


Advanced Routing

<https://campus.barracuda.com/doc/17265/>

Routing tables are used to store the best path to a remote network. Routing tables are processed from top to bottom. If the source of the outgoing packets matches, the routes in the route table are evaluated and the packet forwarded to the correct interface, next hop gateway, or VPN tunnel. Routes are evaluated first by the destination route metric (preference) of an IP packet and then by the scope (network size) to determine which routes matches. Two routes of the same scope (e.g., /24) and metric cannot be created. The management network always uses a preference of 0.

- If two routes with different preferences exist, the route with the lower preference is chosen. E.g., 10.0.10.0/25 (preference 10) is preferred over 10.0.10.0/25 (preference 100)
- If two routes with the same preference exist to a destination, the route with the smaller subnet mask is used. E.g., 10.0.10.0/24 is preferred over 10.0.0.0/16
- VPN routes are placed in a premain, source-based route table by default. If **single routing table** is enabled in the VPN settings, VPN routes are inserted into the main routing table with a preference of 10. For more information, see [Authentication, Encryption, Transport, IP Version and VPN Routing](#).



Directly Attached Network Routes (Direct Routing)

Define how to reach networks that are directly connected to an interface (virtual or physical) of the firewall. To define a directly attached network route, you must enter the following:

- **Target network in CIDR Format** - E.g., 10.0.8.0/24 or 2001:db8::6299::/64
- **Interface** - The network interface on the firewall the network is attached to. E.g., eth2 or port2

After you have introduced the directly attached route and activated the network configuration, the route is in a pending state. Pending routes are marked with the ✘ icon in **CONTROL > Network** and are not active. When a service IP address from this network is introduced, the route becomes active and the ■ icon is displayed for the route.

In the diagram above, you must create a direct route for the ISP issued 62.99.0.0/24. To reach the Internet, a gateway route must be created. If you enter the optional **gateway IP** address when

creating the directly attached route, the default gateway route is created automatically.

You do not need to create a directly attached route for the network the management IP address is in. This route is created automatically when the management IP address is configured.

For setup instructions, see [How to Configure Directly Attached Routes](#) and [How to Configure Directly Attached IPv6 Routes](#).

Using Additional Local IPs to Directly Attach Networks

Additional local IP addresses are a combination of a box level IP address and a directly attached route. The route becomes active when the network configuration is activated; the route is never pending because of the additional IP address. Add the **Additional IP addresses (CONFIGURATION > Configuration Tree > Box > Network)**. IP addresses assigned on box level should not be used on the service layer of a high availability cluster. When using the IP address on box level, the route will remain active even if the assigned services are running on the other firewall in the HA cluster.

Gateway Routes (Next Hop Routing)

To reach networks that cannot be directly accessed, you must define gateway routes. A common gateway route is the default route (0.0.0.0/0), which will forward all packets not belonging to one of the directly attached networks to the remote gateway provided by the ISP. Before adding a gateway route, a direct route must be configured. Otherwise, you cannot contact the next hop IP address. If you are using multiple gateway routes for the same target network, you must give them different route metrics. Gateway routes automatically monitor the gateway IP address. When the gateway is no longer considered healthy, 6535 is added to the metric of the route. Routes with a metric above 65535 are considered to be down. To define a gateway route, you must enter the following:

- **Target network** - Target network in CIDR format. E.g., 0.0.0.0/0 or ::0/0 for the default route
- **Next hop address** - IP address of the gateway device the traffic is sent to. E.g., 62.99.0.254 or 2001:db8:10::ffff

After adding the gateway route, you must initiate a failsafe network activation for the route to become active (in **CONTROL > Network**).

For setup instructions, see [How to Configure Gateway Routes](#) and [How to Configure IPv6 Gateway Routes](#).

Multipath Routing

The Barracuda CloudGen Firewall supports standard Linux multipath routing and Firewall-assisted multipath routing. Standard Linux multipath routing balances do not offer dead next hop detection or session packet balancing. Simple redundancy by next hop detection can be provided by adding multiple routing entries with different route preference numbers. Firewall-assisted multipath routing supports per-packet balancing between next hops and dead next peer detection and is configured in the Forwarding Firewall service.

For setup instructions, see:

- [How to Configure Multipath Routing](#)
- [How to Configure Linux Standard Multipath Routing](#)

Source-based Routes (Policy-based Routing)

Source-based routes can be created manually for static interfaces, or in the configuration for the dynamic interface (DHCP, xDSL, ...). For each route table, you define which source network and then create routes in the source-based route table. The routes can be one of three types:

- **unicast**
- **multipath**
- **throw** - A throw route causes the table lookup to be terminated.

Source-based routes are automatically created for dynamic interfaces. Configure a DHCP, xDSL, or ISDN link. You can disable source-based routing per advanced configuration.

For setup instructions, see [How to Configure Source-Based Routes](#).

Figures

1. routing_over.png
2. route_pending.png
3. route_active.png
4. route_active.png

© Barracuda Networks Inc., 2026 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.