

BGP (Border Gateway Protocol)

Document revision 1.2 (Thu Mar 04 19:34:34 GMT 2004)

This document applies to MikroTik RouterOS V2.8

Table of Contents

[Table of Contents](#)

[General Information](#)

[Summary](#)

[Specifications](#)

[Related Documents](#)

[Description](#)

[Additional Documents](#)

[BGP Setup](#)

[Property Description](#)

[Notes](#)

[Example](#)

[BGP Network](#)

[Description](#)

[Property Description](#)

[Notes](#)

[Example](#)

[BGP Peers](#)

[Description](#)

[Property Description](#)

[Example](#)

[Troubleshooting](#)

[Description](#)

General Information

Summary

The Border Gateway Protocol (BGP) allows setting up an interdomain dynamic routing system that automatically generates the routing table for routing between autonomous systems (AS).

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

The MikroTik RouterOS implementation of the BGP has filtering (using prefix lists) feature

Specifications

Packages required: ***routing***

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

- [Package Management](#)
- [IP Addresses and ARP](#)
- [Routes, Equal Cost Multipath Routing, Policy Routing](#)
- [Prefix Lists](#)

Description

The Border Gateway Protocol (BGP) is an Exterior Gateway Protocol (EGP). It allows setting up an interdomain routing system that automatically guarantees the loop-free exchange of routing information between autonomous systems (AS). It is widely used in companies assigned with a definite IP address ranges and connected to a number of ISPs simultaneously so that if one of the links is down, the IP address ranges are still reachable via another ISP.

The MikroTik RouterOS implementation of the BGP supports filtering with prefix lists, that is used for filtering received and sent routing information.

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>

BGP Setup

Home menu level: */routing bgp*

Property Description

enabled (*yes | no*; default: **no**) - enable or disable BGP

as (*integer*; default: **1**) - autonomous system number

router-id (*IP address*; default: **0.0.0.0**) - the Router identification in form of an IP address

redistribute-connected (*yes | 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-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

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

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

state (*read-only: disabled | running | terminating*) - status of the BGP

- **disabled** - not working, has been disabled
- **running** - working
- **terminating** - shutting down, flushing all route information

Notes

Usually, you want to redistribute connected and static routes, if any. Therefore change the settings for these arguments and proceed to the BGP networks.

Example

To enable BGP protocol specifying that router **192.168.0.206**, that belongs to the **65002** AS, should redistribute the connected routes

```
[admin@MikroTik] routing bgp>
[admin@MikroTik] routing bgp> print
      enabled: yes
      as: 65002
      router-id: 192.168.0.206
      redistribute-static: no
      redistribute-connected: yes
      redistribute-rip: no
      redistribute-ospf: no
      state: running
[admin@MikroTik] routing bgp>
```

BGP Network

Home menu level: */routing bgp network*

Description

BGP Networks is a list of the networks to be advertised.

Property Description

network (*IP address/mask*; default: **0.0.0.0/0**) - network to advertise

Notes

You can add to the list as many networks as required.

The router is not checking whether the network is in the routing table, it always advertises all the routes that are specified here.

Note the difference with OSPF, that use network list for different purpose - to determine where to send updates.

Example

To advertise the network **159.148.150.192/27**:

```
[admin@modux] routing bgp network> add network=159.148.150.192/27
[admin@modux] routing bgp network> print
# NETWORK
0 159.148.150.192/27
[admin@modux] routing bgp network>
```

BGP 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

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

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)

route-reflect (*yes | no*; default: **no**) - defines whether to redistribute further the routes learned from router of the same AS or not. If enabled, can significantly reduce traffic between routers in the same AS

prefix-list-in (*name*; default: **""**) - name of the filtering prefix list for receiving routes

prefix-list-out (*name*; default: **""**) - name of the filtering prefix list for advertising routes

state (*read-only: connected | not-connected*) - the status of the BGP connection to the peer

routes-received - the number of received routes from this peer

Example

To enable routing information exchange with the neighbour (non-multihop) **192.168.0.254** that belongs to **65002** AS:

```
[admin@MikroTik] routing bgp peer> add remote-address=192.168.0.254 remote-as=65002
[admin@MikroTik] routing bgp peer> print
# REMOTE-ADDRESS  REMOTE-AS  MULTIHOPE  ROUTE-REFLECT  PREFIX-LIS...  PREFIX-LI...
0 192.168.0.254  65002      no             no             none          none
[admin@MikroTik] routing bgp> peer print status
# REMOTE-ADDRESS  REMOTE-AS  STATE          ROUTES-RECEIVED
0 192.168.0.254  65002      connected      1
[admin@MikroTik] routing bgp>
```

Troubleshooting

Description

- *The BGP does not learn routes from its peer*

Try to see if the peer is directly attached, or you should use the **multihop** flag when defining the peer and static routing to get the connection between the peers.

- *I can ping from one peer to the other one, but no routing exchange takes place*
Check the status of the peer using **/routing bgp peer print detail** command. See if you do not have firewall that blocks TCP port 179.