

Dynamic Routing Protocols (OSPF/RIP/BGP)

<https://campus.barracuda.com/doc/96026040/>

The Barracuda CloudGen Firewall supports three dynamic routing protocols: Open Shortest Path First (OSPF), Routing Information Protocol (RIP Version 1 and RIP Version 2), and Border Gateway Protocol (BGP). OSPF and RIP are Interior Gateway Protocols (IGP) and distribute routing information within an autonomous system, whereas BGP is a Exterior Gateway Protocol. The routes learned via the dynamic routing protocols are applied to the kernel routing table. Set the route metric instead of the administrative distance to prioritize one route over the other.

OSPF

The CloudGen Firewall supports OSPFv2 and OSPFv3 versions of the OSPF protocol. OSPF is a link state protocol and uses the Dijkstra algorithm to calculate the shortest path tree. A router's interface is the "link". The "state" of this interface is summed up by its IP address, subnet mask, interface type, neighbor state, etc. Every router keeps track of all connected interfaces and states and sends this information with multicasts to its neighbors. These packets are known as LSAs (Link State Advertisements). The router builds its link state database with the information provided by the LSAs. Every time a network change occurs, LSAs containing the new information are sent, thus triggering every router to update its database. After having received all LSAs, the router calculates the loop-free topology. LSAs cannot be filtered within an area because all routers in an area must have the same link state database. If some information is missing, routing loops can occur.

OSPF is a hierarchical IGP and uses Areas to achieve this. The top-level Area is known as Backbone Area, and the number of this Area must always be 0 or 0.0.0.0. All other Areas must be physically connected to this Backbone Area. One very important aspect of OSPF is that Areas must not be split. (If this cannot be avoided, a virtual link must be used to expand Area 0 over any other area.) Routers within an area are known as Area Routers. Routers connected to two or more areas are known as Area Border Routers (ABR) and routers connected to other autonomous systems are called Autonomous System Boundary Routers (ASBR). Routing information can be summarized on ABRs and ASBRs. It is not possible to summarize routing information within an area.

The metric used by OSPF is cost. Every link has an associated cost value, derived from the link bandwidth. The metric to a destination is calculated by adding up all costs. If there are more possible paths to a destination, the route with the lowest cost is chosen as the best route. To advertise LSAs, the router must live in OSPF neighborhood with other routers. When this neighborhood is fully established, the interfaces begin sending the updates (LSAs). To build an adjacency, hello packets are continuously exchanged between neighboring routers. This also keeps track of the existence of the connected OSPF neighbors. To lower the number of updates exchanged on a broadcast medium (for example, Ethernet), LSAs are only sent to a so-called Designated Router (DR). This interface advertises the information to all other routers on the shared medium. Without a DR, an any-to-any neighborhood between all OSPF routers on this segment would be needed. For backup reasons, a

Backup DR (BDR) is elected. Each other router establishes neighborship only with the DR and BDR.

Areas can be configured as stub areas, where external routes are not advertised by ABRs to the Area Routers. Instead, a default route is injected to the area. Area 0 cannot be stub.

OSPF is very CPU and memory intensive. Therefore, be careful when enabling OSPF on low-end interfaces in a large network.

For more information, see:

- [How to Install and Configure the OSPF/RIP/BGP Service](#)
- [How to Configure OSPF Routers and Areas](#)
- [How to Configure the Filter Setup for OSPF and RIP](#)

RIP

The CloudGen Firewall supports RIPv1, RIPv2, and RIPv6 versions of the RIP protocol. RIP is a distance-vector protocol. The expression "distance-vector" can be defined as follows: The vector is the direction to the destination (next hop); the distance is treated as a metric type. Example: Destination A is a distance of 3 hops away, and the direction is via router AA. RIP uses hop count as metric. A maximum of 15 hops are possible; metric 16 means that a network is unreachable. All RIP routers periodically send routing updates. Every update includes the whole routing table. The following techniques have been introduced to prevent routing loops:

- **Split Horizon** – When sending updates out a particular interface, the routes learned from this interface are not included in the update.
- **Split Horizon with Poison reverse** – This method is an extension to Split Horizon. The router includes learned routes in the update but marks these routes as unreachable.
- **Counting to infinity** – To recognize unreachable networks on link failures. Infinity in RIP is defined as 16 hops. Every time a routing update passes a router, the hop count is increased by 1. When the counter reaches 16, the network is considered unreachable.

RIPv1 is classful, which means that subnet information cannot be distributed. RIPv2, on the other hand, is classless. This means the subnet mask is included in the routing update. The maximum route metric for RIP routes is 255. This means it is not possible to use RIP routes as fallback routes if other OSPF or BGP route metrics are over 255.

For more information, see

- [How to Install and Configure the OSPF/RIP/BGP Service](#)

- [How to Configure RIP Router Setup](#)
- [How to Configure the Filter Setup for OSPF and RIP](#)

BGP

The CloudGen Firewall supports BGP4 and BGP4+ versions of the BGP protocol. BGP is an Exterior Gateway Protocol (EGP) and is typically used to connect autonomous systems (AS) of Internet service providers. BGP calculates routing paths based on several pieces of information, such as AS path, IGP metric, multi-exit discriminator, communities, local preferences, next hop, weight, and origin. AS communicate with each other through TCP sessions on port 179. BGP can run between peers in the same AS as well as peers on the border to other AS. It thus acts as an IBGP (Interior Border Gateway Protocol) as well as an Exterior Gateway Protocol (EGP).

- [How to Install and Configure the OSPF/RIP/BGP Service](#)
- [How to Configure BGP Router Setup](#)
- [How to Configure BGP for Inbound Link Failover](#)
- [How to Configure BGP Routing over an IKEv1 IPsec VPN Tunnel](#)

Protocol Comparison

The following table summarizes the feature differences between the supported dynamic routing protocols.

Attribute	OSPF	RIP	BGP
Convergence	Fast	Slow	Slow
Network size	For large and small networks.	Only for small to medium networks due to the fact that max. metric is 15 hops.	For large networks.
Need of device resources	Memory and CPU intensive.	Much less memory and CPU intensive than OSPF.	Depends on the size of the routing table but scales better than OSPF.
Need of network resources	Less than RIP; only small updates are sent.	Bandwidth consuming; whole Routing table is sent (default: every 90 seconds).	Bandwidth consuming while learning network routes from connected AS or while update bursts.

Metric	Is based on bandwidth.	Is based on hop count, no matter how fast the connections are.	Is based on AS Path, IGP-Metric, Multi-Exit Discriminator, Communities, Local Preferences, Next Hop, Weight and Origin.
Design	Hierarchical network possible.	Flat network.	Fully meshed.

HA Operation

The OSPF/RIP service synchronizes externally learned routes with its HA partner. Routes cannot be introduced on the partner while this is "passive" because network routes required to do so are missing. The external routes HA information is thus stored in a file and introduced on the HA system during startup of the OSPF/RIP service. Take-over and startup of the OSPF/RIP service usually take a few seconds. The HA routes are introduced as protocol "extha" (number 245). These routes are then either replaced by newly learned external OSPF or RIP routes (protocols "ospfext" or "ripext") or removed with the HA garbage collection after five minutes.

© Barracuda Networks Inc., 2024 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.