Scaling EIGRP Networks with Stub Routers

By Ivan Pepelnjak

1. 4. 2007

Enhanced Interior Gateway Routing Protocol (EIGRP), Cisco's proprietary yet hugely successful and widely deployed routing protocol, is known to behave disappointingly in inadequately designed networks. Cisco has improved EIGRP's behavior dramatically with the introduction of *stub routers* in Cisco IOS release 12.0(7)T (integrated in IOS release 12.1, thus being available for a number of years). However, this feature has remained a well-hidden mystery, appearing in a short whitepaper, a Networkers presentation and getting five slides (almost identical to the Networkers presentation) in the release 3.0 of the Building Scalable Cisco Internetworks (BSCI) course.

In this article, we'll explore the typical problems that the EIGRP stub routers help to solve, see how the introduction of stub routers improves network stability, and implement a fully redundant remote location (stub site); yet another very common design requirement that is not documented anywhere.

Introduction to Stub Routers

Let's start with an easy example: here is a central site router linked to a large number of remote offices over slow speed or unreliable links (Figure 1).

Figure 1: Simple hub-and-spoke WAN network

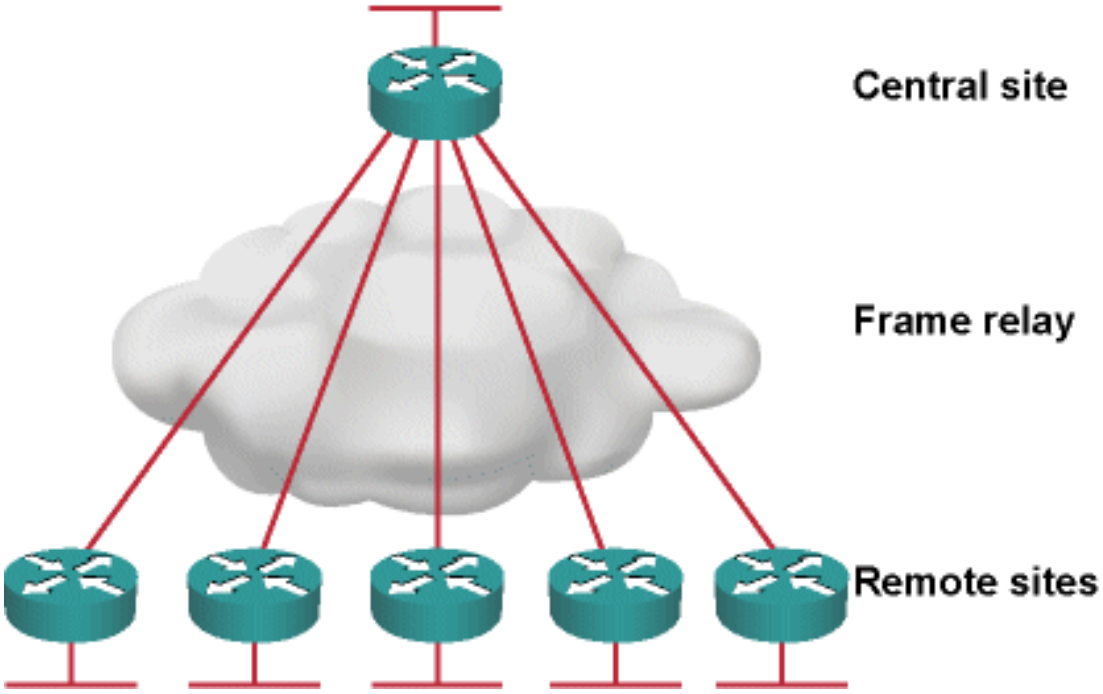

Each time a single remote office becomes unreachable, the central site router starts the Diffusing Update Algorithm (DUAL) process, querying all other remote office routers to determine whether they might have a better path to the lost destination (Figure 2).

