# Configuring Basic IP Routing

SYSTEM ADMINISTRATOR GUIDE

16/1543-CRA 119 1170/1-V1 Uen J

**Disclaimer**

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

**Trademark List**

| | |
|---|---|
| **SmartEdge** | is a registered trademark of Telefonaktiebolaget LM Ericsson. |
| **NetOp** | is a trademark of Telefonaktiebolaget LM Ericsson. |

# Contents

# 1 Overview

This document provides an overview of IP routing and describes the tasks and commands used to configure, monitor, troubleshoot, and administer basic IP routing features through the SmartEdge router. Among the topics discussed are static versus dynamic routing, interior gateway protocols and exterior gateway protocols, protocol distances, and Layer 3 and Layer 4 load balancing.

This document applies to both the Ericsson SmartEdge® and SM family routers. However, the software that applies to the SM family of systems is a subset of the SmartEdge OS; some of the functionality described in this document may not apply to SM family routers.

For information specific to the SM family chassis, including line cards, refer to the SM family chassis documentation.

For specific information about the differences between the SmartEdge and SM family routers, refer to the Technical Product Description *SM Family of Systems* (part number 5/221 02-CRA 119 1170/1) in the `Product Overview` folder of this Customer Product Information library.

IP routing moves information, specifically IPv4 or IPv6 packets, across an internetwork from a source to a destination, typically passing through one or more intermediate nodes along the way. The primary difference between routing and bridging is that the two access different levels of information to determine how to transport packets from source to destination—routing occurs at Layer 3 (the network layer), while bridging occurs at Layer 2 (the link layer) of the Open Systems Interconnection (OSI) reference model.

**Note:** When IP Version 6 (IPv6) addresses are not referenced or explicitly specified, the term IP address can refer generally to IP Version 4 (IPv4) addresses, IPv6 addresses, or IP addressing. In instances where IPv6 addresses are referenced or explicitly specified, the term IP address refers only to IPv4 addresses. For a description of IPv6 addressing and the types of IPv6 addresses, see RFC 3513, *Internet Protocol Version 6 (IPv6) Addressing Architecture*.

In addition to transporting packets through an internetwork, routing involves determining optimal paths to a destination. Routing algorithms use metrics, or standards of measurement, to establish these optimal paths, initializing and maintaining routing tables that contain all route information.

The SmartEdge router routing table stores routes to directly attached devices, static IP routes, and routes learned dynamically from the following routing protocols:

- IPv4: Routing Information Protocol (RIP), the Open Shortest Path First (OSPF) protocol, the Border Gateway Protocol (BGP), and the Intermediate System-to-Intermediate System (IS-IS) routing protocol.

- IPv6: RIPng, OSPFv3, BGP, IS-IS

  **Note:** The terms RIPng and OSPFv3 are not referenced or explicitly specified elsewhere in this document. The term RIP refers to both RIP and RIPng and the term OSPF refers to both OSPF and OSPFv3.

In the routing table, next-hop associations specify that a destination can be reached by sending packets to a next-hop router located on an optimal path to the destination. Routing algorithms must converge rapidly; that is, all routers must agree on optimal routes.

When a network event causes routes either to go down or become unavailable, routers distribute routing update messages that are propagated across networks, causing a universally agreed recalculation of optimal routes. Routing algorithms that converge slowly can cause routing loops or network outages. Many algorithms can quickly select next-best paths and adapt to changes in network topology.

Methods for implementing IP routing, and the protocols used, are described in the sections that follow.

## 1.1 Packet Fragmenting and Reassembling

A SmartEdge node acts as an IP host when an IP packet with one of the local IP addresses of the node in the source address field of the IP packet header is created by an application on a control card, ASE card, or traffic card. IP packets created on a control or ASE card are known as IP control packets; IP packets created on a traffic card are known as IP data packets.

For IPv4 control and data packets, including those created by tunneling applications such as IPsec, GRE, or BGP over MPLS, the Don't Fragment (DF) bit is not set in the IP packet header, allowing interim routers to fragment and reassemble the packets if required enroute, based on the Maximum Transmission Unit (MTU) configured for each hop.

