

BGP Routing Filters

Document revision 1.4 (Fri Sep 23 08:43:17 GMT 2005)

This document applies to V2.9

Table of Contents

[Table of Contents](#)

[General Information](#)

[Summary](#)

[Specifications](#)

[Related Documents](#)

[Description](#)

[Additional Documents](#)

[Filter Rules](#)

[Property Description](#)

General Information

Summary

Border Gateway Protocol (BGP) Routing filters allow to alter attributes of the route for the NLRI prefixes or completely exclude particular NLRI prefixes with routes from the BGP routing update message.

Specifications

Packages required: *routing*

License required: *level3*

Home menu level: */routing filter*

Standards and Technologies: [RFC1771](#)

Hardware usage: *Not significant*

Related Documents

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

Description

BGP filtering refers to the ability of BGP peer to apply administrative policies to incoming and outgoing routing update messages. These policies are implemented as rules organized in chains. The following manual uses terms 'chain' and 'filter' interchangeably. Each rule consists of two parts, one of them specifies to which prefixes the rule applies to and the other tells the router what to do with these prefixes. A rule with no arguments applies to all prefixes and implies **accept** action.

The routing filters may be applied to incoming and outgoing routing update messages for a specific BGP peer and to outgoing BGP update messages for a particular BGP instance. Note, that in case both BGP instance and BGP peer outgoing filters are applied, BGP instance filters take precedence.

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>

Filter Rules

Property Description

action (*accept* | *discard* | *jump* | *none* | *reject* | *return*; default: **none**) - action to perform on route or route attributes for the NLRI prefixes that match the rule

- **accept** - accept the routing information for the matching NLRI prefix
- **discard** - completely exclude matching prefix from the BGP processing. The route will be deleted from the incoming BGP routing update message, thus reducing memory usage on the router. For outgoing BGP update messages the discard action is equal to reject
- **jump** - pass control to another filter list that should be specified as jump-target parameter
- **none** - do not perform any action and pass execution to the next rule in chain. The none action is not displayed by print command
- **reject** - reject the routing information for matching prefix. The prefix from incoming BGP routing update message is shown with R (rejected) flag in the /ip route print command output. The prefix is suppressed from outgoing routing update message
- **return** - return to the previous chain from which a jump to the current chain took place

as-path (*text*) - unanchored pattern to be searched inside AS_PATH attribute of the route. Optional ^ sign preceding parameter value restricts match to the beginning of AS_PATH attribute, while \$ sign, which follows as-path value, restricts the match to the end of AS_PATH

as-path-length (*integer* | *integer*) - length of the AS_PATH attribute, representing the number of ASs that have been traversed. Note that multiple AS_SETs are combined together and counted as 1 AS

atomic-aggregate (*absent* | *present*) - match for the ATOMIC_AGGREGATE BGP attribute

chain (*text*) - chain name to place this rule in. If a chain with the specified name does not exist it will be automatically created

distance (*integer* | *integer*; default: **no**) - protocol-independent administrative distance used to compare routes obtained from different sources

jump-target (*name*) - name of the target chain to jump to, if the action=jump is used

local-pref (*integer* | *integer*) - match for the LOCAL_PREF BGP attribute

match-chain (*name*) - the name of the chain which is used to evaluate the route. If the chain accepts the route, match-chain property produces a true match

med (*integer | integer*) - match for the MULTI_EXIT_DISC BGP attribute

origin (*igp | egp | incomplete*) - match for the ORIGIN BGP attribute

prefix (*IP address | netmask | IP address | IP address*) - match for the NLRI prefix

prefix-length (*integer | integer*) - match for the NLRI prefix length

prefsrc (*IP address | netmask | IP address | IP address*) - match for the preferred source IP address of the route

route-comment (*text*) - match for the route comment

routing-mark (*text*) - match for the routing mark. A routing mark identifies certain routes for successive processing

scope (*integer: 0..255 | integer: 0..255*) - scope and target-scope are used to recursively lookup next hop address for the route. Routes that are used to lookup the next hop address for a given route should have scope value equal or less then the target-scope value of this route

set-check-gateway (*ping | arp*) - specifies that the router should check whether the gateway for the particular route is reachable by using either ping or arp request prior to sending anything using this route

set-disabled - disables the route. Disabled routes are not considered by BGP best path selection algorithm

set-distance (*integer: 0..255*) - sets administrative distance for a route. The distance is protocol-independent and is used to compare routes obtained from different sources

set-localpref (*integer: 0..4294967295*) - specifies LOCAL_PREF BGP attribute value for the route

set-med (*integer: 0..4294967295*) - sets MULTI_EXIT_DISC BGP attribute

set-nexthop (*IP address*) - sets next hop IP address for the route

set-prefsrc (*IP address*) - sets preferred source address for the route

set-prepend (*integer: 0..16*) - specifies how many times the router should prepend its AS number to the AS_PATH BGP attribute value for this route

set-route-comment (*text*) - specifies comment for the route

set-routing-mark (*text*) - sets routing mark for the route

set-scope (*integer: 0..255*) - sets scope for the route. Scope and target-scope are used to recursively lookup next hop address for the route. Routes that are used to lookup the next hop address for a given route should have scope value equal or less then the target-scope value of this route

set-target-scope (*integer: 0..255*) - sets target scope for the route. Scope and target-scope are used to recursively lookup next hop address for the route. Routes that are used to lookup the next hop address for a given route should have scope value equal or less then the target-scope value of this route

set-weight (*integer: -2147483648..2147483647*) - specifies weight for the route. Route weight is used by BGP best path selection algorithm to select the best route towards destination

target-scope (*integer: 0..255 | integer: 0..255*) - scope and target-scope are used to recursively lookup next hop address for the route. Routes that are used to lookup the next hop address for a given route should have scope value equal or less then the target-scope value of this route

type (*absent | present*) - match for the ATOMIC_AGGREGATE BGP attribute

unset (*multiple choice: prefsrc | routing-mark | check-gateway | disabled*) - unsets specified parameters of the route

weight (*integer: -2147483648..2147483647*) - match for the weight of the route