Figure 2: DUAL query process in the hub-and-spoke network

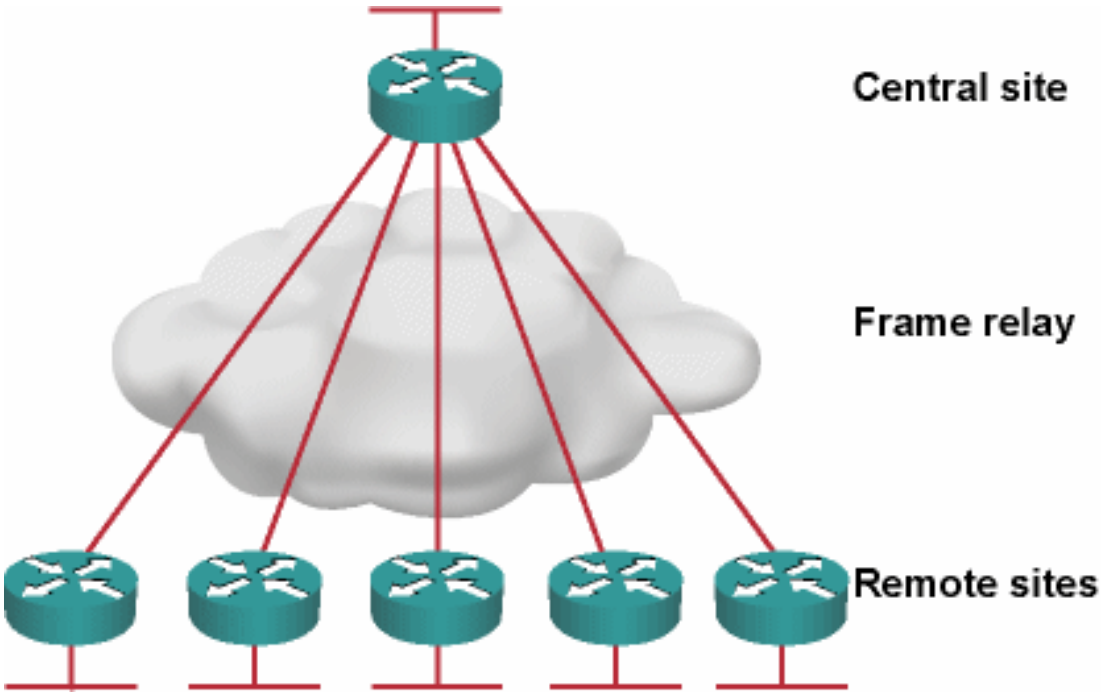

As the remote offices have no connectivity apart from their upstream link to the central router, these queries are obviously a waste of bandwidth and processing power. Even worse, in larger networks they might cause a Stuck-in-Active (SIA) event, potentially bringing down an EIGRP adjacency between core routers, thus resulting in a massive network blackout (SIA events are a major cause of network outages in poorly designed EIGRP networks).

Note: You can get an in-depth description of the DUAL process and its scalability limitations in the EIGRP Network Design Solutions book. The book is out-of-print for a few years, but you can still read it on-line with a Safari subscription.

The EIGRP stub router functionality, introduced in IOS release 12.0(7)T, gives you exactly what its name implies – the ability to declare a remote office router as a stub router (a router with no further connectivity). You configure a stub router with the **eigrp stub [connected | static | redistributed | summary | receive-only]** router configuration command. Apart from the **receive-only** option, which prevents the stub router from announcing any routes (not really useful), the other keywords define which routes inserted into the EIGRP routing process from other sources the router should announce to its neighbors.

Note: A stub router will not advertise routes received from an EIGRP neighbor to another EIGRP neighbor. Furthermore, to advertise an external route to EIGRP neighbors, the route has to be inserted in the EIGRP topology table first (with the ***redistribute*** *command) and allowed to be advertised with the* ***eigrp stub*** *command.*

In our example, the remote office routers have no external connectivity, so they just need to announce connected routes. The relevant configuration commands to transform a remote office router into an EIGRP stub router are shown in Listing 1.

Listing 1: Configuring a remote office router as an EIGRP stub router

```
router eigrp 1
 eigrp stub connected
```

The EIGRP stub routers announce their status in a new TLV (type-length-value triple) in the EIGRP hello messages. If their neighbors understand the new TLV, they stop sending queries to the stub router; instead the queries are responded to with the inaccessible message (infinite reply), while the stub routers get notified about the change with an infinite update message. This results in improved convergence time as the core routers don't have to wait for query responses from the remote offices.

Note: The stub routing feature by itself does not prevent core routes from being advertised to the stub peer. Use summarization or filter lists combined with default routing to avoid unnecessary updates.

Routers running an IOS release older than 12.0(7)T simply ignore the new TLV. They would thus still query the stub routers, but the stub routers would reply immediately without propagating the query (still resulting in marginally improved performance in meshed networks).