For IPv6 packets, the end systems are responsible for fragmenting the packets on ingress and reassembling the packets on egress, based on the MTU for the data path. If the size of an IPv6 control packet exceeds the Path Maximum Transmission Unit (PMTU) determined by Neighbor Discovery (ND) for the data path, the control or ASE card fragments the packet to match the PMTU value. If the size of an IPv6 data packet exceeds the MTU set on the egress port on the traffic card, the traffic card fragments the IPv6 data packet to match the MTU value. In the current release of the SmartEdge OS, IPv6 data packets are created on traffic cards for BGF over IPv6.

**Warning!**

Risk of data loss for IPv6 data packets created by a traffic card.

To avoid this risk, ensure that the MTU set on the traffic card with the **`mtu`** (Ethernet) command in port configuration mode is set to be less than the PMTU of the data path. If a network-wide PMTU policy is used, set matching port-level MTU and network-wide PMTU values.

## 1.2 Static Versus Dynamic Routing

Static routing involves packet forwarding on the basis of static routes configured by the system administrator. Static routes work well in environments where network traffic is relatively predictable and network topology is relatively simple.

In contrast, dynamic routing algorithms adjust to changing network circumstances by analyzing incoming routing update messages. RIP, OSPF, BGP, and IS-IS all use dynamic routing algorithms. A dynamic routing algorithm can also be supplemented with static routes where appropriate. For example, a router of last resort (to which all unroutable packets are sent) can store information on such packets for troubleshooting purposes.

Some routing algorithms operate in a flat, hierarchy-free space, while others use routing hierarchies. In a flat routing system such as RIP, all routers are peers of all other routers. As networks increase in size, flat routing systems encounter scaling limitations. To address this, some routing protocols allow the administrator to partition the network into hierarchical levels, which facilitates the summary of topology information for anyone located outside the immediate level or area. An example is the OSPF protocol, which supports a two-level hierarchy where area 0 is the backbone area that interconnects all other areas.

## 1.3 IGPs Versus EGPs

Another group of protocols that works to optimize network performance is the Interior Gateway Protocols (IGPs). These optimize the route between points within a network. Examples of commonly used IGPs are RIP, OSPF, and IS-IS.

Exterior Gateway Protocols (EGPs) support route information exchange between different networks. An example of a commonly used EGP is BGP-4. The choice of an optimal path is made based on the cost of the path measured by metrics associated with each link in the network.

IGPs and EGPs have slightly differing administrative designs. An IGP typically runs in an area under a single administrative control; this area is referred to as an autonomous system (AS) or a routing domain. In contrast, an EGP allows two different autonomous systems to exchange routing information and send data across the AS border. Policy decisions in EGPs can be shaped to decide which routing information crosses the border between the two autonomous systems.

## 1.4 IP Routing Protocols

The following IP routing protocols are supported:

- The Virtual Router Redundancy Protocol (VRRP) eliminates the single point of failure that is common in a static default routed environment. A VRRP router controls IP addresses associated with a virtual router. Any of the virtual router's IP addresses on a LAN can then be used as the default first hop router by end hosts, providing a dynamic failover in forwarding responsibility should the VRRP router become unavailable. The main advantage of using VRRP is having a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end host; see *Configuring VRRP*.

- RIP is a distance-vector IGP that uses hop count as its metric. Each router sends all or some of the portion of its routing table, but only to its neighbors. The RIP is widely used for routing traffic in the global Internet; see *Configuring RIP*.

- OSPF is a link-state IGP that uses link-state advertisements (LSAs) to inform other routers of the state of the sender's links. Each router sends only the portion of the routing table that describes the state of its own links to all nodes in the internetwork. LSAs are used to build a complete picture of the network topology, enabling other routers to determine optimal routes to destinations. In OSPF, the autonomous system can be hierarchically organized by partitioning it into areas. Each area contains a group of contiguous networks and hosts. An area border router (ABR) communicates routing information between the areas; see *Configuring OSPF*.

- BGP-4 is a distance-vector EGP, and uses the Transmission Control Protocol (TCP) as its transport protocol. With BGP, a TCP connection is established over which two BGP peers exchange routing information. Routers that belong to the same autonomous system run internal BGP (iBGP), while routers that belong to different autonomous systems run external BGP (eBGP); see *Configuring BGP*.

- IS-IS is an OSI link-state hierarchical routing protocol that floods the network with link-state information. This builds a complete and consistent picture of network topology. Hierarchical routing simplifies backbone design, and the backbone routing protocol can also change without impacting the intra-area routing protocol; see *Configuring IS-IS*.

