Configure the BGP Local-AS Feature

Contents

Introduction

This document describes the Border Gateway Protocol (BGP) local-AS feature, which was initially available in Cisco IOS®Software Release 12.0(5)S.

Prerequisites

Requirements

This document recommends knowledge of the BGP routing protocol and its operations. For more information, refer to the Examine Border Gateway Protocol Case Studies.

Components Used

The information in this document was created with this software version, but is not restricted to specific software and hardware versions:

- Cisco IOS Software Release 12.2(28)

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Background Information

The local-AS feature allows a router to appear to be a member of a second autonomous system (AS), in addition to its original AS. This feature can only be used for true eBGP peers. You cannot use this feature for two peers that are members of different confederation sub-ASs.

The local-AS feature is a useful tool when one ISP acquires another, as it allows routers in the acquired ISP to seamlessly join the acquiring ISP's AS. This is particularly beneficial when the acquired ISP's customers prefer to maintain their existing peering arrangements and configurations. With the local-AS feature, routers in the acquired ISP can appear as members of the acquiring ISP's AS, while still retaining their original AS number in the eyes of their customers. This ensures a smooth transition and minimizes disruptions for the acquired ISP's customers.

For example, refer to the network diagrams, where in Figure 1, ISP-A has not yet purchased ISP-B and inFigure 2, ISP-A has purchased ISP-B, and ISP-B uses the local-AS feature. If ISP-A purchases ISP-B, but ISP-B customers do not want to modify any configuration, the local-AS feature allows routers in ISP-B to become members of ISP-A (100) AS. At the same time, these routers appear to their customers to retain their ISP-B (200) AS number.

In Figure 2, ISP-B now belongs to AS 100, and ISP-C (customer) to AS 300. When peering with ISP-C, ISP-B uses AS 200 as its AS number with the use of the **neighbor <ISP-C> local-as 200** command. In the BGP updates sent from ISP-B to ISP-C, the AS_SEQUENCE in the AS_PATH attribute contains "200 100". The "200" is prepended by ISP-B due to the **local-as 200** command configured for ISP-C.

Normally a combined ISP-A/B renumbers the routers in ISP-B to be part of AS 100. What if ISP-C is unable to change its eBGP configurations with ISP-B? Prior to the local-AS feature, the combined ISP-A/B has to maintain two AS numbers. The **local-as** command allows ISP-A/B to physically be one AS while it appears to be two ASs to ISP-C.

Command Syntax

This list shows the syntax of the commands that the configurations in this document use:

- neighbor <neighbor IP address> local-as <local-AS-number>
- neighbor peer-group local-as <local-AS-number>

Local-AS cannot be customized for individual peers in a peer group.

Local-AS cannot have the local BGP protocol AS number or the AS number of the remote peer.

The **local-as** command is valid only if the peer is a true eBGP peer. It does not work for two peers in different sub-ASs in a confederation.

Configure

This section presents you with the information to configure the features this document describes.

Network Diagrams

This document uses these network setups.