Note: As the stub status is carried in EIGRP hello messages, any change to the stub status causes EIGRP adjacency teardown and reestablishment.

You can check whether an EIGRP neighbor is a stub router with the **show ip eigrp neighbor detail** command. The printout generated on our core router is displayed in Listing 2.

Listing 2: EIGRP neighbors displayed on the core router

```
a1#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 1
H   Address      Interface    Hold Uptime   SRTT   RTO  Q  Seq
                              (sec)         (ms)       Cnt Num
0   172.16.1.33  Se0/0/0.401    11 00:01:03 1031  5000  0  9
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 1
   Stub Peer Advertising ( CONNECTED ) Routes
   Suppressing queries
1   172.16.1.2   Se0/0/0.100    10 00:03:32  753  4518  0  4
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 2
   Stub Peer Advertising ( CONNECTED ) Routes
   Suppressing queries
```

Dual-homed Remote Sites

In most mission-critical networks, the remote sites have at least two upstream links to two core routers (Figure 3). In these designs, the amount of DUAL traffic generated by an outage is significantly higher than in the previous scenario, as both central routers start querying the remote offices.

Figure 3: Dual-homed hub-and-spoke WAN network

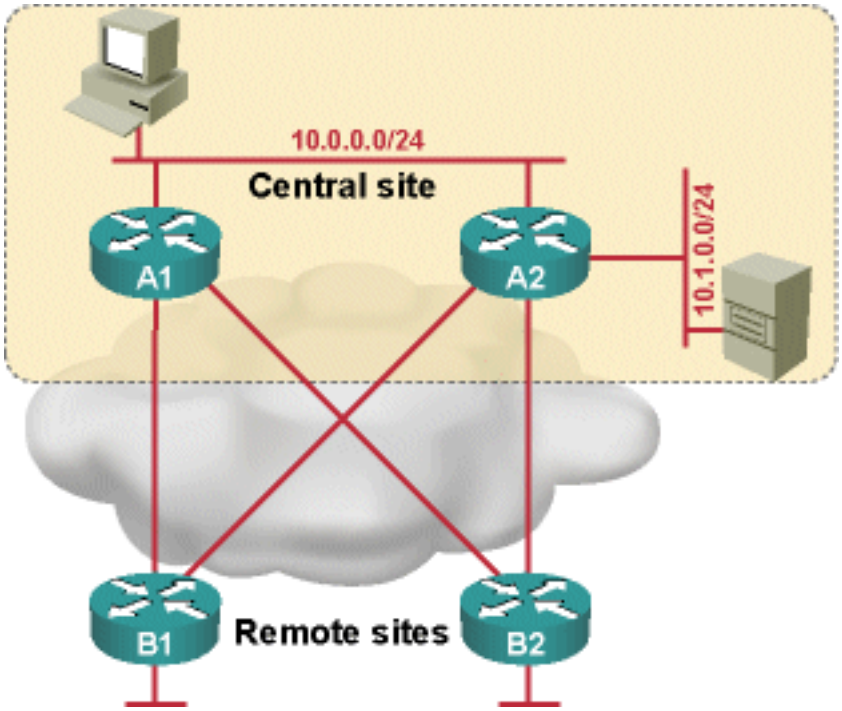

Furthermore, if the core routers lose their direct connectivity (for example, the LAN interface of one of them fails), the transit traffic might start flowing across remote sites, congesting their slow-speed links (Figure 4).

