 reachability
 delay down 10 up 20
!
interface FastEthernet0/0
 ip address 192.168.0.1 255.255.255.0
 ip nat inside
!
interface Serial0/0/0
 description *** Link to ISP 1 ***
 ip address 172.16.1.1 255.255.255.252
 ip nat outside
!
interface Serial0/0/1
 description *** Link to ISP 2 ***
 ip address 172.17.3.1 255.255.255.252
 ip nat outside
!
ip local policy route-map LocalPolicy
!
ip route 0.0.0.0 0.0.0.0 Serial0/0/0 10 track 100
ip route 0.0.0.0 0.0.0.0 Serial0/0/1 11 track 101
ip route 0.0.0.0 0.0.0.0 Serial0/0/0 250
ip route 0.0.0.0 0.0.0.0 Serial0/0/1 251
!
```

NIL LEARNING ©2014 All rights reserved.

```
!
ip nat inside source route-map ISP_A interface Serial0/0/0 overload
ip nat inside source route-map ISP B interface Serial0/0/1 overload
!
ip access-list extended PingISP_A
 permit icmp host 172.16.1.1 host 172.29.0.1
ip access-list extended PingISP_B
 permit icmp host 172.17.3.1 host 172.29.0.1
!
route-map ISP_A permit 10
 match interface Serial0/0/0
!
route-map ISP_B permit 10
 match interface Serial0/0/1
!
route-map LocalPolicy permit 10
 match ip address PingISP_A
 set interface Serial0/0/0
!
route-map LocalPolicy permit 20
 match ip address PingISP_B
 set interface Serial0/0/1
!
!
event manager applet ISP_A_down
 event track 100 state down
 action 1.0 syslog msg "ping to central site from Serial 0/0/0 failed"
event manager applet ISP_A_up
 event track 100 state up
 action 1.0 syslog msg "central site is reachable"
event manager applet ISP_B_down
 event track 101 state down
 action 1.0 syslog msg "ping to central site from Serial 0/0/1 failed"
```

NIL LEARNING ©2014 All rights reserved.

```
event manager applet ISP_B_up
 event track 101 state up
 action 1.0 syslog msg "central site is reachable"
!
end
```

Summary

With the ever faster replacement of traditional WAN networks with MPLS VPN- or Internet-based solutions, it's increasingly important to have a good design and implementation strategy for small multi-homed sites. While it's easy to implement multi-homed sites whenever you are able to run a routing protocol between the customer edge (CE) and provider edge (PE) router, as is the case with most MPLS VPN implementations, the static default routing imposed on most Internet customers by their ISPs make reliable multi-homing almost impossible in modern networks that are not able to signal loss of layer-2 connectivity reliably.

The *Reliable Static Routing Using Object Tracking* feature available in Cisco IOS release 12.4 allows you to tie static route viability to a tracked object (interface, another route ...). If you track the state of the next-hop router, it's possible to detect layer-3 failures reliably, triggering a reroute to the backup ISP. You can improve this design, track the end-to-end availability of the central site and reroute to the backup ISP whenever you cannot reach the central site through the primary ISP. Even more, you don't have to rely on ICMP echo packets; IP SLA feature of Cisco IOS can track availability of a large number of applications (for example, your company's central web server).

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 ©2014 All rights reserved.

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** certifica- tions, 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 innova- tive 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.**

©2014 All rights reserved.

More Info

Slovenia

T: +386 1 4746 500

E: sales-support@nil.com

Botswana

T: +267 318 1684

E: training@it-iq.bw

Croatia

T: +385 (0)51 583 255

E: info-nilcroatia@nil.com

Kenya

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Morocco

T: +212(0) 660 808 394

E: info-nilmorocco@nil.com

Nigeria

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Saudi Arabia

T: +966 1 465 4641

E: info.nilme@nil.com

Serbia

T: +381 11 2282 818

E: info-nilserbia@nil.co.rs

South Africa

T: +27 (0)11 575 4637

E: mea_sales@nil.com

Turkey

T: +902 123 81 8639

E: info-nilturkey@nil.com

USA

T: +1 612 886 3900

E: info-nilusa@nil.com

www.learning.nil.com

NIL LEARNING ©2014 All rights reserved.