

Manual:IP/Route

Applies to RouterOS: v3, v4, v5+

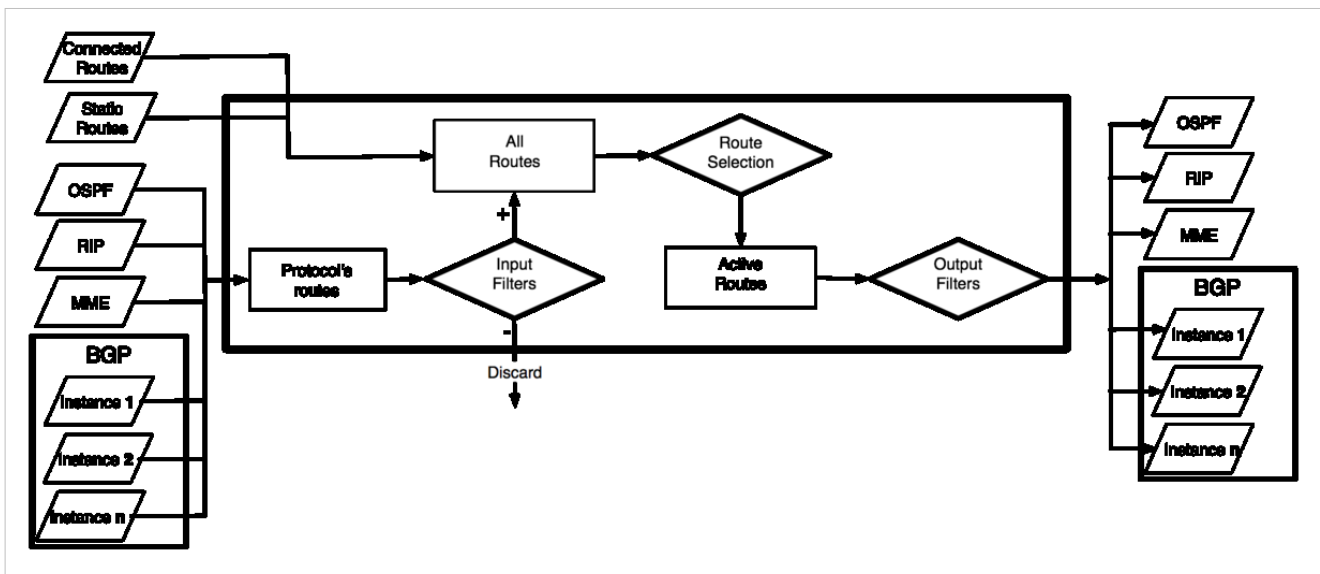


Overview

Router keeps routing information in several separate spaces:

- FIB (Forwarding Information Base), that is used to make packet forwarding decisions. It contains a copy of the necessary routing information.
- Each routing protocol (except BGP) has its own internal tables. This is where per-protocol routing decisions are made. BGP does not have internal routing tables and stores complete routing information from all peers in the RIB.
- RIB contains routes grouped in separate routing tables based on their value of **routing-mark**. All routes without **routing-mark** are kept in the *main* routing table. These tables are used for best route selection. The *main* table is also used for nexthop lookup.

Routing Information Base



RIB (Routing Information Base) contains complete routing information, including static routes and policy routing rules configured by the user, routing information learned from routing protocols, information about connected networks. RIB is used to filter routing information, calculate best route for each destination prefix, build and update Forwarding Information Base and to distribute routes between different routing protocols.

By default forwarding decision is based only on the value of destination address. Each route has **dst-address** property, that specifies all destination addresses this route can be used for. If there are several routes that apply to a particular IP address, the most specific one (with largest netmask) is used. This operation (finding the most specific route that matches given address) is called *routing table lookup*.

If routing table contains several routes with the same **dst-address**, only one of them can be used to forward packets. This route is installed into FIB and marked as *active*.

When forwarding decision uses additional information, such as a source address of the packet, it is called *policy routing*. Policy routing is implemented as a list of policy routing rules, that select different routing table based on destination address, source address, source interface, and routing mark (can be changed by firewall mangle rules) of the packet.

All routes by default are kept in the *main* routing table. Routes can be assigned to specific routing table by setting their **routing-mark** property to the name of another routing table. Routing tables are referenced by their name, and are created automatically when they are referenced in the configuration.

Each routing table can have only one active route for each value of **dst-address** IP prefix.