Figure 4: Dual-homed hub-and-spoke WAN network

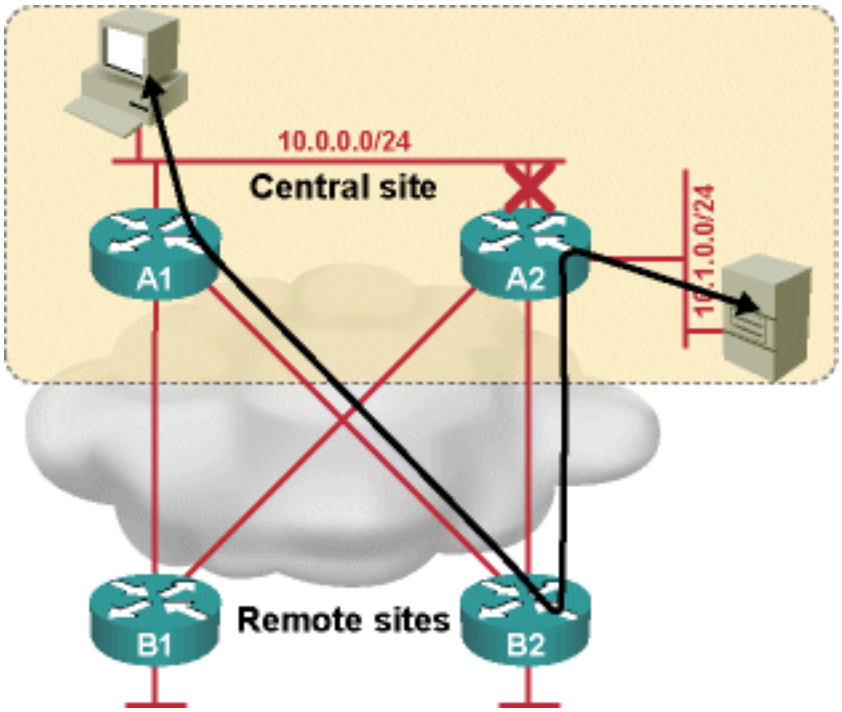

Both undesired effects are solved very effectively by the EIGRP stub routers. The stub routers are not queried in the DUAL process, thus the WAN link congestion is avoided. Furthermore, as these routers never propagate EIGRP routes to other EIGRP neighbors, they cannot become transit routers.

For example, the Listing 3 displays the EIGRP topology table on router A1 when the connectivity between A1 and A2 was lost and the remote office routers were not configured as stub routers. The routes beyond A2 are reachable through all remote office routers (B1 and B2).

Listing 3: Core routes reachable through remote office routers

```
a1#show ip eigrp topology 10.1.0.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 10.1.0.0 255.255.255.0
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 2684416
  Routing Descriptor Blocks:
  172.16.1.33 (Serial0/0/0.401), from 172.16.1.33, Send flag is 0x0
      Composite metric is (2684416/2172416), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 40100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
  172.16.1.2 (Serial0/0/0.100), from 172.16.1.2, Send flag is 0x0
      Composite metric is (2684416/2172416), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 40100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
```

When the remote office routers are configured as stub routers, A1 loses its connectivity to 10.1.0.0/24 when its connection to A2 is lost (Listing 4).

Listing 4: Core route no longer available via remote offices

```
%DUAL-5-NBRCHANGE: IP-
EIGRP(0) 1: Neighbor 10.0.0.6 (FastEthernet0/0) is down: holding time expire
d

a1#show ip eigrp topology 10.1.0.0 255.255.255.0

% IP-EIGRP (AS 1): Route not in topology table
```

Resilient Remote Sites

Network designers that aim to increase overall network reliability, usually deploy two routers at the remote offices, each one of them being connected to one of the upstream links (Figure 5). To reduce the amount of routing information exchange, heavy route summarization is usually deployed on the links between the core routers and remote offices. In our example, the core routers announce only a summary default route toward the remote offices (see Listing 5 and Listing 6 for the typical core router and remote office router configuration).

