

# Network Configuration Example

## Configuring BGP Confederations

Release

NCE0042



Modified: 2017-01-20

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
www.juniper.net

Copyright © 2017, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

*Network Configuration Example Configuring BGP Confederations*

NCE0042

Copyright © 2017, Juniper Networks, Inc.

All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

## **END USER LICENSE AGREEMENT**

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

<b>Chapter 1</b>	<b>Configuring BGP Confederations . . . . .</b>	<b>5</b>
	About This Network Configuration Example . . . . .	5
	Understanding BGP Confederations . . . . .	5
	Example: Configuring BGP Confederations . . . . .	6



## CHAPTER 1

# Configuring BGP Confederations

- [About This Network Configuration Example on page 5](#)
- [Understanding BGP Confederations on page 5](#)
- [Example: Configuring BGP Confederations on page 6](#)

### About This Network Configuration Example

---

This network configuration example provides an overview of BGP confederations and describes how to configure a confederated BGP network. A BGP confederation breaks up a large autonomous system (AS) into subautonomous systems (sub-ASs). This is one method that can be used to resolve scaling issues that can result from the full mesh requirement of BGP.

### Understanding BGP Confederations

---

BGP confederations are another way to solve the scaling problems created by the BGP full mesh requirement. BGP confederations effectively break up a large autonomous system (AS) into subautonomous systems (sub-ASs). Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64,512 and 65,535.

Within a sub-AS, the same internal BGP (IBGP) full mesh requirement exists. Connections to other confederations are made with standard external BGP (EBGP), and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

The confederation AS appears whole to other confederation ASs. The AS path received by other ASs shows only the globally assigned AS number. It does not include the confederation sequence or the privately assigned sub-AS numbers. The sub-AS numbers are removed when the route is advertised out of the confederation AS. [Figure 1 on page 6](#) shows an AS divided into four confederations.

Figure 1: BGP Confederations

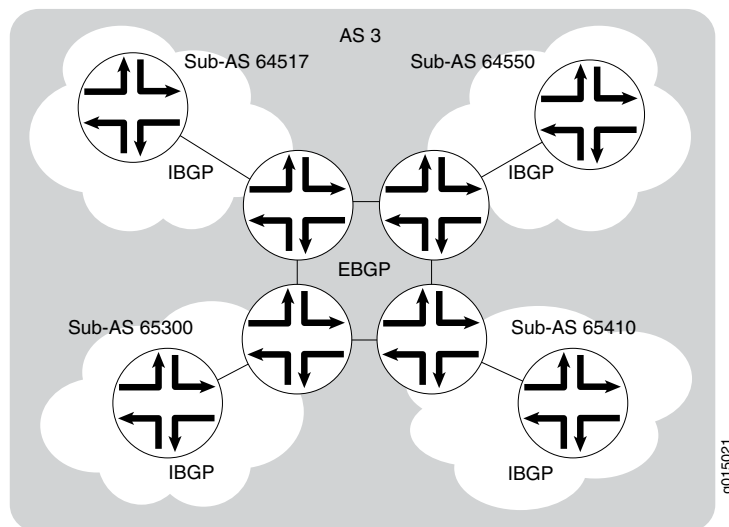


Figure 1 on page 6 shows AS 3 divided into four sub-ASs, 64517, 64550, 65300, and 65410, which are linked through EBGP sessions. Because the confederations are connected by EBGP, they do not need to be fully meshed. EBGP routes are readvertised to other sub-ASs.

#### Related Documentation

- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)
- [Understanding BGP](#)
- [Example: Configuring BGP Confederations on page 6](#)

## Example: Configuring BGP Confederations

This example shows how to configure BGP confederations.

- [Requirements on page 6](#)
- [Overview on page 7](#)
- [Configuration on page 7](#)
- [Verification on page 9](#)

### Requirements

- Configure network interfaces.
- Configure external peer sessions. See *Example: Configuring External BGP Point-to-Point Peer Sessions*.
- Configure interior gateway protocol (IGP) sessions between peers.
- Configure a routing policy to advertise the BGP routes.