Figure 1

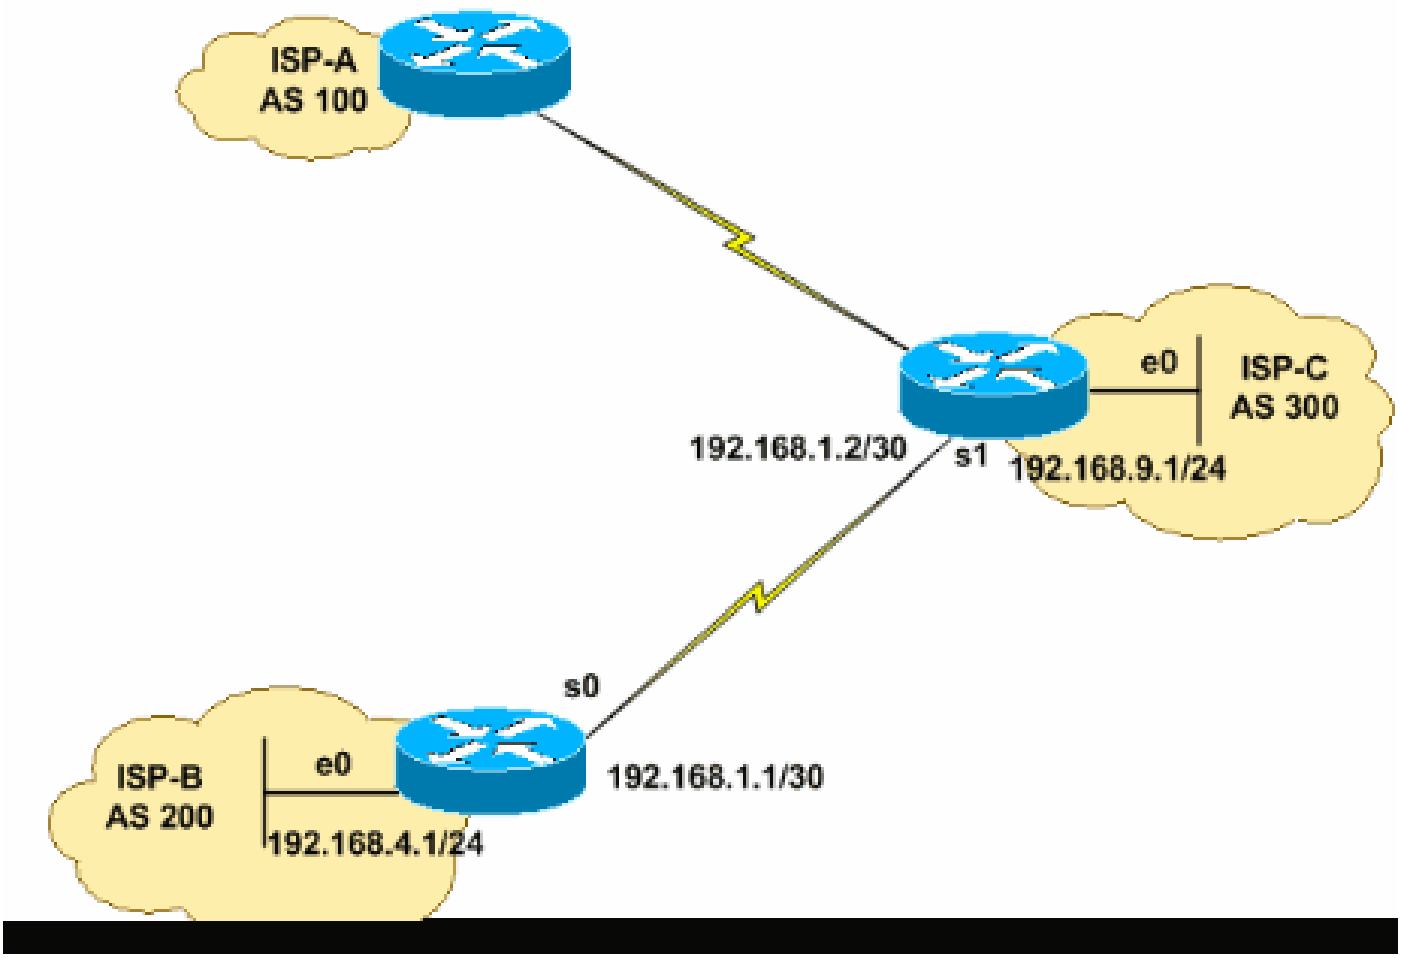
ISP-A
AS 100
e0
ISP-C
AS 300
192.168.1.2/30
s1
192.168.9.1/24
s0
ISP-B
AS 200
e0
192.168.1.1/30
192.168.4.1/24

