

Troubleshoot Border Gateway Protocol Routes that Do Not Advertise

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Conventions](#)

[Background Information](#)

[Routes Announced With a Basic Network Statement](#)

[Use these Steps to Troubleshoot](#)

[Routes Announced With the Network Statement with a Mask](#)

[Use these Steps to Troubleshoot](#)

[Routes Announced With the aggregate-address Command](#)

[Use these Steps to Troubleshoot this Problem](#)

[Unable to Announce iBGP-Learned Routes](#)

[Use these Steps to Troubleshoot](#)

[Routes Announced with Redistribute Static](#)

[Related Information](#)

Introduction

This document describes how to troubleshoot when a Border Gateway Protocol (BGP) router does not announce BGP routes to peers.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions. However, the outputs shown in the document are based on Cisco 2500 series routers that run Cisco IOS® software version 12.2(24)a.

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 [Cisco Technical Tips Conventions](#).

Background Information

This document provides a systematic approach to help troubleshoot situations when a Border Gateway Protocol (BGP) router does not announce BGP routes to peers.

There are multiple ways in which a prefix is added to a BGP table and announced to peers:

- Issue the basic **network** command under router BGP. This method is used to originate BGP routes from the autonomous system (AS). For more information, refer to the network command section of the [BGP Case Studies 1](#).
- Redistribute Interior Gateway Protocol (IGP) or a static configuration.
- Propagate BGP routes learned from other internal BGP (iBGP) or external BGP (eBGP) peers.



Note: Only the best paths received from BGP peers are propagated. Refer to [BGP Best Path Selection Algorithm](#) for more information on best path selection.

- Issue the **aggregate-address** command. Refer to [Understanding Route Aggregation in BGP](#) for more information.

Routes Announced With a Basic Network Statement

When routes are announced with a basic network statement, the behavior of the **network** command varies based on whether **auto-summary** is enabled or disabled. When auto-summary is enabled, it summarizes the locally originated BGP networks (network x.x.x.x) to their classful boundaries (auto-summary is enabled by default in BGP). If a subnet exists in the routing table and these three conditions are satisfied, any subnet (component route) of that classful network in the local routing table prompts BGP to install the classful network into the BGP table:

- Auto-summary enabled
- Classful network statement for a network in the routing table
- Classful mask on that network statement

When auto-summary is disabled, the routes introduced locally into the BGP table are not summarized to their classful boundaries.


For example, BGP introduces the classful network 10.0.0.0 mask 255.0.0.0 in the BGP table if these conditions are met:

- The subnet in the routing table is 10.75.75.0 mask 255.255.255.0.
- You configure network 10.0.0.0 under the **router bgp** command.
- Auto-summary is enabled.

If these conditions are not all met, BGP does not install an entry in the BGP table unless there is an exact match in the IP routing table.

Use these Steps to Troubleshoot

With auto-summary enabled on R101, the router is not able to announce classful network 172.16.0.0/16 to R102.



Router Can Not Announce

1. Check to see if R101 announces 172.16.0.0/16 to R102. The output shown confirms that R101 does not announce 172.16.0.0/16 to R102.

```
<#root>
R101#
show ip bgp neighbors 10.10.10.2 advertised-routes
```

```
Total number of prefixes 0
R101#
```

2. Check to see if the configuration runs. The example shown illustrates that R101 is configured with classful network statement. Auto-summary is enabled by default in the Cisco IOS software version used for this scenario.

```
<#root>
R101#
show running-config | begin bgp
router bgp 1
network 172.16.0.0
neighbor 10.10.10.2 remote-as 2
[...]
```

3. Check to see if you have a component route (a classful route or a subnet route) of network 172.16.0.0/16 in the routing table.

```
<#root>
R101#
```

```
show ip route 172.16.0.0 255.255.0.0 longer-prefixes
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
```

Gateway of last resort is not set

```
R101#sh ip route 172.16.0.0 255.255.0.0 longer-prefixes
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
```

