

BGP Command Reference

Document revision 1.5 (Thu Sep 22 12:50:17 GMT 2005)

This document applies to V2.9

Table of Contents

[Table of Contents](#)

[General Information](#)

[Summary](#)

[Quick Setup Guide](#)

[Specifications](#)

[Related Documents](#)

[Description](#)

[Additional Documents](#)

[Instances](#)

[Description](#)

[Property Description](#)

[Peers](#)

[Description](#)

[Property Description](#)

General Information

Summary

The Border Gateway Protocol (BGP) allows setting up an interdomain dynamic routing system that automatically updates routing tables of devices running BGP in case of network topology changes.

MikroTik RouterOS supports BGP Version 4, as defined in RFC1771.

Starting from version v2.9 MikroTik RouterOS has a brand new BGP implementation, which provides advanced functionality not available in the previous versions.

Quick Setup Guide

To configure a BGP instance with AS number of **200** and establish a BGP session to the **10.0.11.11** peer from the AS **100**, redistributing connected and static routes only, you should do the following:

- Configure default BGP instance:

```
[admin@rb12] > /routing bgp instance set default as=200 redistribute-static=yes
redistribute-connected=yes [admin@rb12] > /routing bgp instance print Flags: X -
disabled 0 as=200 router-id=0.0.0.0 redistribute-static=yes
redistribute-connected=yes redistribute-rip=no redistribute-ospf=no
redistribute-other-bgp=no name="default" out-filter="" [admin@rb12] >
```
- Add BGP peer:

```
[admin@rb12] > /routing bgp peer add remote-address=10.0.11.11 remote-as=100
instance=default [admin@rb12] > /routing bgp peer print Flags: X - disabled 0
```

```
remote-address=10.0.11.11 remote-as=100 multihop=no in-filter="" out-filter=""
keepalive-time=0s hold-time=0s ttl=1 [admin@rb12] >
```

Note, that the peer should be configured accordingly in order BGP to work.

Attention! In this scenario the router has no input or output filters configured. This means that it can redistribute lots of unnecessary or harmful information to its peers. Always consider configuring proper routing filters before you configure BGP peering.

Specifications

Packages required: *routing-test*

License required: *level3*

Home menu level: */routing bgp*

Standards and Technologies: [RFC1771](#)

Hardware usage: *requires additional RAM for storing routing information (128MB recommended)*

Related Documents

- [Software Package Management](#)
- [IP Addresses and ARP](#)
- [Routes, Equal Cost Multipath Routing, Policy Routing](#)
- [BGP Routing Filters](#)

Description

The Border Gateway Protocol (BGP) is the core routing protocol of the Internet. It maintains a table of routes 'prefixes', which specify network layer reachability information (NLRI) between autonomous systems (AS). BGP is described as path vector protocol or policy routing protocol, referring to the way it chooses the best route towards destination. Unlike many other routing protocols, BGP does not use technical metrics to select the best path but rather administrative policies. The current version of BGP, Border Gateway Protocol 4, is specified in RFC 1771.

The routes learned by BGP protocol are installed in the route list with the distance of 200 for iBGP (Internal BGP) routes and of 20 for eBGP (External BGP) routes.

Additional Documents

- <http://www.ietf.org/rfc/rfc1771.txt>
- <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm>
- <http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2003.htm>

Instances

Description

Home menu level: */routing bgp instance*

Property Description

as (*integer*: 0..65535) - BGP autonomous system number

name (*name*; default: "") - BGP instance name

out-filter (*name*; default: "") - output routing filter used by this BGP instance

redistribute-connected (*yes | no*; default: **no**) - if enabled, the router will redistribute the information about all connected routes, i.e., routes to the networks that can be directly reached

redistribute-ospf (*yes | no*; default: **no**) - if enabled, the router will redistribute the information about all routes learned by the OSPF protocol

redistribute-other-bgp (*yes | no*; default: **no**) - specifies whether this BGP instance should redistribute to its peers routes learned by other BGP instances

redistribute-rip (*yes | no*; default: **no**) - if enabled, the router will redistribute the information about all routes learned by RIP protocol

redistribute-static (*yes | no*; default: **no**) - if enabled, the router will redistribute the information about all static routes added to its routing database, i.e., routes that have been created using the `/ip route add` command on the router

router-id (*IP address*; default: **0.0.0.0**) - the router identification string in form of an IP address. If no router-id is specified, it will be selected automatically based on the routing information

Peers

Home menu level: */routing bgp peer*

Description

You need to specify the BGP peer with whom you want to exchange the routing information. The BGP exchanges routing information only if it can establish a TCP connection to its peer. You can add as many peers as required.

Property Description

hold-time (*time*) - specifies the BGP Hold Time value to use when negotiating with peers. According to BGP specifications, if router does not receive successive KEEPALIVE and/or UPDATE and/or NOTIFICATION messages within the period specified in the Hold Time field of the OPEN message, then the BGP connection to the peer will be closed

in-filter (*name*; default: "") - name of the routing filter that is applied to incoming routing update messages

keepalive-time (*time*) - specifies the time interval between successive KEEPALIVE messages. BGP process will negotiate the keepalive time with the neighbour upon connection establishment

multihop (*yes | no*; default: **no**) - if enabled, allows BGP sessions, even when the neighbour is not on a directly connected segment. The multihop session is not established if the only route to the multi-hop peer's address is the default route (0.0.0.0/0)

out-filter (*name*; default: "") - name of the routing filter that is applied to outgoing routing update messages

remote-address (*IP address*; default: **0.0.0.0**) - address of the remote peer

remote-as (*integer*; default: **0**) - AS number of the remote peer