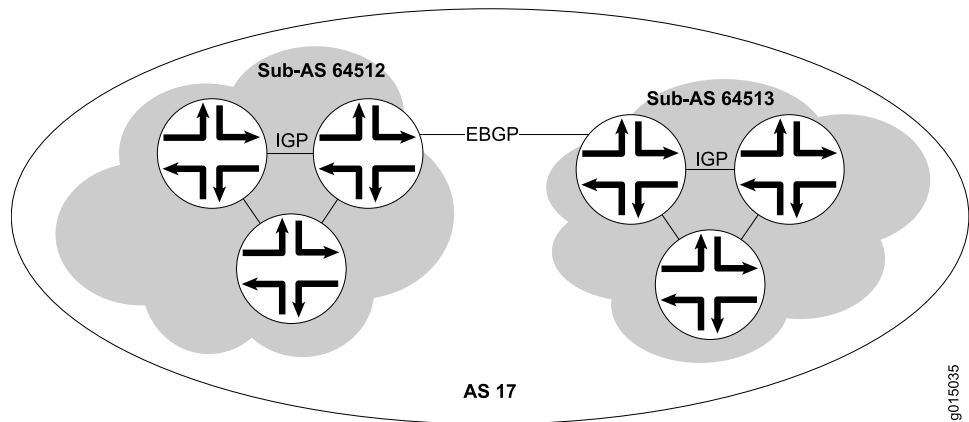
## Overview

Within a BGP confederation, the links between the confederation member autonomous systems (ASs) must be external BGP (EBGP) links, not internal BGP (IBGP) links.

Similar to route reflectors, BGP confederations reduce the number of peer sessions and TCP sessions to maintain connections between IBGP routing devices. BGP confederation is one method used to solve the scaling problems created by the IBGP full mesh requirement. BGP confederations effectively break up a large AS into subautonomous systems. Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64512 and 65535. Within a sub-AS, the same IBGP full mesh requirement exists. Connections to other confederations are made with standard EBGP, and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

Figure 2 on page 7 shows a sample network in which AS 17 has two separate confederations: sub-AS 64512 and sub-AS 64513, each of which has multiple routers. Within a sub-AS, an IGP is used to establish network connectivity with internal peers. Between sub-ASs, an EBGP peer session is established.

Figure 2: Typical Network Using BGP Confederations



## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

### All Devices in Sub-AS 64512

```
set routing-options autonomous-system 64512
set routing-options confederation 17 members 64512
set routing-options confederation 17 members 64513
set protocols bgp group sub-AS-64512 type internal
set protocols bgp group sub-AS-64512 local-address 192.168.5.1
set protocols bgp group sub-AS-64512 neighbor 192.168.8.1
set protocols bgp group sub-AS-64512 neighbor 192.168.15.1
```

Border Device in Sub-AS 64512	<pre>set protocols bgp group to-sub-AS-64513 type external set protocols bgp group to-sub-AS-64513 peer-as 64513 set protocols bgp group to-sub-AS-64513 neighbor 192.168.5.2</pre>
All Devices in Sub-AS 64513	<pre>set routing-options autonomous-system 64513 set routing-options confederation 17 members 64512 set routing-options confederation 17 members 64513 set protocols bgp group sub-AS-64513 type internal set protocols bgp group sub-AS-64513 local-address 192.168.5.2 set protocols bgp group sub-AS-64513 neighbor 192.168.9.1 set protocols bgp group sub-AS-64513 neighbor 192.168.16.1</pre>
Border Device in Sub-AS 64512	<pre>set protocols bgp group to-sub-AS-64512 type external set protocols bgp group to-sub-AS-64512 peer-as 64512 set protocols bgp group to-sub-AS-64512 neighbor 192.168.5.1</pre>
Step-by-Step Procedure	<p>This procedure shows the steps for the devices that are in sub-AS 64512.</p> <p>The <b>autonomous-system</b> statement sets the sub-AS number of the device.</p> <p>The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i>.</p> <p>To configure BGP confederations:</p> <ol style="list-style-type: none"><li>1. Set the sub-AS number for the device. <pre>[edit routing-options] user@host# set autonomous-system 64512</pre></li><li>2. In the confederation, include all sub-ASs in the main AS. <p>The number 17 represents the main AS. The <b>members</b> statement lists all the sub-ASs in the main AS.</p><pre>[edit routing-options confederation] user@host# set 17 members 64512 user@host# set 17 members 64513</pre></li><li>3. On the border device in sub-AS 64512, configure an EBGp connection to the border device in AS 64513. <pre>[edit protocols bgp group to-sub-AS-64513] user@host# set type external user@host# set neighbor 192.168.5.2 user@host# set peer-as 64513</pre></li><li>4. Configure an IBGP group for peering with the devices within sub-AS 64512. <pre>[edit protocols bgp group sub-AS-64512] user@host# set type internal user@host# set local-address 192.168.5.1 user@host# neighbor 192.168.8.1 user@host# neighbor 192.168.15.1</pre></li></ol>