There are different groups of routes, based on their origin and properties.

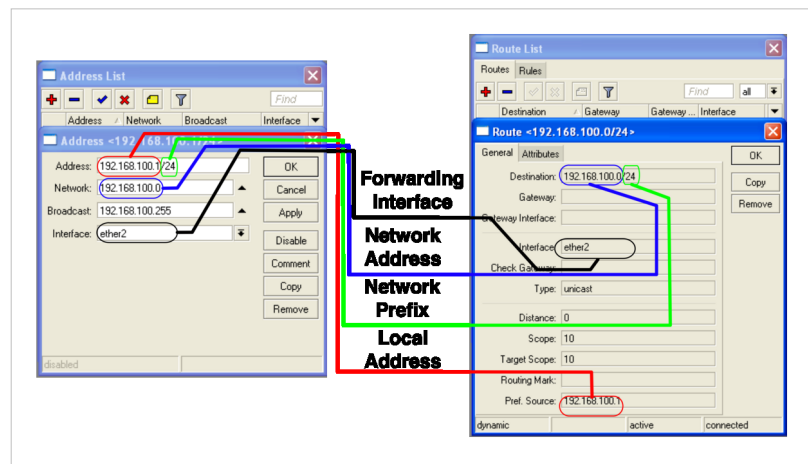
Default route

Route with **dst-address** *0.0.0.0/0* applies to every destination address. Such route is called the *default route*. If routing table contains an active default route, then routing table lookup in this table will never fail.

Connected routes

Connected routes are created automatically for each IP network that has at least one enabled interface attached to it (as specified in the `/ip address` configuration). RIB tracks status of connected routes, but does not modify them. For each connected route there is one *ip address* item such that:

- address part of **dst-address** of connected route is equal to **network** of *ip address* item.
- netmask part of **dst-address** of connected route is equal to netmask part of **address** of *ip address* item.
- **pref-src** of connected route is equal to address part of **address** of *ip address* item.
- **interface** of connected route is equal to **actual-interface** of *ip address* item (same as **interface**, except for bridge interface ports).



Multipath (ECMP) routes

Because results of the forwarding decision are cached, packets with the same **source address**, **destination address**, **source interface**, **routing mark** and **ToS** are sent to the same gateway. This means that one connection will use only one link in each direction, so ECMP routes can be used to implement per-connection load balancing. See interface bonding if you need to achieve per-packet load balancing.

To implement some setups, such as load balancing, it might be necessary to use more than one path to given destination. However, it is not possible to have more than one active route to destination in a single routing table.

ECMP (Equal cost multi-path) routes have multiple **gateway** nexthop values. All reachable nexthops are copied to FIB and used in forwarding packets.

OSPF protocol can create ECMP routes. Such routes can also be created manually.

Routes with interface as a gateway

Value of **gateway** can be specified as an interface name instead of the nexthop IP address. Such route has following special properties:

- Unlike connected routes, routes with interface nexthops are not used for nexthop lookup.
- It is possible to assign several interfaces as a value of **gateway**, and create ECMP route. It is not possible to have connected route with multiple **gateway** values.

Route selection

Each routing table can have one active route for each destination prefix. This route is installed into FIB. Active route is selected from all candidate routes with the same **dst-address** and **routing-mark**, that meet the criteria for becoming an active route. There can be multiple such routes from different routing protocols and from static configuration. Candidate route with the lowest **distance** becomes an active route. If there is more than one candidate route with the same **distance**, selection of active route is arbitrary (except for BGP routes).

BGP has the most complicated selection process (described in separate article). Notice that this protocol-internal selection is done only after BGP routes are installed in the main routing table; this means there can be one candidate route from each BGP peer. Also note that BGP routes from different BGP instances are compared by their distance, just like other routes.

Criteria for selecting candidate routes

To participate in route selection process, route has to meet following criteria:

