

How to Use HSRP to Provide Redundancy in a Multihomed BGP Network

Document ID: 13768

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

Configure

- Network Diagram
- Configurations

Verify

- Packets Going from the Local Network Toward the Destination
- Packets Coming from the Destination Toward the Local Network

Troubleshoot

Related Information

Introduction

This document describes how to provide redundancy in a multihomed Border Gateway Protocol (BGP) network where you have connections to two separate Internet service providers (ISPs). In the event of a failure of connectivity toward one ISP, the traffic is rerouted dynamically through the other ISP with the BGP `set as-path {tag | prepend as-path-string}` command and Hot Standby Router Protocol (HSRP).

Prerequisites

Requirements

Readers of this document should have knowledge of these topics:

- Cisco HSRP [↗](#)
- Configuring HSRP
- BGP Best Path Selection Algorithm
- Configuring BGP

Components Used

This document is not restricted to specific software and hardware versions.

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, make sure 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 objective of the configuration in this document is to achieve this network policy:

- All outbound traffic sourced from hosts on network 192.168.21.0/24 and destined to the Internet must be routed through R1 to ISP-A. However, if that link fails or R1 fails, all outbound traffic must be rerouted through R2 to ISP-B (and then to the Internet) without manual intervention.
- All inbound traffic destined to an autonomous system, AS 100, from the Internet must be routed by way of R1. In the event that the link from ISP-A to R1 fails, the inbound traffic must automatically be rerouted through ISP-B to R2.

These requirements can be met with two technologies: BGP and HSRP.

The first objective of a fully redundant outbound path can be realized with HSRP. Typically PCs do not have the capability to collect and exchange routing information. The IP address of the default gateway is statically configured on a PC and if the gateway router goes down, the PC loses connectivity to any device beyond its local network segment. This is the case even if an alternate gateway exists. HSRP has been designed to meet these requirements. Refer to HSRP Features and Functionality for more information.

The second objective can be achieved with the BGP **set as-path prepend** command, which allows BGP to propagate a longer AS path (by way of prepending its own AS number more than once) through the R2 to ISP-B link for prefix 192.168.21.0/24. Thus, all traffic destined for 192.168.21.0/24 that comes from the outside of AS 100 takes the shorter AS path through the ISP-A to R1 link. If the primary path (ISP-A to R1) fails, all traffic takes the longer AS path (ISP-B to R2) in order to reach network 192.168.21.0/24. In order to learn more about the BGP **set as-path prepend** command, refer to the AS_PATH Attribute diagram in the BGP Case Studies document.

Configure

In this section, you are presented with the information to configure the features described in this document.


Note: Use the Command Lookup Tool (registered customers only) to find more information on the commands used in this document.

Network Diagram

This document uses the network setup shown here:

In this diagram, Router 1 (R1) and Router 2 (R2) are in AS 100, which has external BGP (eBGP) peering with ISP-A (AS 300) and ISP-B (AS 400) respectively. Router 6 (R6) is a part of AS 600, which has eBGP peering with ISP-A and ISP-B. R1, R2 has iBGP peering, which is necessary to ensure optimal routing. For example, when you try to reach AS 400 internal routes, R1 does not employ the longer path over AS 300. R1 forwards the traffic to R2 instead.

R1 and R2 are also configured for HSRP over a common Ethernet segment. Hosts on the same Ethernet segment have a default route that points towards the HSRP standby IP address 192.168.21.10.



Configurations

R1

Current configuration

```
hostname R1
!
interface serial 0
ip address 192.168.31.1 255.255.255.0
!
interface Ethernet1
ip address 192.168.21.1 255.255.255.0
standby 1 priority 105
standby 1 preempt delay minimum 60
standby 1 ip 192.168.21.10
standby 1 track Serial0

!--- The
standby track serial command tracks the state of
!--- the Serial0 interface and brings down the
!--- priority of standby group 1, if the interface goes down.
!--- The
standby preempt delay minimum 60 command makes sure that
!--- R1 preempts and takes over as active router again. This command also ensures that
!--- the router waits 60 seconds before doing so in order to give BGP time enough
!--- to converge and populate the routing table. This avoids
!--- traffic being sent to R1 before it is ready to forward it.

!
!
router bgp 100
no synchronization
network 192.168.21.0
neighbor 192.168.21.2 remote-as 100
neighbor 192.168.21.2 next-hop-self
neighbor 192.168.31.3 remote-as 300
no auto-summary
!
```

R2

Current configuration:

```
hostname R2
!
interface serial 0
ip address 192.168.42.2 255.255.255.0
```

```

!
interface Ethernet1
 ip address 192.168.21.2 255.255.255.0
 standby 1 priority 100
 standby 1 preempt
 standby 1 ip 192.168.21.10
!
!
router bgp 100
 no synchronization
 network 192.168.21.0
 neighbor 192.168.21.1 remote-as 100
 neighbor 192.168.21.1 next-hop-self
 neighbor 192.168.42.4 remote-as 400
 neighbor 192.168.42.4 route-map foo out

!--- It appends AS 100 to the BGP updates sent to AS 400
!--- in order to make it a backup for the ISP-A to R1 path.

 no auto-summary
!
access-list 1 permit 192.168.21.0
route-map foo permit 10
 match ip address 1
 set as-path prepend 100
end

```

Verify

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

Certain **show** commands are supported by the Output Interpreter Tool (registered customers only) , which allows you to view an analysis of **show** command output.