**Results** From configuration mode, confirm your configuration by entering the **show routing-options** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show routing-options
autonomous-system 64512;
confederation 17 members [ 64512 64513 ];

user@host# show protocols
bgp {
  group to-sub-AS-64513 { # On the border devices only
    type external;
    peer-as 64513;
    neighbor 192.168.5.2;
  }
  group sub-AS-64512 {
    type internal;
    local-address 192.168.5.1;
    neighbor 192.168.8.1;
    neighbor 192.168.15.1;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps for sSub-AS 64513.

## Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 9](#)
- [Verifying BGP Groups on page 10](#)
- [Verifying BGP Summary Information on page 11](#)

### Verifying BGP Neighbors

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From the CLI, enter the **show bgp neighbor** command.

## Sample Output

```

user@host> show bgp neighbor
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: Sync
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh

  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90

```

```
Keepalive Interval: 30
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 4
  Received prefixes: 6
  Suppressed due to damping: 0
Table inet6.0 Bit: 20000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 2
  Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10
```

**Meaning** The output shows a list of the BGP neighbors with detailed session information. Verify the following information:

- Each configured peering neighbor is listed.
- For **State**, each BGP session is **Established**.
- For **Type**, each peer is configured as the correct type (either internal or external).
- For **AS**, the AS number of the BGP neighbor is correct.

---

### Verifying BGP Groups

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From the CLI, enter the **show bgp group** command.

### Sample Output

```
user@host> show bgp group
Group Type: Internal  AS: 10045      Local AS: 10045
Name: pe-to-asbr2
Export: [ match-all ]
Total peers: 1      Established: 1
10.0.0.4+179
bgp.l3vpn.0: 1/1/0
vpn-green.inet.0: 1/1/0
Flags: Export Eval
```

```

Groups: 1   Peers: 1   External: 0   Internal: 1   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
bgp.13vpn.0      1      1      0      0      0      0

```

**Meaning** The output shows a list of the BGP groups with detailed group information. Verify the following information:

- Each configured group is listed.
- For **AS**, each group's remote AS is configured correctly.
- For **Local AS**, each group's local AS is configured correctly.
- For **Group Type**, each group has the correct type (either internal or external).
- For **Total peers**, the expected number of peers within the group is shown.
- For **Established**, the expected number of peers within the group have BGP sessions in the **Established** state.
- The IP addresses of all the peers within the group are present.

### Verifying BGP Summary Information

**Purpose** Verify that the BGP configuration is correct.

**Action** From the CLI, enter the **show bgp summary** command.

### Sample Output

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0      6      4      0      0      0      0
Peer      AS      InPkt  OutPkt  OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2   65002   88675   88652    0      2      42:38 2/4/0
0/0/0
10.0.0.3   65002   54528   54532    0      1      2w4d22h 0/0/0
0/0/0
10.0.0.4   65002   51597   51584    0      0      2w3d22h 2/2/0
0/0/0

```

**Meaning** The output shows a summary of BGP session information. Verify the following information:

- For **Groups**, the total number of configured groups is shown.
- For **Peers**, the total number of BGP peers is shown.
- For **Down Peers**, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.
- Under **Peer**, the IP address for each configured peer is shown.
- Under **AS**, the peer AS for each configured peer is correct.
- Under **Up/Dwn State**, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of

routes being damped (such as 0/0/0). If the field is **Active**, it indicates a problem in the establishment of the BGP session.

**Related  
Documentation**

- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*
- [Understanding BGP Confederations on page 5](#)
- *BGP Configuration Overview*