Figure 2

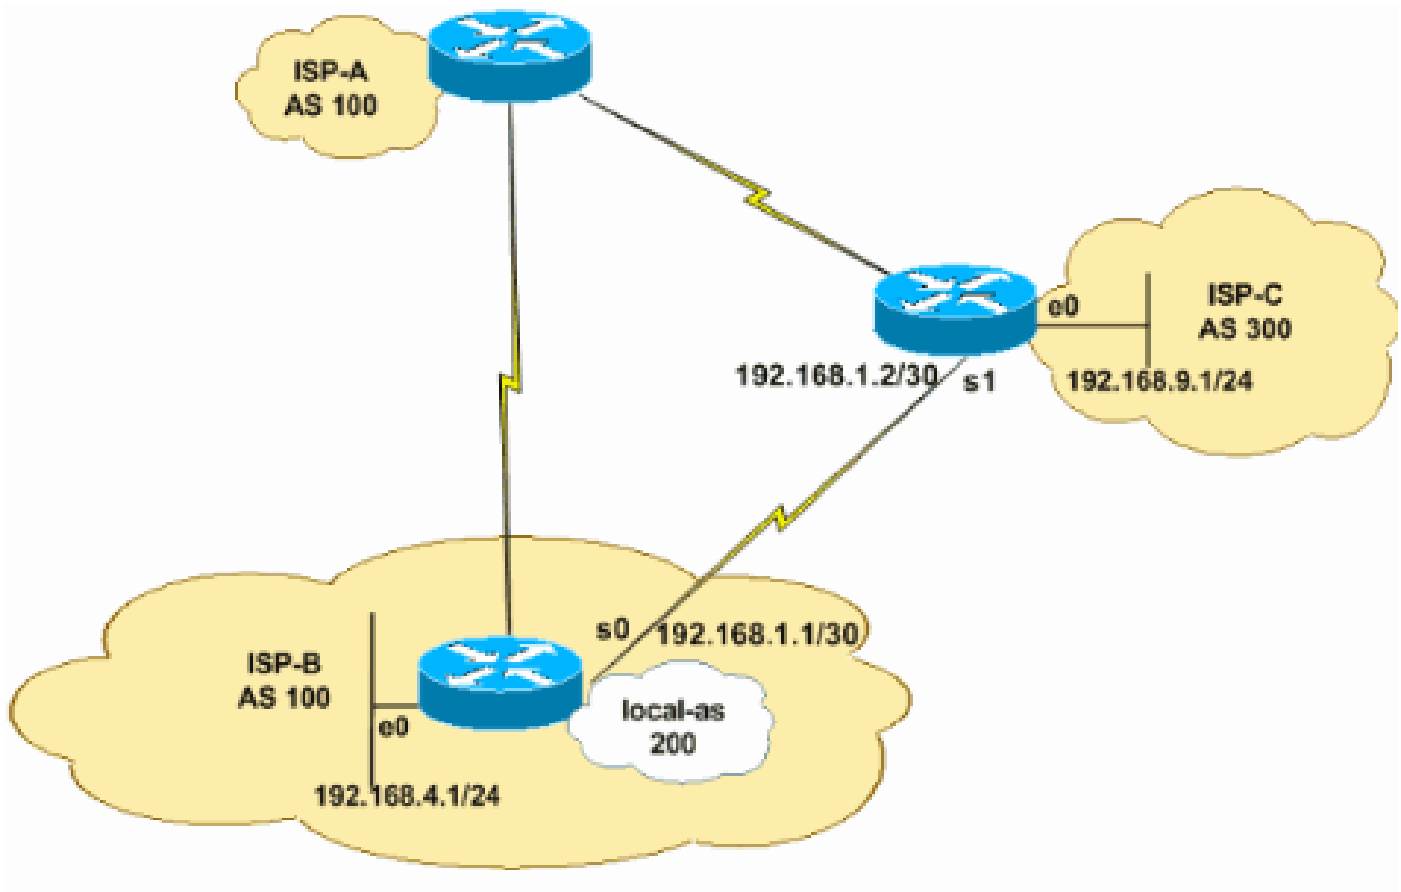
ISP-A
AS 100
e0
ISP-C
AS 300
192.168.1.2/30
s1
192.168.9.1/24
s0
192.168.1.1/30
ISP-B
AS 100
e0
local-as
200
192.168.4.1/24

Configurations

This document uses these configurations:

- ISP-B (AS 100, local-as 200)
- ISP-C (AS 300)

ISP-B (AS 100, local-as 200)

```
hostname ISP-B
!
interface serial 0
ip address 192.168.1.1 255.255.255.252
!
interface ethernet 0
ip address 192.168.4.1 255.255.255.0
!
router bgp 100

!--- Note the AS number 100. This is the AS number of ISP-A, which is now
!--- used by all routers in ISP-B after its acquisition by ISP-A.

 neighbor 192.168.1.2 remote-as 300

!--- Defines the e-BGP connection to ISP-C.

 neighbor 192.168.1.2 local-as 200

!--- This command makes the remote router in ISP-C to see this
!--- router as belonging to AS 200 instead of AS 100.
!--- This also make this router to prepend AS 200 in
!--- all updates to ISP-C.

network 192.168.4.0
!
!
```

ISP-C (AS 300)

```
hostname ISP-C
!
interface serial 1
ip address 192.168.1.2 255.255.255.252
!
interface ethernet 0
ip address 192.168.9.1 255.255.255.0
!
router bgp 300
 neighbor 192.168.1.1 remote-as 200

!--- Defines the e-BGP connection to ISP-B.
!--- Note AS is 200 and not AS 100.

 network 192.168.9.0
!
!
```

Verify

This section provides information you can use to confirm your configuration properly works.

