

Reload Delay for Hot Standby Router Protocol Standby Commands Does Not Work on ASR920

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Problem](#)

[Workaround](#)

[Solution](#)

Introduction

This document describes the behaviour of Hot Standby Router Protocol (HSRP) reload delay commands on ASR920 series routers. The difference in interface behaviours across IOS-XE versions have been highlighted in order to correctly deploy an HSRP solution and to obtain predictable performance.

Prerequisites

Requirements

The reader should be familiar with Bridge-domains, Hot Standby Router Protocol (HSRP) and its related commands.

- [Cisco IOS First Hop Redundancy Protocols Command Reference](#)
- [First Hop Redundancy Protocols Configuration Guide, Cisco IOS XE Release 3S \(Cisco ASR 920\)](#)

Components Used

The information in this document is based on these software and hardware versions specified below:

- Cisco ASR 920 Series Aggregation Services Router
- Cisco IOS XE® Software Release that supports the ASR920 Series Routers

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

Problem


The ASR920 Series Routers are aggregation routers that are designed for Carrier Ethernet deployments and supports the HSRP feature. HSRP is deployed in groups of routers to select an active and a standby router to provide redundancy in the network. An active router is the router of choice for routing packets, a standby router is a router that takes over the routing duties when an active router fails, or when preset conditions are met. To ensure predictability and manageability, network administrators want a specific node to be active if that node is operational. This is achieved by the “standby preempt” feature of HSRP.

In large deployments, where the routing protocols could take a longer time to converge, the HSRP standby node pre-empting the active immediately as it boots up can cause traffic drops in the network. Ideally, the standby should takeover as active when it is ready to forward traffic. i.e., after its control-plane is up and upstream routing has converged. The below two commands can be used to delay the initialisation of the HSRP groups and to delay the preemption until the control plane is up. The reload keyword specifies additional delay in seconds which takes effect only after the reload of the router

- **standbydelayminimum** min-seconds [**reload** reload-seconds]
- **standby**[group-number] **preempt** [**delay**{ [**minimum** seconds] [**reload** seconds]]

A standby ASR920 router running IOS-XE 16.8.1c in a HSRP group boots up and preempts the active node immediately even with the reload-delay commands configured. This causes a traffic outage on large networks while HSRP is supposed to provide high network resiliency.

The issue was recreated with the router topology in image 1.



ge 1

Ima

Configuration

```
ASR-920-A configuration:
interface GigabitEthernet0/0/5
no ip address
negotiation auto
service instance 150 ethernet
 encapsulation dot1q 150
 rewrite ingress tag pop 1 symmetric
 bridge-domain 150
```

```
interface BDI150
ip address 10.0.1.2 255.255.255.0
standby delay minimum 5 reload 90
standby 80 ip 10.0.1.1
standby 80 priority 250
standby 80 preempt delay minimum 30 reload 90
```

ASR-920-B configuration:

```
interface GigabitEthernet0/0/5
no ip address
negotiation auto
service instance 150 ethernet
 encapsulation dot1q 150
 rewrite ingress tag pop 1 symmetric
 bridge-domain 150
```

```
interface BDI150
ip address 10.0.1.3 255.255.255.0
standby delay minimum 5 reload 90
standby 80 ip 10.0.1.1
standby 80 preempt delay minimum 30 reload 90
```

ASR-920-B is the active and once reloaded we get the logs as below which indicates that the delay timers did not work as expected. The timestamp in the logs indicate the router transitioned to active without the 90 seconds delay.


Logs

```
*Jul 27 01:17:11.493: %LINK-3-UPDOWN: Interface GigabitEthernet0/0/5, changed state to down
*Jul 27 01:17:15.805: %LINK-3-UPDOWN: Interface GigabitEthernet0/0/5, changed state to up
*Jul 27 01:17:16.506: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/0/5,
changed state to up
*Jul 27 01:17:34.166: %LINK-3-UPDOWN: Interface GigabitEthernet0/0/5, changed state to down
*Jul 27 01:17:36.802: %LINK-3-UPDOWN: Interface GigabitEthernet0/0/5, changed state to up
*Jul 27 01:17:44.818: %HSRP-5-STATECHANGE: BDI150 Grp 80 state Standby -> Active
```

Workaround

Use Tengig interface as the physical interface. If HSRP communication is over a tengig link i.e. the MAC addresses of both BDIs in the bridge-domain mac address table is learnt via a tengig interface, the HSRP timers will work as expected.