Figure 5: Dual-homed hub-and-spoke WAN network

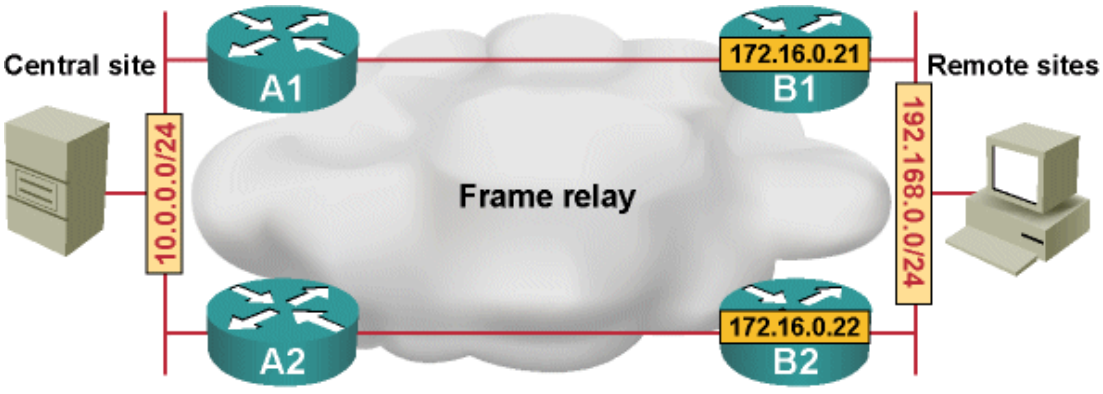

Listing 5: IP routing and EIGRP configuration on A1

```
hostname a1

!

interface Loopback0

 ip address 172.16.0.11 255.255.255.255

!
```

```
interface FastEthernet0/0
 ip address 10.0.0.5 255.255.255.0
!
interface Serial0/0/0.100 point-to-point
 ip address 172.16.1.1 255.255.255.252
 ip summary-address eigrp 1 0.0.0.0 0.0.0.0 5
 frame-relay interface-dlci 100
!
router eigrp 1
 network 0.0.0.0
 no auto-summary
```

Listing 6: IP routing and EIGRP configuration on B1

```
hostname b1
!
interface Loopback0
 ip address 172.16.0.21 255.255.255.255
!
interface FastEthernet0/0
 ip address 192.168.0.5 255.255.255.0
 standby 1 ip 192.168.0.1
 standby 1 preempt
!
interface Serial0/0/0.100 point-to-point
 ip address 172.16.1.2 255.255.255.252
 frame-relay interface-dlci 100
!
router eigrp 1
 network 0.0.0.0
 no auto-summary
 eigrp stub connected
```

The introduction of the EIGRP stub functionality in this design poses an interesting challenge: while the DUAL traffic is reduced (as expected), the routing might stop working when one of the upstream connections fail. For

example, consider the scenario when the link between A2 and B2 fails. B2 will still advertise its loopback interface to B1, but the host route will not be propagated to the core router (A1), making B2's loopback inaccessible. If you use loopback interfaces for network management (and you should do that in any well-designed network), you'd lose access to B2 from the network management station.

Likewise, the default route advertised to B1 from A1 is not propagated to B2. If the end station on remote office LAN sends its packets to B2, they would be dropped even though there is still a working link between the remote office and the central site (see Figure 6). The routing table on B2 after the WAN link failure is also included in Listing 7. As you can see, the only routes advertised from B1 to B2 are its directly connected interfaces.

Figure 6: Remote site routing stops working after a WAN link failure

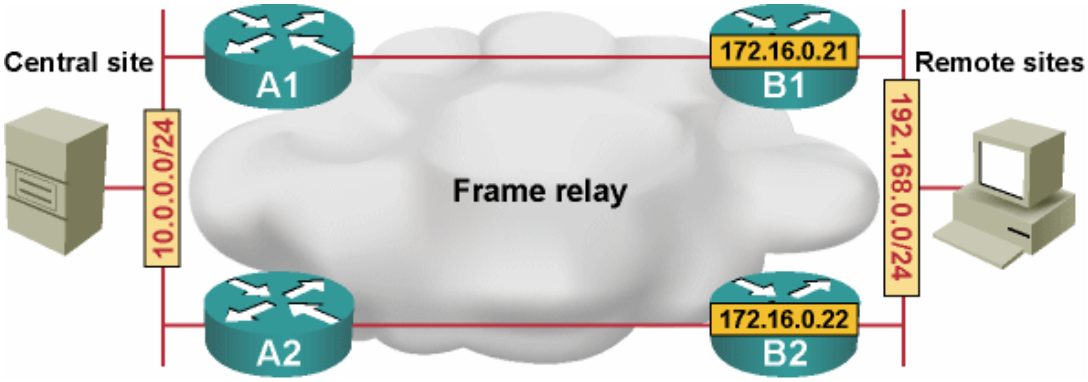

Listing 7: Routes advertised from B1 to B2

```
b2#show ip route eigrp
     172.16.0.0 255.255.0.0 is variably subnetted, 3 subnets, 2 masks
D       172.16.0.21 255.255.255.255
           [90/156160] via 192.168.0.5, 00:02:54, FastEthernet0/0
D       172.16.1.0 255.255.255.252
           [90/2172416] via 192.168.0.5, 00:02:54, FastEthernet0/0
```

You could solve this routing problem by introducing another routing protocol (for example, RIPv2) at the remote site and performing two-way redistribution between EIGRP and RIPv2, resulting in a pretty complex design. Fortunately, Cisco provided just the tool we need in IOS release

12.3(11)T – the **eigrp stub leak-map** option. With this option, you can specify a **route-map** that selects the routes a stub router propagates to its neighbors (effectively turning it into a not-so-stubby-router).