Gateway of last resort is not set

```
R101#
```

4. Because there is no component route (no classful route or subnet route) in the R101 IP routing table, the network 172.16.0.0 in not installed in the BGP table. The minimum requirement for a prefix configured under the **network** command to be installed in a BGP table is to have a component route in the IP routing table. So, ensure that R101 has a component route for network 172.16.0.0/16. This is done if the route learns it through IGP or through a static configuration. In the example shown, the static route is configured to null 0.

```
<#root>
```

```
R101(config)#
```

```
ip route 172.16.10.0 255.255.255.0 null0 200
```

5. As soon as the IP routing table has a component route for 172.16.0.0/16, BGP installs a classful network in the BGP table.

```
<#root>
```

```
R101#
```

```
show ip route 172.16.0.0 255.255.0.0 longer-prefixes
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PFR

Gateway of last resort is not set

172.16.0.0/24 is subnetted, 1 subnets

s 172.16.10.0 is directly connected, Null0

To bring the change into effect in BGP and start announcing network 172.16.0.0/16 to R102, you must either clear the BGP neighbor or do a soft reset to peer. This example shows a soft reset outbound to peer 10.10.10.2 to bring the changes into effect.


```
<#root>
R101#
clear ip bgp 10.10.10.2 soft out
R101#
```

6. The **show ip bgp** command confirms that classful network 172.16.0.0/16 is introduced into BGP.

```
<#root>
R101#
show ip bgp | include 172.16.0.0
*>
172.16.0.0
          0.0.0.0          0          32768 i
R101#
```

7. Confirm that R101 announces routes to R102.

```
<#root>
R101#
show ip bgp neighbors 10.10.10.2 advertised-routes | include 172.16.0.0
*> 172.16.0.0          0.0.0.0          0          32768 i
```

 **Note:** With auto-summary disabled, BGP installs network 172.16.0.0/16 when there is an exact match to the route in the routing table. If there are subnet routes, but no exact match to the route (172.16.0.0/16) in the routing table, then BGP does not install the network 172.16.0.0/16 in the BGP table.


Routes Announced With the Network Statement with a Mask

Networks that fall on a major net boundary (255.0.0.0, 255.255.0.0, or 255.255.255.0) do not need to have a mask included. For example, the **network 172.16.0.0** command is sufficient to send the prefix 172.16.0.0/16 into the BGP table. However, networks that do not fall on major net boundaries are required to have a network statement with a mask, such as **network 172.16.10.0 mask 255.255.255.0**.

An exact route in the routing table is required for a network statement with a mask in order for it to be installed into a BGP table.

Use these Steps to Troubleshoot

R101 is unable to announce network 172.16.10.0/24 to R102.



Cannot Announce Network 172.16.10.0/24 to R102

1. Check to see if R101 announces the 172.16.10.0/24 prefix to R102.

```
<#root>
R101#
show ip bgp neighbors 10.10.10.2 advertised-routes
R101#
```

OR

This command can be used to check whether the routes are advertised:

```
<#root>
R101#
show ip bgp 172.16.10.0/24
R101# BGP routing table entry for 172.16.10.0/24, version 24480684
```


```
Bestpath Modifiers: deterministic-med
Paths: (4 available, best #3)
```

```
Not advertised to any peer
```

This output confirms that R101 is not announcing 192.168.32.0/22 to R102.

2. Check to see if the configuration runs.

```
<#root>
R101#
show run | begin bgp
router bgp 1
network 172.16.10.0
```

 **Note:** You need to originate network 172.16.10.0/24. This network does not fall on the boundary of a Class B network (255.255.0.0). A network statement with mask 255.255.255.0 needs to be configured to make it work.

3. After a network statement with mask is configured, the **show run** command shows output similar to this:

```
<#root>
R101#
show run | begin bgp
router bgp 1
network 172.16.10.0 mask 255.255.255.0
```

1. Check to see if the route is in the BGP routing table.