When you configure redundancy in any network, you must consider two things:

- The creation of a redundant path for packets going from a local network to a destination network.
- The creation of a redundant path for packets coming back from a destination to a local network.

Packets Going from the Local Network Toward the Destination

In this example, the local network is 192.168.21.0/24. Router R1 and R2 are running HSRP on the Ethernet segment connected to interface Ethernet1. R1 is configured as the HSRP active router with a standby priority of 105, and R2 is configured with a standby priority of 100. The **standby 1 track Serial0** (s0) command on R1 allows the HSRP process to monitor that interface. If the interface state goes down, the HSRP priority is reduced. When interface s0's line protocol goes down, the HSRP priority is reduced to 95 (the default value by which priority is reduced is 10). This makes the other HSRP router, R2, have a higher priority (a priority of 100). R2 becomes the HSRP active router and attracts traffic destined to the active HSRP address 192.169.21.10.

Issue the **show standby** command in order to see the active HSRP router when the interface s0 on R1 is up:

```

R1#show standby
 Ethernet1 - Group 1
   Local state is Active, priority 105, may preempt
   Hellotime 3 sec, holdtime 10 sec
   Next hello sent in 0.338
   Virtual IP address is 192.168.21.10 configured

```

```

Active router is local
Standby router is 192.168.21.2 expires in 8.280
Virtual mac address is 0000.0c07.ac01
13 state changes, last state change 00:46:10
IP redundancy name is "hsrp-Et0-1"(default)
Priority tracking 1 interface, 1 up:
Interface                Decrement    State
Serial0                   10           Up

```

R2#show standby

```

Ethernet1 - Group 1
  State is Standby
  56 state changes, last state change 00:05:13
  Virtual IP address is 192.168.21.10
  Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (default)
  Hello time 3 sec, hold time 10 sec
  Next hello sent in 1.964 secs
  Preemption enabled
  Active router is 192.168.21.1, priority 105 (expires in 9.148 sec)
  Standby router is local
  Priority 100 (default 100)
  IP redundancy name is "hsrp-Et0-1" (default)

```

R1#show standby ethernet 1 brief

```

          P indicates configured to preempt.
          |
Interface  Grp Prio P State    Active addr    Standby addr    Group addr
Et1        1  105 P Active   local          192.168.21.2    192.168.21.10
R1#

```

R2#show standby ethernet 1 brief

```

          P indicates configured to preempt.
          |
Interface  Grp Prio P State    Active          Standby          Virtual IP
Et1        1  100 P Standby  192.168.21.1   local            192.168.21.10
R2#

```

The **show standby** command shows R1 as the active HSRP router because of the higher priority of 105. Since R1 is the active router, R1 owns the standby IP address 192.168.21.10. All IP traffic from the host configured with the default gateway to 192.168.21.10 routes through R1.

If you bring the s0 interface on Router R1 down, the HSRP active router changes since HSRP on R1 is configured with the **standby track serial 0** command. When the Serial 0 interface protocol goes down, HSRP reduces the priority of R1 by 10 (default) to 95. R1 changes its state to "Standby". R2 takes over as the active router, and thus owns the standby IP address 192.168.21.10. Accordingly, all traffic destined from hosts in the 192.168.21.0/24 segment routes traffic through R2. The **debug** and **show** command output confirms the same.

```

R1(config)# interface s0
R1(config-if)# shut
%STANDBY-6-STATECHANGE: Standby: 1: Ethernet1 state Active      -> Speak
%LINK-5-CHANGED: Interface Serial0, changed state to administratively down
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to down
%STANDBY-6-STATECHANGE: Standby: 1: Ethernet1 state Speak     -> Standby
%LINK-3-UPDOWN: Interface Serial0, changed state to down:
%STANDBY-6-STATECHANGE: Standby: 1: Ethernet1 state Active      -> Speak
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to down
%STANDBY-6-STATECHANGE: Standby: 1: Ethernet1 state Speak     -> Standby

```

Notice that R1 becomes a standby router.