*Note: The **leak-map** option configured on the stub router does not change the DUAL behavior on its neighbors (core routers). You must thus ensure through proper network design that the exclusion of the stub router from the DUAL querying process does not lead to undesired side effects.*

To configure the route leaking on the stub router, you have to configure a **route-map** that would usually refer to an **access-list** which would in turn select the routes you want to leak (but you could, for example, also decide to leak only external EIGRP routes with the proper **match** keywords in the **route-map**). For example, in the final configuration of the B1 router (included in Listing 8), the *EigrpLeakedRoutes* access list permits the default route (this one has to be leaked from A1 to B2) and the loopback interfaces of the remote office routers (these addresses have to be leaked to A1). The final EIGRP routing table on B2 is shown in Listing 9; note that it also contains the default route advertised from B1 due to the **leak-map**.

Listing 8: Leak-map configuration on B1

```
router eigrp 1
 network 0.0.0.0
 no auto-summary
 eigrp stub connected leak-map EigrpLeakedRoutes
!
ip access-list standard EigrpLeakedRoutes
 permit 0.0.0.0
 permit 172.16.0.16 0.0.0.15
!
route-map EigrpLeakedRoutes permit 10
 match ip address EigrpLeakedRoutes
```

Listing 9: Correct routing table on B2

```
b2#show ip route eigrp
     172.16.0.0 255.255.0.0 is variably subnetted, 3 subnets, 2 masks
D       172.16.0.21 255.255.255.255
```

```
            [90/156160] via 192.168.0.5, 00:00:28, FastEthernet0/0
D       172.16.1.0 255.255.255.252
            [90/2172416] via 192.168.0.5, 00:00:28, FastEthernet0/0
D*   0.0.0.0 0.0.0.0 [90/2174976] via 192.168.0.5, 00:00:27, FastEthernet0/0
```

Tight Control over Route Leaking

The introduction of EIGRP route leaking on the stub routers has to be based on a thoroughly checked network design. For example, if the route leaking is configured on B1 and B2 with the configuration from Listing 8 and the link bandwidth between A2 and B2 is significantly lower than the one between A1 and B1 (Figure 7), B2 will leak the default route back to A2, which will store it in its EIGRP topology database (Listing 10). Due to limitations on querying the stub routers (DUAL queries are never sent to these routers), the leaked routes could persist in the EIGRP topology tables even after the original IP prefix on which they were based is no longer reachable, resulting in temporary routing loops.

Note: EIGRP routers eventually arrive to the correct final routing topology even when stub routers nondiscriminatory leak core routes. However, the convergence time is increased as is the chance of introducing temporary loops (which never occur when using the regular DUAL algorithm).