```
<#root>
R101#
show ip bgp | include 172.16.10.0
R101#
```

Network 172.16.10.0/24 does not exist in the BGP table.

2. Check to see if there is an exact route in the IP routing table. The output shown confirms that there is

not an exact route in the routing table.

```
<#root>
R101#
show ip route 172.16.10.0 255.255.255.0

% Network not in table
R101#
```

3. Decide which routes you want to originate. Then either fix the IGP, or configure static routes.

```
<#root>
R101(config)#
ip route 172.16.10.0 255.255.255.0 null0 200
```

4. Check the IP routing table.

```
<#root>
R101#
show ip route 172.16.10.0 255.255.255.0 longer-prefixes

[..]
    172.16.0.0/24 is subnetted, 1 subnets
S
    172.16.10.0 is directly connected, Null0
```

5. Verify that the routes are in the BGP table.

```
<#root>
R101#
show ip bgp | include 172.16.10.0

*> 172.16.10.0/24    0.0.0.0                0                32768 i
```

6. To bring the change into effect in BGP and start announcing network 172.16.10.0/24 to R102, you must either clear the BGP neighbor, or do a soft reset to the peer. This example uses a soft reset outbound to peer 10.10.10.2.


```
<#root>
R101#
clear ip bgp 10.10.10.2 soft out
```

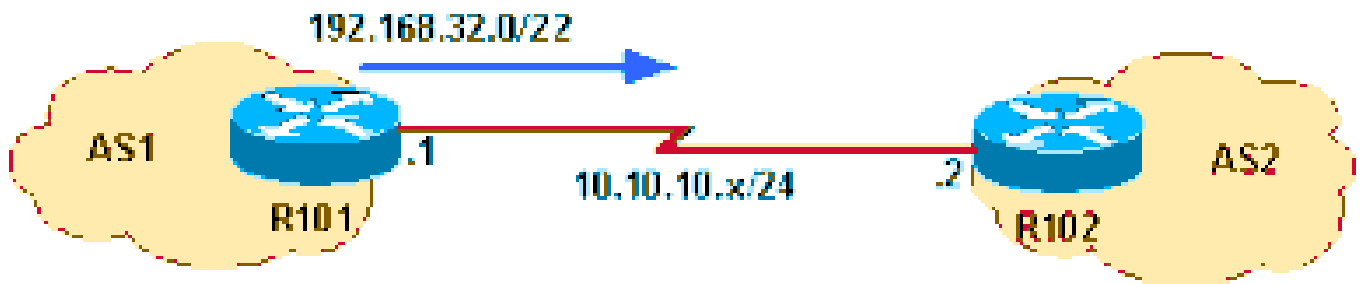
7. Confirm that routes are advertised to R102.

```
<#root>
R101#
show ip bgp neighbors 10.10.10.2 advertised-routes | include 172.16.10.0
*> 172.16.10.0/24 0.0.0.0 0 32768 i
```

Routes Announced With the aggregate-address Command

BGP allows the aggregation of specific routes into one route with the **aggregate-address <address> <mask>** command. Aggregation applies to routes that exist in the BGP routing table. This is in contrast to the **network** command, which applies to the routes that exist in IP routing table. Aggregation can be performed if at least one or more of the specific routes of the aggregate address exists in the BGP routing table. Refer to [Understanding Route Aggregation in BGP](#) for more information on BGP aggregation and associated attributes.

Use these Steps to Troubleshoot this Problem



Routes Announced With the aggregate-address Command

In this network diagram, R101 is unable to announce the aggregate address 192.168.32.0/22 to R102. Network 192.168.32.0/22 aggregates these three Class C address spaces:

- 192.168.33.0/24
- 192.168.35.0/24
- 192.168.35.0/24

1. Confirm that R101 is not announcing 192.168.32.0/22 to R102.

```
<#root>
```

```
R101#  
show ip bgp neighbors 10.10.10.2 advertised-routes | include 192.168.32.0  
  
R101#
```

2. Check to see if the configuration runs.