## 1.5 Administrative (Protocol) Distances

When determining a single optimal route among multiple routes within a single routing protocol, the SmartEdge router selects the route that has the shortest distance. When deciding a best path among routes originating from multiple protocols, the system uses a more complex methodology, in which protocols are assigned an administrative distance (AD). The SmartEdge routing table stores the AD for directly connected, static, eBGP, OSPF, IS-IS, RIP, iBGP, and unknown routes.

Table 1 lists the protocols and their default AD values for routes learned through various protocols.

*Table 1 Administrative Distance Defaults*

| Protocol | Distance Value |
|---|---|
| Directly connected | 0 |
| Static IP | 1 |
| eBGP | 20 |
| OSPF | 110 |
| IS-IS | 115 |
| RIP | 120 |
| iBGP | 200 |
| Unknown | 255 |

## 1.6 Layer 3 and Layer 4 Load Balancing

For increased bandwidth and resilience, multiple links can be bundled into a single logical link by using link aggregation group (LAG), merge point (MP), or equal-cost multipath (ECMP). At the same time, a single traffic flow should always take the same path.

The SmartEdge router maps traffic from multiple flows onto one of the constituent links of the logical link based on a multipath forwarding hash function that includes a tuple of the following Layer 3 information:

- Source IP address of a packet
- Destination IP address of a packet

If you are configuring Link Aggregation Control Protocol (LACP) and ECMP over IP Version 4 (IPv4), then you can use the **`service load-balance ip layer4`** command to include the following Layer 4 information in the load balancing hash algorithm:

- TCP or User Datagram Protocol (UDP) source and destination ports

- Layer 4 protocol value from the IP packet header

Including the source and destination ports balances the load of the traffic among the available paths while keeping packets for a particular data flow in order and preserving the same path for all packets in a given flow.

The inclusion of Layer 4 information is supported only for LACP and ECMP and only for complete IPv4 packets. It is not supported for IP Version 6 (IPv6) or when IP fragmentation is used.

Please be aware that the SmartEdge router automatically reverts to Layer 3 hashing for the following packet types:

- Non-UDP or TCP packets
- Packets with IP options
- Fragmented packets

Use the **`show ip route summary`** command to verify whether Layer 4 load balancing is enabled on a router.

# 2 Configuration and Operations Tasks

**Note:** In this section, the command syntax in the task tables displays only the root command.

To configure basic IP routing, perform the tasks described in the sections that follow.

## 2.1 Configuring Static Routes

Rather than dynamically selecting the best route to a destination, you can configure one or more static routes to the destination. Once configured, a static route stays in the routing table indefinitely. When multiple static routes are configured for a single destination and the outbound interface of the current static route goes down, a backup route is activated, improving network reliability.

You can configure up to eight static routes for a single destination. Each route is assigned a default distance value and cost value. Modifying these values allows you to set a preference for one route over the next. A static route can be overridden by a dynamically learned route with a lower administrative distance.

Among multiple routes with the same destination, preferred routes are selected in the following order:

1. The route with the shortest distance value is preferred first.

2. If two or more routes have the same distance and cost values, the equal cost multipath (ECMP) is preferred.

3. When redistributing static routes, routing protocols ignore the cost value assigned to those static routes. If static routes are redistributed through dynamic routing protocols, only the active static route to a destination is advertised.

### 2.1.1 Configuring Static IPv4 Routing

To configure a static route, perform either of the tasks described in Table 2. Enter all commands in context configuration mode.

*Table 2 Configure Static IP Routing*

| Task | Root Command |
|---|---|
| Configure one or more IP static routes to the same destination. | *ip route* |

**Note:** When configuring a static IPv4 route with the *ip route* command, if you use the ***next-hop-if-name*** argument (instead of the IP address) to specify the next-hop, ensure that the MAC address can be resolved by doing one of the following:

- Manually configuring an ARP entry for the MAC address
- Configuring a proxy ARP server on the LAN the next hop is on

### 2.1.2 Configuring Static IPv6 Routing

To configure a static route, perform either of the tasks described in Table 3. Enter all commands in context configuration mode.

*Table 3 Configure Static IPv6 Routing*

| Task | Root Command |
|---|---|
| Configure one or more IPv6 static routes to the same destination. | *ipv6 route* |

**Note:** Do not configure a link-local address as the next-hop IPv6 address (set with the ***next-hop-ipv6-addr*** argument in the **ipv6 route** command); any attempts to do so are rejected by the OS.