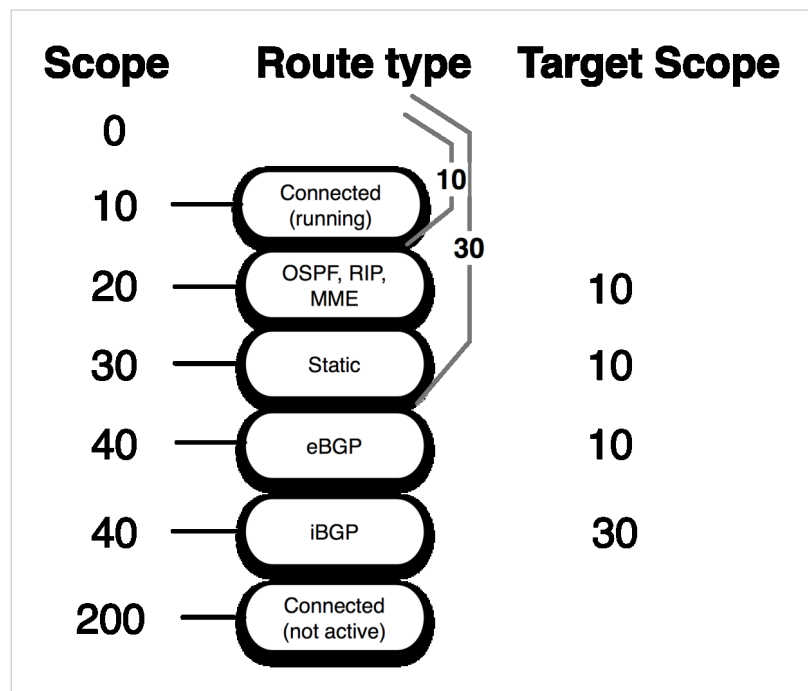
- route is not disabled.
- **distance** is not 255. Routes that are rejected by route filter have **distance** value of 255.
- **pref-src** is either not set or is a valid local address of the router.
- **routing-mark** is either not set or is referred by firewall or policy routing rules.
- If type of route is *unicast* and it is not a connected route, it must have at least one reachable nexthop.

Nexthop lookup

Nexthop lookup is a part of the route selection process.

Routes that are installed in the FIB need to have interface associated with each gateway address. Gateway address (nexthop) has to be directly reachable via this interface. Interface that should be used to send out packets to each gateway address is found by doing nexthop lookup.

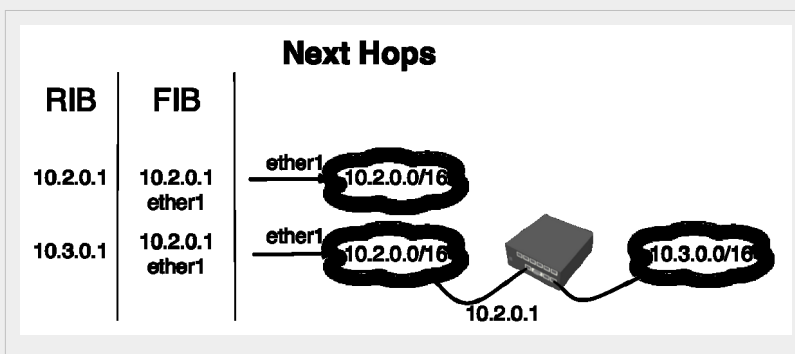
Some routes (e.g. iBGP) may have gateway address that is several hops away from this router. To install such routes in the FIB, it is necessary to find the address of the directly reachable gateway (an immediate nexthop), that



should be used to reach the gateway address of this route. Immediate nextop addresses are also found by doing nextop lookup.

Nextop lookup is done only in the *main* routing table, even for routes with different value of **routing-mark**. It is necessary to restrict set of routes that can be used to look up immediate nextops. Nextop values of RIP or OSPF routes, for example, are supposed to be directly reachable and should be looked up only using connected routes. This is achieved using **scope** and **target-scope** properties.

- Routes with interface name as the value of gateway are not used for nextop lookup. If route has both interface nextops and active IP address nextops, then interface nextops are ignored.
- Routes with **scope** greater than the maximum accepted value are not used for nextop lookup. Each route specifies maximum accepted **scope** value for it's nextops in the **target-scope** property. Default value of this property allows nextop lookup only through connected routes, with the exception of iBGP routes that have larger default value and can lookup nextop also through IGP and static routes.



- nextop 10.2.0.1 is resolved through a connected route, it's status is *reachable*.
- nextop 10.3.0.1 is resolved recursively through a 10.3.0.0/16 route, it's status is *recursive*, and it uses 10.2.0.1 as the immediate nextop value that is installed in the FIB.