```
<#root>  
router bgp 1  
  [..]  
  
  aggregate-address 192.168.32.0 255.255.252.0 summary-only  
    neighbor 10.10.10.2 remote-as 2
```

R101 is configured to announce only the aggregate address to R102 with the summary-only attribute.

3. Check the IP routing table.

```
<#root>  
  
R101#  
show ip route 192.168.32.0 255.255.252.0 longer-prefixes  
  
[..]  
S    192.168.33.0/24 is directly connected, Null0
```

The IP routing table has the component route of aggregate 192.168.32.0/22; however, for an aggregate address to be announced to a peer, a component route must exist in the BGP routing table rather than in the IP routing table. The IP routing table has the component route of aggregate 192.168.32.0/22; however, for an aggregate address to be announced to a peer, a component route must exist in the BGP routing table rather than in the IP routing table.

4. Check to see if a component route exists in the BGP routing table.

```
<#root>  
  
R101#  
show ip bgp 192.168.32.0 255.255.252.0 longer  
  
R101#
```

The output confirms that the BGP table does not have a component route, so the next logical step is to ensure that a component route exists in the BGP table.

1. In this example, a component route 192.168.33.0 is installed into the BGP table with the **network** command.

```
<#root>
R101(config)#
router bgp 1
R101(config-router)#
network 192.168.33.0
```

2. Check to see if the component route exists in the BGP table.

```
<#root>
R101#
show ip bgp 192.168.32.0 255.255.252.0 longer-prefixes

BGP table version is 8, local router ID is 10.10.20.1
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
*> 192.168.32.0/22  0.0.0.0                    32768 i
s> 192.168.33.0     0.0.0.0              0          32768 i
R101#
```

The **s** means that the component route is suppressed due to the summary-only argument.

3. Confirm that the aggregate is announced to R102.

```
<#root>
R101#
show ip bgp neighbors 10.10.10.2 advertised-routes | include 192.168.32.0/22


*> 192.168.32.0/22  0.0.0.0
```

Unable to Announce iBGP-Learned Routes

A BGP router with synchronization enabled does not advertise iBGP-learned routes to other eBGP peers if it is not able to validate those routes in its IGP. The IGP has a route to iBGP-learned routes, the router announces the iBGP routes to eBGP peers. Otherwise, the router treats the route as not synchronized with IGP and does not advertise it. To prevent BGP from iBGP routes validation in IGP, disable synchronization with the **no synchronization** command under router BGP. Refer to the [Synchronization](#) section of [BGP Case Studies](#) for more information.

Use these Steps to Troubleshoot

In the diagram shown, R101 learns prefix 10.130.130.0/24 from R103 through iBGP and is unable to announce it to eBGP peer R102.



R101 is Unable to Announce to eBGP peer R102.

1. First check R101.

```
<#root>
R101#
show ip bgp neighbors 10.10.20.2 advertised-routes | include 10.130.130.0
R101#
```

This output confirms that R101 is not announcing prefix 10.254.255.255/24 to R102.

Look at the BGP table on R101:

```
<#root>
R101#
show ip bgp 10.130.130.0 255.255.255.0 longer

BGP table version is 4, local router ID is 10.10.20.1
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
* i10.130.130.0/24  10.10.20.3         0      100      0  i

R101#
```

Network 10.130.130.0/24 exists in the BGP table. However, the network 10.130.130.0/24 does not have the status code of best route (>). This means that the [BGP Best Path Selection Algorithm](#) did not choose this prefix as the best path. Since only the best paths are announced to BGP peers, network 10.130.130.0/24 is not announced to R102. Next, you need to troubleshoot why the BGP path selection criteria did not select


this network as the best route.

1. Examine the output of the **show ip bgp prefix** command to give you more details on why the prefix was not chosen as the best route nor installed in IP routing table.