If R2 goes into the active state, then you see output similar to this:

```
R2#
%STANDBY-6-STATECHANGE: Standby: 1: Ethernet1 state Standby -> Active
```

If you execute the **show standby** command on R1 and R2, observe the standby priorities after the interface s0 goes down on R1:

```
R1#show standby
Ethernet1 - Group 1
  Local state is Standby, priority 95 (configd 105), may preempt
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 0.808
  Virtual IP address is 192.168.21.10 configured
  Active router is 192.168.21.2, priority 100 expires in 9.008
  Standby router is local
  15 state changes, last state change 00:00:40
  IP redundancy name is "hsrp-Et0-1" (default)
  Priority tracking 1 interface, 0 up:
    Interface          Decrement    State
    Serial0            10          Down (administratively down)
R1#
```

```
R2#show standby
Ethernet1 - Group 1
  State is Active
  57 state changes, last state change 00:00:33
  Virtual IP address is 192.168.21.10
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (bia)
  Hello time 3 sec, hold time 10 sec
  Next hello sent in 2.648 secs
  Preemption enabled
  Active router is local
  Standby router is 192.168.21.1, priority 95 (expires in 7.096 sec)
  Priority 100 (default 100)
  IP redundancy name is "hsrp-Et0-1" (default)
R2#
```

```
R2#
```

```
R1#sh standby ethernet 1 brief
          P indicates configured to preempt.
          |
Interface  Grp Prio P State   Active addr   Standby addr   Group addr
Et0        1  95  P Standby 192.168.21.2   local          192.168.21.10
R1#
```

```
R2#sh standby ethernet 1 brief
          P indicates configured to preempt.
          |
Interface  Grp Prio P State   Active         Standby         Virtual IP
Et0        1  100 P Active  local         192.168.21.1   192.168.21.10
R2#
```

Notice that R1's standby priority has reduced from 105 to 95, and R2 has become the active router.

Summary

In the event of a failure of connectivity between ISP-A and R1, HSRP reduces the priority of the standby group on R1. R1 goes from an active state to a standby state. R2 goes from a standby state to an active state. The standby IP address 192.168.21.10 becomes active on R2, and hosts sending traffic to the Internet use R2 and ISP-B, providing an alternate path for outbound traffic.

For more information on the HSRP **standby track** command, refer to How to Use the **standby preempt** and **standby track** Commands.

Packets Coming from the Destination Toward the Local Network

According to the network policy defined in the Background Information section, since ISP-A is your primary path and ISP-B is the backup path for traffic coming towards 192.168.21.0/24 (for reasons such as a larger bandwidth connection towards ISP-A), you can append your own AS number in the BGP updates announced towards ISP-B in R2 to make the AS path through ISP-B appear longer. In order to do this, configure a route map for the BGP neighbor 192.168.42.4. In that route map, append your own AS with the **set as-path prepend** command. Apply this route-map to outbound updates to neighbor 192.168.42.4.

Note: In production, you must append the AS number more than once to ensure that the announced route becomes less preferred.

This is the BGP table in R6 for network 192.168.21.0 when the BGP connectivity between R1 to ISP-A and R2 to ISP-B is up:

```
R6#
show ip bgp 192.168.21.0
BGP routing table entry for 192.168.21.0/24, version 30
Paths: (2 available, best #1)
  Advertised to non peer-group peers:
    192.168.64.4
  300 100
    192.168.63.3 from 192.168.63.3 (10.5.5.5)
      Origin IGP, localpref 100, valid, external, best, ref 2
  400 100 100
    192.168.64.4 from 192.168.64.4 (192.168.64.4)
      Origin IGP, localpref 100, valid, external
```

BGP selects the best path as AS {300 100} through ISP-A because it has a smaller AS path length in comparison with the AS path {400 100 100} from ISP-B. The reason there is a longer AS path length from ISP-B is because of the AS path prepend configuration in R2.

When connectivity breaks between R1 and ISP-A, R6 must choose the alternate path through ISP-B to reach network 192.168.21.0/24 in AS 100:

```
R1(config)#interface s0
R1(config-if)#shut
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0, changed state to down
```

This is the BGP table in R6 for the network 192.168.21.0/24:

```
R6#show ip bgp 192.168.21.0
BGP routing table entry for 192.168.21.0/24, version 31
Paths: (1 available, best #1)
  Advertised to non peer-group peers:
    192.168.63.3
  400 100 100
    192.168.64.4 from 192.168.64.4 (192.168.64.4)
      Origin IGP, localpref 100, valid, external, best
```

Refer to Sample Configuration for BGP with Two Different Service Providers (Multihoming) for more information on BGP configurations in a multihomed network.

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- [Load Sharing with BGP in Single and Multihomed Environments: Sample Configurations](#)
 - [How BGP Routers Use the Multi-Exit Discriminator for Best Path Selection](#)
 - [Load Sharing with HSRP](#)
 - [HSRP Technology Support Page](#)
 - [BGP Technology Support Page](#)
 - [IP Routing Technology Support Page](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: May 10, 2006

Document ID: 13768