Figure 7: Remote site with a primary and a backup uplink

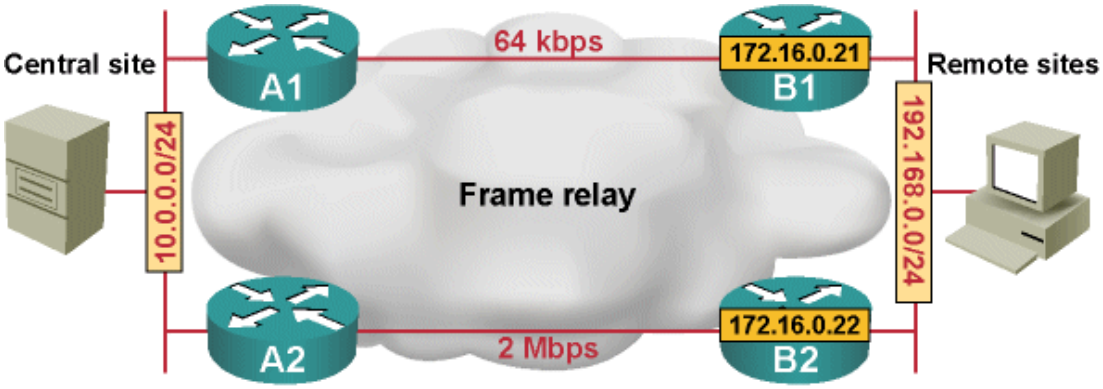

Listing 10: Default route is leaked back to the core router

```
a2#show ip eigrp topology 0.0.0.0
IP-EIGRP (AS 1): Topology entry for 0.0.0.0 0.0.0.0
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 28160
```

```
Routing Descriptor Blocks:
0.0.0.0 (Null0), from 0.0.0.0, Send flag is 0x0
    Composite metric is (28160/0), Route is Internal
    Vector metric:
      Minimum bandwidth is 100000 Kbit
      Total delay is 100 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 0
    Exterior flag is set
172.16.1.6 (Serial0/0/0.101), from 172.16.1.6, Send flag is 0x0
    Composite metric is (41029120/2174976), Route is Internal
    Vector metric:
      Minimum bandwidth is 64 Kbit
      Total delay is 40200 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 3
    Exterior flag is set
```

To alleviate the possibility of a routing loop in EIGRP networks with stub routers leaking routes to their neighbors, the route leaking has to be tightly controlled – you have to match a combination of IP prefixes and the outgoing interface in the **route-map** attached to the EIGRP stub router with the **leak-map** keyword. In our sample network, the route map would thus perform the following route selection:

- Default route (0.0.0.0/0) would only be leaked to the LAN interface;
- Stub router loopback addresses would only be leaked to the WAN interface (uplink to the central site).

This setup prevents the default route from being leaked back to the core routers. It also prevents the loopback addresses from the remote site from

being readvertised back into the same remote site. The relevant parts of the router configuration are included in Listing 11.

Listing 11: Tightly controlled EIGRP route leaking on B1 and B2

```
router eigrp 1
 network 0.0.0.0
 no auto-summary
 eigrp stub connected leak-map TightEigrpLeak
!
ip access-list standard EigrpDefault
 permit 0.0.0.0
ip access-list standard EigrpLoopback
 permit 172.16.0.16 0.0.0.15
!
route-map TightEigrpLeak permit 100
 match ip address EigrpDefault
 match interface FastEthernet0/0
!
route-map TightEigrpLeak permit 200
 match ip address EigrpLoopback
 match interface Serial0/0/0.100
```

As expected, with the corrected route, leaking the default route is no longer advertised from B2 to A2, while the loopback address of B1 is advertised to A2 from A1 (as it receives the IP prefix from B1) and B2 (which leaks the IP prefix received from B1). The relevant printouts are shown in Listing 12.

Listing 12: Corrected EIGRP topology on A2

```
a2#show ip eigrp topology 0.0.0.0
IP-EIGRP (AS 1): Topology entry for 0.0.0.0 0.0.0.0
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 28160
  Routing Descriptor Blocks:
  0.0.0.0 (Null0), from 0.0.0.0, Send flag is 0x0
      Composite metric is (28160/0), Route is Internal
      Vector metric:
```

```
        Minimum bandwidth is 100000 Kbit
        Total delay is 100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 0
      Exterior flag is set
a2#show ip eigrp topology 172.16.0.21 255.255.255.255
IP-EIGRP (AS 1): Topology entry for 172.16.0.21 255.255.255.255
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2300416
  Routing Descriptor Blocks:
  10.0.0.5 (FastEthernet0/0), from 10.0.0.5, Send flag is 0x0
      Composite metric is (2300416/2297856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 25100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
  172.16.1.6 (Serial0/0/0.101), from 172.16.1.6, Send flag is 0x0
      Composite metric is (40642560/156160), Route is Internal
      Vector metric:
        Minimum bandwidth is 64 Kbit
        Total delay is 25100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
```

Summary

In this article, you've seen how you can use the EIGRP stub router functionality (introduced in IOS release 12.0T and 12.1) to decrease bandwidth utilization on slow links and dramatically improve EIGRP convergence times and its robustness. This functionality is easiest to deploy in stub sites (remote offices) with a single router.

As soon as you're deploying more than one router on the stub site, each one of them could potentially become a transit router. Cisco has introduced the **leak-map** enhancement to the **eigrp stub** configuration command in IOS release 12.3(11)T that allows you to specify which EIGRP routes a stub router should leak to its neighbor (or, in network design terms, for which IP prefixes the stub router should act as a transit router). Proper deployment of the **leak-map** feature guarantees correct routing even in various failure scenarios, while retaining all the benefits introduced with the EIGRP stub router concept.

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** 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.**

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