```
<#root>
R101#
show ip bgp 10.130.130.0

BGP routing table entry for 10.130.130.0/24, version 4
Paths: (1 available, no best path)
  Not advertised to any peer
  Local
    10.10.20.3 from 10.10.20.3 (10.130.130.3)
      Origin IGP, metric 0, localpref 100, valid, internal, not synchronized
```

The output shows that prefix 10.130.130.0/24 is not synchronized.

 **Note:** Before the identification of Cisco bug ID [CSCdr90728](#) (BGP paths are not marked as not synchronized), the **show ip bgp prefix** command did not show the paths marked as not synchronized. This problem is corrected in Cisco IOS Software Releases 12.1(4) and later.

2. Check to see if the BGP configuration runs.

This output shows that BGP synchronization is enabled. BGP synchronization is enabled by default in Cisco IOS software.

```
<#root>
R101#
show ip protocols

Routing Protocol is "bgp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set

IGP synchronization is enabled

Automatic route summarization is disabled
Neighbor(s):
  Address           FiltIn FiltOut DistIn DistOut Weight RouteMap
  10.10.10.2
  10.10.20.3
Maximum path: 1
Routing for Networks:
Routing Information Sources:
  Gateway           Distance     Last Update
  10.10.20.3         200          01:48:24
Distance: external 20 internal 200 local 200
```

3. Configure BGP to disable synchronization. Issue the **no synchronization** command under router BGP.

```
<#root>
R101(config)#
router bgp 1
R101(config-router)#
no synchronization

R101#
show ip protocols

Routing Protocol is "bgp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set

IGP synchronization is disabled

Automatic route summarization is disabled
Neighbor(s):
  Address           FiltIn FiltOut DistIn DistOut Weight RouteMap
  10.10.10.2
  10.10.20.3
Maximum path: 1
Routing for Networks:
Routing Information Sources:
  Gateway           Distance      Last Update
  10.10.20.3         200           01:49:24
Distance: external 20 internal 200 local 200
```

During the next run of the BGP scanner, which scans the BGP table every 60 seconds and makes decision based on BGP path selection criteria, network 10.130.130.0 is installed (since the synchronization is disabled). This means that the maximum time for the route to be installed is 60 seconds, but it can be less, based on when the **no synchronization** command is configured, and when the next instance of the BGP scanner occurs. So, it is best to wait 60 seconds before the next step of verification.

1. Verify that the route has been installed.

The output shown confirms that prefix 10.130.130.0/24 is the best route; therefore, it is installed into the IP routing table and is propagated to peer 10.10.10.2.

```
<#root>
R101#
show ip bgp 10.130.130.0

BGP routing table entry for 10.254.255.255/24, version 5
Paths: (1 available, best #1, table Default-IP-Routing-Table)
```

Advertised to non peer-group peers:

```
10.10.10.2
Local
 10.10.20.3 from 10.10.20.3 (10.130.130.3)
  Origin IGP, metric 0, localpref 100, valid, internal, best
```

R101#

```
show ip bgp neighbors 10.10.10.2 advertised-routes | include 10.130.130.0/24
*>i10.130.130.0/24          10.10.20.3          0    100    0 i
```

Routes Announced with Redistribute Static

If the routers are connected with two links, and the routes are learned through BGP and floating static routes, the floating static routes are installed in the routing table. This occurs if the static routes are redistributed in the case of BGP route failure. If the BGP routes come back online, the floating static routes in the routing table are not changed to reflect the BGP routes.

This issue can be solved if you remove the **redistribute static** command under the BGP process to avoid the prioritization of floating static routes over BGP routes.

Related Information

- [Why Do BGP Neighbors Toggle Between Idle, Connect, and Active States?](#)
- [What Does the "#%BGP-3-INSUFCHUNKS: Insufficient chunk pools for aspath" Error Message Mean?](#)
- [BGP: Frequently Asked Questions](#)
- [Troubleshoot Common BGP Issues](#)
- [Routing-BGP Support Page](#)
- [Support and Downloads](#)