Interface and immediate nextop are selected based on the result of nextop lookup:

- If most specific active route that nextop lookup finds is connected route, then interface of this connected route is used as the nextop interface, and this gateway is marked as *reachable*. Since gateway is directly reachable through this interface (that's exactly what connected route means), the gateway address is used as the immediate nextop address.
- If most specific active route that nextop lookup finds has nextop that is already resolved, immediate nextop address and interface is copied from that nextop and this gateway is marked as *recursive*.
- If most specific active route that nextop lookup finds is ECMP route, then it uses first gateway of that route that is not *unreachable*.
- If nextop lookup does not find any route, then this gateway is marked as *unreachable*.

Forwarding Information Base

FIB (Forwarding Information Base) contains copy of information that is necessary for packet forwarding:

- all active routes
- policy routing rules

By default (when no **routing-mark** values are used) all active routes are in the *main* table, and there is only one hidden implicit rule ("catch all" rule) that uses the *main* table for all destination lookups.

Routing table lookup

FIB uses following information from packet to determine it's destination:

- source address
- destination address
- source interface
- routing mark
- ToS (not used by RouterOS in policy routing rules, but it is a part of routing cache lookup key)

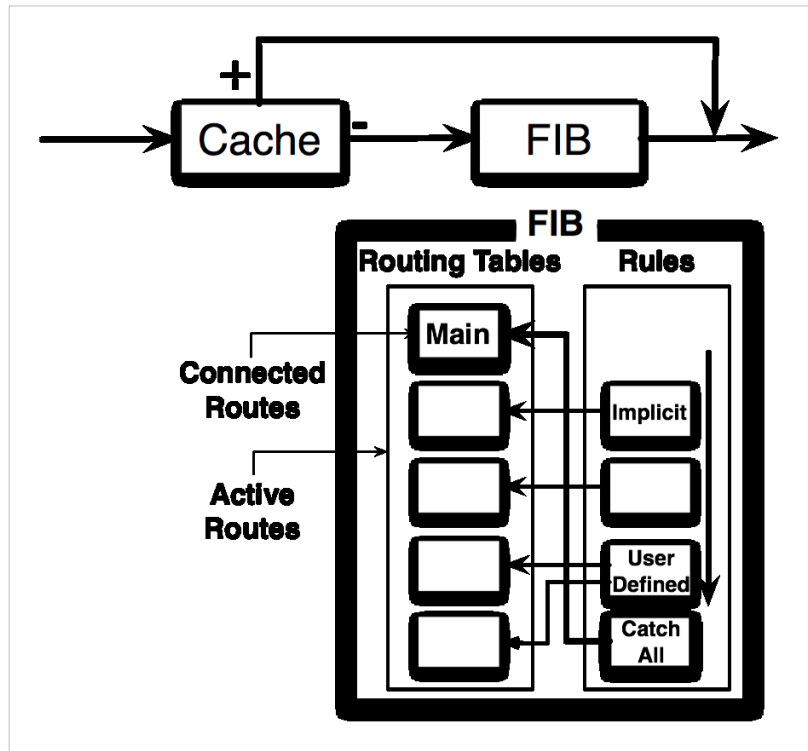
Possible routing decisions are:

- receive packet locally
- discard packet (either silently or by sending ICMP message to the sender of the packet)
- send packet to specific IP address on specific interface

Results of routing decision are remembered in the routing cache. This is done to improve forwarding performance. When another packet with the same **source address, destination address, source interface, routing mark** and **ToS** is routed, cached results are used. This also allows to implement per-connection load balancing using ECMP routes, because values used to lookup entry in the routing cache are the same for all packets that belong to the same connection and go in the same direction.

If there is no routing cache entry for this packet, it is created by running routing decision:

- check that packet has to be locally delivered (destination address is address of the router)
- process implicit policy routing rules
- process policy routing rules added by user
- process implicit catch-all rule that looks up destination in the *main* routing table
- return result is "network unreachable"



Result of routing decision can be:

- IP address of nexthop + interface
- point-to-point interface
- local delivery
- discard
- ICMP prohibited
- ICMP host unreachable
- ICMP network unreachable

Rules that do not match current packet are ignored. If rule has action *drop* or *unreachable*, then it is returned as a result of the routing decision process. If action is *lookup* then destination address of the packet is looked up in routing table that is specified in the rule. If lookup fails (there is no route that matches destination address of packet), then FIB proceeds to the next rule. Otherwise:

- if type of the route is *blackhole*, *prohibit* or *unreachable*, then return this action as the routing decision result;
- if this is a connected route, or route with an interface as the **gateway** value, then return this interface and the destination address of the packet as the routing decision result;
- if this route has IP address as the value of **gateway**, then return this address and associated interface as the routing decision result;
- if this route has multiple values of nexthop, then pick one of them in round robin fashion.