## 2.2 Configuring Additional Basic IP Routing Parameters

To configure basic IP routing parameters, perform the tasks described in Table 4. Enter all commands in context configuration mode, unless otherwise noted.

*Table 4 Configure Additional Basic IP Routing Parameters*

| Step | Task | Root Command | Notes |
|---|---|---|---|
| 1. | Add custom IP martian addresses in the routing table to configure an upper limit for the number of routes installed in an IP routing table. | *ip martian* | — |
| 2. | Configure an upper limit for the number of routes installed in an IP routing table. | *ip maximum-routes* | — |
| 3. | Configure a static route for multicast RPF lookup. | *ip mstatic* | Enter this command in interface configuration mode. |
| 4. | Perform a reverse path forwarding (RPF) check to verify the source IP address on all incoming unicast packets at the specified interface. | *ip verify unicast source* | — |
| 5. | Configure a global router ID for the SmartEdge router. | *router-id (contexts)* | The global router ID must be configured for RSVP to operate correctly. |

*Table 4 Configure Additional Basic IP Routing Parameters*

| Step | Task | Root Command | Notes |
|---|---|---|---|
| 6. | Enable intercontext static routing among non-local contexts. | *service inter-context routing* | Enter this command in global configuration mode.<br><br>This command can only be disabled when there is no instance of non-local context static routing configured on th router. |
| 7. | Specify whether the load-balancing hash algorithm includes Layer 3 information or both Layer 3 and Layer 4 information. The default is Layer 3 information only. | *service load-balance ip* | Enter this command in global configuration mode.<br><br>The inclusion of Layer 4 information in the hash algorithm is supported for forwarding decisions for LACP and ECMP only. It is supporte for just IPv4 and only for complete IP packets. It is not supported for IPv6 or when IP fragmentation is used. |

*Table 4 Configure Additional Basic IP Routing Parameters*

| Step | Task | Root Command | Notes |
|---|---|---|---|
| 8. | Enable the negotiation of the maximum transmission unit (MTU) for IPv4 Transmission Control Protocol (TCP) sessions. | *tcp path-mtu-discovery* | Enter this command in global configuration mode.<br><br>Enabling MTU negotiation has no effect on existing TCP sessions.<br><br>Both the SmartEdge router and the remote router must be configured for MTU negotiation to work properly. |
| 9. | Globally enable IPv6 path maximum transmission unit (PMTU) negotiation on a router and configure the timeout value used for aging PMTUs. | *ipv6 path-mtu-discovery discovery-interval* | Enter this command in global configuration mode.<br><br>When IPv6 PMTU discovery is enabled, the source router fragments any IPv6 packet that exceeds the MTU of the receiving node into multiple smaller packets . In addition, the source router reduces the MTU for that particular hop (the path between the source and the receiving routers) to match the MTU of the receiving router. This process is repeated for all hops in a path. If all nodes support the MTU discovery feature, the MTU of the entire path is eventually discovered, ensuring that all packets sent on that path reach their destination. |

## 2.3 Basic IP Routing Operations

To manage basic IP routing functions, perform the appropriate tasks described in Table 5. Enter the **show** commands in any mode; enter the **clear** and **debug** commands in exec mode.

*Table 5 Basic IP Routing Operations Tasks*

| Task | Root Command |
|---|---|
| Remove routes and a maximum route flag from the IP routing table. | *clear ip maximum-routes* |
| Remove routes from the IP routing or IP multicast table. | *clear ip route* |