View the BGP routing table to see how the **local-as** command changed the AS_PATH. What you observe is that ISP-B prepends AS 200 to updates that are sent to and received from ISP-C. In addition, note that ISP-B is in AS number 100.

```
<#root>

ISP-B#

show ip bgp summary

  BGP router identifier 192.168.4.1,

local AS number 100


  BGP table version is 3, main routing table version 3
  2 network entries and 2 paths using 266 bytes of memory
  2 BGP path attribute entries using 104 bytes of memory
  1 BGP AS-PATH entries using 24 bytes of memory
  0 BGP route-map cache entries using 0 bytes of memory
  0 BGP filter-list cache entries using 0 bytes of memory
  BGP activity 2/6 prefixes, 2/0 paths, scan interval 15  secs

  Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
  192.168.1.2     4   300      29      29        3    0    0 00:25:19       1
```

In this output, note that ISP-C sees ISP-B as part of AS 200.

```
<#root>

ISP-C#

show ip bgp summary


  BGP table version is 3, main routing table version 3
  2 network entries (2/6 paths) using 480 bytes of memory
  2 BGP path attribute entries using 192 bytes of memory
  0 BGP route-map cache entries using 0 bytes of memory
  0 BGP filter-list cache entries using 0 bytes of memory

  Neighbor        V    AS MsgRcvd MsgSent  TblVer InQ OutQ   Up/Down  State/PfxRcd
  192.168.1.1     4

200

       34      34       3   0    0  00:30:19   1
```

Notice in this output that ISP-B prepends "200" to the routes learned from ISP-C.

```
<#root>

ISP-B#

show ip bgp


  BGP table version is 3, local router ID is 192.168.4.1
  Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
  Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop            Metric LocPrf  Weight   Path
  *> 192.168.4.0      0.0.0.0               0          32768   i
  *> 192.168.9.0      192.168.1.2           0              0

200

 300 i
```

Notice that ISP-C sees routes from ISP-B with an AS_PATH of "200 100".

```
<#root>

ISP-C#

show ip bgp


  BGP table version is 3, local router ID is 192.168.1.2
  Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
  Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop          Metric LocPrf  Weight Path
  *> 192.168.4.0      192.168.1.1          0              0

200 100

 i
  *> 192.168.9.0      0.0.0.0              0          32768 i
```

These commands show the configured **local-as** values in their output:

- show ip bgp neighbor <neighbor IP address>
- show ip bgp peer-group peer <group name>

```
<#root>

ISP-B#

show ip bgp neighbors 192.168.1.2

BGP neighbor is 192.168.1.2,  remote AS 300,

local AS 200
```

```
, external link
  BGP version 4, remote router ID 192.168.9.1
  BGP state = Established, up for 00:22:42
  Last read 00:00:42, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                2          1
    Keepalives:            25         25
    Route Refresh:          0          1
    Total:                 28         28
  Default minimum time between advertisement runs is 30 seconds
```

Troubleshoot

The **debug ip bgp updates** command displays the received prefixes with its attributes from the neighbor. This output shows that the prefix 192.168.4.0/24 is received with AS PATH 200, 100.

<#root>

```
ISP-C#
*May 10 12:45:14.947: BGP(0): 192.168.1.1 computing updates, afi 0, neighbor version 0, table vers
*May 10 12:45:14.947: BGP(0): 192.168.1.1 send UPDATE (format) 192.168.9.0/24, next 192.168.1.2, m
*May 10 12:45:14.947: BGP(0): 192.168.1.1 1 updates enqueued (average=52, maximum=52)
*May 10 12:45:14.947: BGP(0): 192.168.1.1 update run completed, afi 0, ran for 0ms, neighbor versi
*May 10 12:45:14.947: BGP: 192.168.1.1 initial update completed
*May 10 12:45:15.259: BGP(0): 192.168.1.1 rcvd UPDATE w/ attr: nexthop 192.168.1.1, origin i, metr

 200 100

*May 10 12:45:15.259: BGP(0): 192.168.1.1 rcvd

192.168.4.0/24

*May 10 12:45:15.279: BGP(0): Revise route installing 192.168.4.0/24 -> 192.168.1.1 to main IP tab
ISP-C#
```

Related Information

- **Examine Border Gateway Protocol Frequently Asked Questions**
- **BGP Technical Support**
- **Cisco Technical Support & Downloads**