Result of this routing decision is stored in new routing cache entry.

Properties

Route flags

Property(Flag)	Description
disabled (X)	Configuration item is disabled. It does not have any effect on other routes and is not used by forwarding or routing protocols in any way.
active (A)	Route is used for packet forwarding. See route selection.
dynamic (D)	Configuration item created by software, not by management interface. It is not exported, and cannot be directly modified.
connect (C)	connected route.
static (S)	static route.
rip (r)	RIP route.
bgp (b)	BGP route.
ospf (o)	OSPF route.
mme (m)	MME route.
blackhole (B)	Silently discard packet forwarded by this route.
unreachable (U)	Discard packet forwarded by this route. Notify sender with ICMP host unreachable (type 3 code 1) message.
prohibit (P)	Discard packet forwarded by this route. Notify sender with ICMP communication administratively prohibited (type 3 code 13) message.

General properties

Property	Description
check-gateway (<i>arp</i> <i>ping</i> ; Default: <code>""</code>)	Periodically (every 10 seconds) check gateway by sending either ICMP echo request (<i>ping</i>) or ARP request (<i>arp</i>). If no response from gateway is received for 10 seconds, request times out. After two timeouts gateway is considered unreachable. After receiving reply from gateway it is considered reachable and timeout counter is reset.
comment (<i>string</i> ; Default: <code>""</code>)	Description of particular route
distance (<i>integer</i> [1..255]; Default: <code>"1"</code>)	Value used in route selection. Routes with smaller distance value are given preference. If value of this property is not set, then the default depends on route protocol: <ul style="list-style-type: none"> • connected routes: 0 • static routes: 1 • eBGP: 20 • OSPF: 110 • RIP: 120 • MME: 130 • iBGP: 200
dst-address (<i>IP prefix</i> ; Default: <code>"0.0.0.0/0"</code>)	IP prefix of route, specifies destination addresses that this route can be used for. Netmask part of this property specifies how many of the most significant bits in packet destination address must match this value. If there are several active routes that match destination address of packet, then the most specific one (with largest netmask value) is used.
gateway (<i>IP IP%interface</i> <i>IP@table</i> , <i>IP</i> <i>string</i> , [...]; Default: <code>""</code>)	Array of IP addresses or interface names. Specifies which host or interface packets should be sent to. Connected routes and routes with <i>blackhole</i> , <i>unreachable</i> or <i>prohibit</i> type do not have this property. Usually value of this property is a single IP address of a gateway that can be directly reached through one of router's interfaces (but see nexthop lookup). ECMP routes have more than one gateway value. Value can be repeated several times.
pref-src (<i>IP</i> ; Default: <code>""</code>)	Which of the local IP addresses to use for locally originated packets that are sent via this route. Value of this property has no effect on forwarded packets. If value of this property is set to IP address that is not local address of this router then the route will be inactive. If pref-src value is not set, then for locally originated packets that are sent using this route router will choose one of local addresses attached to the output interface that match destination prefix of the route (an example).
route-tag (<i>integer</i> ; Default: <code>""</code>)	Value of route tag attribute for RIP or OSPF. For RIP only values 0..4294967295 are valid.
routing-mark (<i>string</i> ; Default: <code>""</code>)	Name of routing table that contains this route. Not set by default which is the same as <i>main</i> . Packets that are marked by firewall with this value of routing-mark will be routed using routes from this table, unless overridden by policy routing rules. Not more than 250 routing marks are possible per router.
scope (<i>integer</i> [0..255]; Default: <code>"30"</code>)	Used in nexthop resolution. Route can resolve nexthop only through routes that have scope less than or equal to the target-scope of this route. Default value depends on route protocol: <ul style="list-style-type: none"> • connected routes: 10 (if interface is running) • OSPF, RIP, MME routes: 20 • static routes: 30 • BGP routes: 40 • connected routes: 200 (if interface is not running)
target-scope (<i>integer</i> [0..255]; Default: <code>"10"</code>)	Used in nexthop resolution. This is the maximum value of scope for a route through which a nexthop of this route can be resolved. See nexthop lookup. For iBGP value is set to 30 by default.
type (<i>unicast</i> <i>blackhole</i> <i>prohibit</i> <i>unreachabl</i> ; Default: <code>unicast</code>)	Routes that do not specify nexthop for packets, but instead perform some other action on packets have type different from the usual <i>unicast</i> . <i>blackhole</i> route silently discards packets, while <i>unreachable</i> and <i>prohibit</i> routes send ICMP Destination Unreachable message (code 1 and 13 respectively) to the source address of the packet.
vrf-interface (<i>string</i> ; Default: <code>"10"</code>)	VRF interface name