*Table 5 Basic IP Routing Operations Tasks*

| Task | Root Command |
|---|---|
| Enable the generation of IP routing debug messages. | *debug ip routing* |
| Enable the generation of IP Version 6 (IPv6) routing debug messages. | *debug ipv6 routing* |
| Enable the generation of debug messages for static IP routes. | *debug static* |
| Display the current static route configuration for the current context. | *show configuration static* |
| Display information about IP routes for specific IP addresses or IP prefixes. | *show ip route* |
| Display information about all IP routes. | *show ip route all* |
| Display information about BGP routes. | *show ip route bgp* |
| Display routes for Routing Information Base (RIB) clients. | *show ip route client* |
| Display routes for directly connected networks. | *show ip route connected* |
| Display routes for Forwarding Information Base (FIB) clients. | *show ip route fib-client* |
| Display hidden IP routes, which are added internally. | *show ip route hidden* |
| Display information about IS-IS routes. | *show ip route isis* |
| Display information about IP martian routes. | *show ip route martian* |
| Display all unicast-dependent multicast routing table information. | *show ip route multicast* |
| Display information about next-hop IP route information. | *show ip route next-hop* |
| Display information about OSPF routes. | *show ip route ospf* |
| Display next-hop or prefix information registered in the Routing Information Base (RIB). | *show ip route registered* |
| Display information about RIP routes. | *show ip route rip* |
| Display information about static routes. | *show ip route static* |
| Display information about subscriber routes. | *show ip route subscriber* |
| Display information about IP routes. | *show ip route summary* |
| Display summary information about all IP routes. | *show ip route summary all-context* |
| Display IP route information for the controller card. | *show ip route xcrp* |
| Display IPv6 route information. | *show ipv6 route* |
| Display static route information. | *show static route* |

# 3 Configuration Examples

The following example routes packets for network **10.10.0.0/16** via interface, **enet1:**

```
[local]Redback(config-ctx)#ip route 10.10.0.0/16 enet1
```

The following example defines a default route through interface **atm5**. Because no cost is defined, this route uses a cost of 0, and is therefore used as the active route. If this route goes away, the second and third routes alternate because they have the same distance and cost:

```
[local]Redback(config-ctx)#ip route 0.0.0.0/0 atm5

[local]Redback(config-ctx)#ip route 0.0.0.0/0 10.1.1.1 cost 2

[local]Redback(config-ctx)#ip route 0.0.0.0/0 172.21.200.254 cost 2
```

The following example displays the routing table for the routes configured in the previous examples:

```
[local]Redback>show ip route

Codes: C - connected, S - static, R - RIP, e B - EBGP, i B - IBGP
       O   - OSPF, IA - OSPF inter area, N1  - OSPF NSSA external type 1
       N2  - OSPF NSSA external type 2,  E1  - OSPF external type 1
       E2  - OSPF external type 2
       i   - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2
       >   - Active Route

Type    Network             Next Hop        Dist  Metric    UpTime  Interface

> S     0.0.0.0/0                              1       0      3w0d  atm5
> S     10.10.0.0/16                           1       0      3w0d  enet
```

The following example shows the routing table after the default route through interface **atm5** is removed:

```
[local]Redback>show ip route

Codes: C  - connected, S - static, R - RIP, e B - EBGP, i B - IBGP
       O  - OSPF, IA - OSPF inter area, N1  - OSPF NSSA external type 1
       N2 - OSPF NSSA external type 2,  E1  - OSPF external type 1
       E2 - OSPF external type 2
       i  - IS-IS, L1 - IS-IS level-1,  L2  - IS-IS level-2
       >  - Active Route

Type    Network             Next Hop         Dist  Metric    UpTime  Interface

> S     0.0.0.0/0           10.1.1.1            1       2      3w0d
> S                         172.21.200.254
> S     0.10.0.0/16                             1       0      3w0d  enet
```

**Note:** Only the default route for interface **atm5** displays.

The following example shows how to configure the load balancing algorithm to include both Layer 3 and Layer 4 information:

```
[local]Redback# configure
[local]Redback(config)# service load-balance ip layer-4
```

# Glossary

**ABR**
area border router

**BGP**
Border Gateway Protocol

**eBGP**
external BGP

**EGPs**
Exterior Gateway Protocols

**FIB**
Forwarding Information Base

**iBGP**
internal BGP

**IGPs**
Interior Gateway Protocols

**IPv4**
IP Version 4

**IPv6**
Internet Protocol Version 6

**IPv6**
IP Version 6

**IS-IS**
Intermediate System-to-Intermediate System

**LACP**
Link Aggregation Control Protocol

**OSI**
Open Systems Interconnection

**OSPF**
Open Shortest Path First

**PMTU**
IPv6 path maximum transmission unit

**RIB**
Routing Information Base

**RIP**
Routing Information Protocol

**TCP**
Transmission Control Protocol

**UDP**
User Datagram Protocol

**VRRP**
Virtual Router Redundancy Protocol

# Reference List


[1] *General Troubleshooting Guide*

[2] *Data Collection Guideline for the SmartEdge Router*