A working configuration is explained here and it uses the Topology in Image 2.



Ima

ge 2

Configuration

ASR-920-A configuration:

```
interface BDI20
ip address 10.0.2.2 255.255.255.0
standby delay minimum 5 reload 90
standby 21 ip 10.0.2.1
standby 21 timers msec 300 msec 900
standby 21 priority 250
standby 21 preempt delay minimum 30 reload 90
```

```
interface TenGigabitEthernet0/0/12
no ip address
service instance 20 ethernet
 encapsulation dot1q 20
 rewrite ingress tag pop 1 symmetric
 bridge-domain 20
```

ASR-920-B configuration:

```
interface BDI20
ip address 10.0.2.3 255.255.255.0
standby delay minimum 5 reload 90
standby 21 ip 10.0.2.1
standby 21 timers msec 300 msec 900
standby 21 preempt delay minimum 30 reload 90
```

```
interface TenGigabitEthernet0/0/12
no ip address
service instance 20 ethernet
 encapsulation dot1q 20
 rewrite ingress tag pop 1 symmetric
 bridge-domain 20
```

ASR-920-B is the active and once reloaded we get the logs as below which indicates that the delay timers worked as expected. The timestamp in the logs indicate the router transitioned to standby. After a delay of 90 seconds the router again takes over as active.

Logs

```
*Jul 22 21:53:35.735: %BDI_IF-5-CREATE_DELETE: Interface BDI20 is created
*Jul 22 21:53:36.497: %LINEPROTO-5-UPDOWN: Line protocol on Interface BDI20, changed state to
```

down

```
*Jul 22 21:54:21.850: %LINK-3-UPDOWN: Interface BDI20, changed state to up
*Jul 22 21:54:22.552: %LINEPROTO-5-UPDOWN: Line protocol on Interface BDI20, changed state to up
*Jul 22 21:55:54.346: %HSRP-5-STATECHANGE: BDI20 Grp 21 state Speak -> Standby
*Jul 22 21:57:22.430: %HSRP-5-STATECHANGE: BDI20 Grp 21 state Standby -> Active
```

Solution

The reload delay timer starts at the first interface-up event. If the interface goes down while the timer is counting down, the timer is killed and the minimum delay timer will take over. Cisco have identified that in certain IOS versions, the interface flaps twice during the router boot up. The first interface down event kills the reload timer and hence as the interface comes up second time the reload-delay does take effect.

The root cause of the issue is the physical interface flap event at the time of router boot up. This is documented by the defect [CSCuh56657](#) and is fixed from IOS-XE 16.9.1a onwards.

Troubleshoot commands

- **show standby BDI <int num>**
- **show standby brief**
- **show standby delay**
- **show standby neighbors**
- **Show logging**

show standby BDI command can be used to confirm which HSRP timer is currently running on the Bridge Domain Interface (BDI) interface. The command output shows that in the problematic state when the interface flaps, the reload timer is overridden by the minimum timer. This causes the pre-emption to occur beforehand.

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if reload delay, 72 secs remaining)
Virtual IP address is 10.0.1.1
```

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if min delay, 1 secs remaining)
Virtual IP address is 10.0.1.1
```

show standby brief command displays the router role.

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if reload delay, 72 secs remaining)
Virtual IP address is 10.0.1.1
```

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if min delay, 1 secs remaining)
Virtual IP address is 10.0.1.1
```

show standby delay command displays timer values.

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
```

```
State is Init (if reload delay, 72 secs remaining)
Virtual IP address is 10.0.1.1
```

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if min delay, 1 secs remaining)
Virtual IP address is 10.0.1.1
```

- **show standby neighbors** command displays HSRP neighbours.

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if reload delay, 72 secs remaining)
Virtual IP address is 10.0.1.1
```

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if min delay, 1 secs remaining)
Virtual IP address is 10.0.1.1
```

- **Show logging** command will display the HSRP logs.

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if reload delay, 72 secs remaining)
Virtual IP address is 10.0.1.1
```

```
ASR-920-A#show standby bdi 150
BDI150 - Group 80
State is Init (if min delay, 1 secs remaining)
Virtual IP address is 10.0.1.1
```