Other Read-only properties

Property	Description
gateway-status (array)	Array of gateways, gateway states and which interface is used for forwarding. Syntax "IP state interface", for example "10.5.101.1 reachable bypass-bridge". State can be <i>unreachable</i> , <i>reachable</i> or <i>recursive</i> . See nexthop lookup for details.
ospf-metric (integer)	Used OSPF metric for particular route
ospf-type (string)	

BGP Route Properties

These properties contain information that is used by BGP routing protocol. However, values of these properties can be set for any type of route, including static and connected. It can be done either manually (for static routes) or using route filters.

Property	Description
bgp-as-path (string; Default: "")	Value of BGP AS_PATH attribute. Comma separated list of AS numbers with confederation AS numbers enclosed in () and AS_SETs enclosed in {}. Used to check for AS loops and in BGP route selection algorithm: routes with shorter AS_PATH are preferred (but read how AS_PATH length is calculated).
bgp-atomic-aggregate (yes no; Default:)	Value of BGP ATOMIC_AGGREGATE attribute.
bgp-communities (array of (integer:integer internet no-advertise no-export local-as; Default:)	Value of BGP communities list. This attribute can be used to group or filter routes. Named values have special meanings: <ul style="list-style-type: none"> <i>internet</i> - advertise this route to the Internet community (i.e. all routers) <i>no-advertise</i> - do not advertise this route to any peers <i>no-export</i> - do not advertise this route to EBGp peers <i>local-as</i> - same as <i>no-export</i>, except that route is also advertised to EBGp peers inside local confederation
bgp-local-pref (integer; Default:)	Value of BGP LOCAL_PREF attribute. Used in BGP route selection algorithm: routes with greater LOCAL_PREF value are preferred. If value is not set then it is interpreted as 100.
bgp-med (integer; Default:)	Value of BGP MULTI_EXIT_DISC BGP attribute. Used in BGP route selection algorithm: routes with lower MULTI_EXIT_DISC value are preferred.. If value is not set then it is interpreted as 0.
bgp-origin (igp egp incomplete; Default:)	Value of BGP ORIGIN attribute. Used in BGP route selection algorithm: <i>igp</i> routes are preferred over <i>egp</i> and <i>egp</i> over <i>incomplete</i> .
bgp-prepend (integer [0..16]; Default:)	How many times to prepend router's own AS number to AS_PATH attribute when announcing route via BGP. Affects only routes sent to eBGp peers (for iBGp value 0 is always used).

Read-only

Property	Description
bgp-ext-communities (<i>string</i>)	Value of BGP extended communities attribute
bgp-weight (<i>integer</i>)	Additional value used by BGP best path selection algorithm. Routes with higher weight are preferred. It can be set by incoming routing filters and is useful only for BGP routes. If value is not set then it is interpreted as 0.
received-from (<i>string</i>)	Name of the BGP peer from which route is received.

Article Sources and Contributors

Manual:IP/Route *Source:* <http://wiki.mikrotik.com/index.php?oldid=20436> *Contributors:* Eep, Janisk, Marishb

Image Sources, Licenses and Contributors

Image:Version.png *Source:* <http://wiki.mikrotik.com/index.php?title=File:Version.png> *License:* unknown *Contributors:* Normis

Image:rib.png *Source:* <http://wiki.mikrotik.com/index.php?title=File:Rib.png> *License:* unknown *Contributors:* Eep

Image:conn_route_and_address.png *Source:* http://wiki.mikrotik.com/index.php?title=File:Conn_route_and_address.png *License:* unknown *Contributors:* Eep

Image:scope_and_target_scope.png *Source:* http://wiki.mikrotik.com/index.php?title=File:Scope_and_target_scope.png *License:* unknown *Contributors:* Eep

Image:nh-lookup.png *Source:* <http://wiki.mikrotik.com/index.php?title=File:Nh-lookup.png> *License:* unknown *Contributors:* Eep

Image:fib.png *Source:* <http://wiki.mikrotik.com/index.php?title=File:Fib.png> *License:* unknown *Contributors:* Eep