

# Load Balancing

---

## SYSTEM ADMINISTRATOR GUIDE

## **Copyright**

© Ericsson AB 2009–2011. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

## **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

<b>1</b>	<b>Load-Balancing Operations</b>	<b>1</b>
1.1	Load Balancing for Each Link Group	1
1.2	Load Balancing for Each Configuration	2
1.3	Load Balancing for Each Traffic Source	3
<b>2</b>	<b>Load-Balancing Commands</b>	<b>9</b>





# 1 Load-Balancing Operations

This guide describes how to enable, monitor, and administer load balancing of traffic on the SmartEdge router.

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.

## 1.1 Load Balancing for Each Link Group

Some link groups do not support load balancing, such as the access link group. Traffic in an access link group is distributed among the active ports but is not load balanced. Circuits in the link bundle are assigned group IDs and distributed per group ID (not by source and destination address as in load balancing) so that all traffic in the same group goes through the same port. This section describes load balancing link groups support.

Load balancing is supported by the following link-group types:

- The Multilink Frame Relay (MFR) and MultiLink Point to Point Protocol (MLPPP) link groups use round-robin to transmit packets, so the load balancing among link-group members is packet-by-packet.
- The `dot1q`, `ether`, and `hdlc` link-groups all work the same way in terms of packet load balancing. These link-groups are referred generically as the LG link-group, or LAG.

**Note:** If the packets are coming from a routed domain and going to an L2 domain through a BVI port, and if an `ether` or a `dot1q` link-group is part of the bridge group associated with the BVI port, either Layer 3 or Layer 4 (five-tuple hashing) load balancing can be used.



## 1.2 Load Balancing for Each Configuration

Certain configurations affect load balancing. For example, the following types of configurations change the load-balancing behavior:

- The **service load-balance ip layer-3** configuration command causes the load balancing calculation to only include the source and destination IP addresses. For more information on this type of configuration, see the *Layer 3 and Layer 4 Load Balancing* section in *Configuring Basic IP Routing* and the **service load-balance ip** command.
- The **service load-balance ip layer-4** configuration command (5-tuple hashing) causes the load balancing calculation to include the source and destination IP addresses, the IP protocol, and the source and destination ports for User Datagram Protocol (UDP) and Transmission Control Protocol (TCP) streams. The **show ip route summary** command lists `Load-balance: Layer 4` on the top line of the summary when both Layer 3 and Layer 4 load balancing are enabled on the router. For more information on this type of configuration, see the *Layer 3 and Layer 4 Load Balancing* section in *Configuring Basic IP Routing* and the **service load-balance ip** command.
- The **pseudowire multi-path** configuration command impacts load-balancing behavior on ingress or egress Provider Edge (PE) node for Layer 2 Virtual Private Network (L2VPN) and Virtual Private LAN Services (VPLS) pseudowire (PW) traffic. By default, the legacy load balancing scheme is enabled and a pseudowire travels over one member link in a LAG bundle so that PW traffic is sent based on the inner virtual circuit label (pseudowire label) as a single link in the link-group. When **pseudowire multi-path** is enabled, PW traffic is distributed among the links in the link-group, spanning over different member links, taking multiple paths based on the IP source and destination address. The **show pseudowire** command returns `multi-path enabled` when pseudowire multipath load balancing is enabled on the router. For more information on this type of configuration, see the *Enabling Load Balancing on Pseudowires* section in *Configuring L2VPN* and the **pseudowire multi-path** command.
- In a **router ldp** configuration (where a Label Distribution Protocol (LDP) routing instance for a context is enabled and the router is in LDP router configuration mode), the **ecmp-transit** configuration command enables equal cost multipath (ECMP) on label-switched path (LSP) transit nodes. The **show config ldp** command lists `ecmp-transit` when ECMP load balancing is enabled on the router. For more information on this type of configuration, see the *Configuring an LDP Routing Instance* section in *Configuring LDP* and the **ecmp-transit** command.

To enable load-balancing, perform the appropriate task listed in Table 1 (For information on the commands used to do these tasks, see the *Command List*). Enter the **service load-balance**, **pseudowire multi-path**, and **ecmp-transit** commands in global configuration mode; enter the **show** commands in any mode.



Table 1 Load Balancing Tasks

Task	Command
Enable Layer 3 load balancing.	<i>service load-balance ip layer-3</i>
Enable Layer 3 and Layer 4 load balancing.	<i>service load-balance ip layer-4</i>
Enable pseudowire multi-path load balancing.	<i>pseudowire multi-path</i>
Enable ECMP load balancing.	<i>ecmp-transit</i>
Display whether Layer 4 load balancing is enabled.	<i>show ip route summary</i>
Display whether pseudowire multi-path load balancing is enabled.	<i>show pseudowire</i>
Display whether ecmp-transit load balancing is enabled.	<i>show configuration ldp</i>

### 1.3 Load Balancing for Each Traffic Source

The following tables are divided by load balancing for traffic on the Ingress Edge Router and P (transit), meaning the Provider Label Switch Router in the center of a Multiprotocol Label Switching (MPLS) network (this router is not a PE router, this router gets routes from a third party that is not in the VPN, and this router gives those routes to a PE router). The tables show how load balancing is accomplished depending on the source of the traffic flow and applications, traffic type, paths, link groups, and other aspects of the configuration described in the legend.

The following tables show that the traffic load balancing is completed in the same manner for the following applications for all paths:

- For IP traffic from all applications for both P (transit) and Ingress Edge Router: Plain IP, Generic Routing Encapsulation (GRE) with IP over IP (IPIP), Soft GRE, and IP-based MPLS (L3VPN) applications traffic is balanced by a load balancing hash algorithm that uses information from a tuple of the source IP address and the destination IP address of a packet (ip src + dst) for all paths by default.
- L2TP application traffic is balanced by a load balancing hash algorithm that uses information from the tunnel and session ID (tunnel + session id) for all paths.
- L2 Bridge application traffic is balanced by a load balancing hash algorithm that uses information from the source and destination medium access control (MAC) address.

Following is the legend for the numbers used (Legend: n) in Table 2, Table 3, Table 4, and Table 5:



1. Multiple single physical links.
2. One path is a link group (`dot1q`, `ether`, or `hdlc`) with more than one physical links. The LG is treated as a single link first, and load-balancing first uses the method in Legend: 1. Then, distribution among LG members uses the method in Legend: 3.
3. The path is a link-group (`dot1q`, `ether`, or `hdlc`) with more than one member, and the load is distributed among group members.
4. If 5-tuple hashing is configured, load balancing is based on ip src, dst, ip protocol, udp/tcp src port, and dst port.
5. If queuing is enabled on a L2TP network server (LNS), the virtual node ID is used for load balancing. If L2TP ECMP is not enabled, the LNS uses inner ip src + dst for load balancing. (Only for ingress PE: The L2TP access concentrator (LAC) uses outer ip src + dst for load balancing.)
6. For P (transit): If the first nibble following the MPLS label stack is 0x4 or 0x6, it is treated respectively as an IPv4 or IPv6 payload.

Table 2 *Ingress Edge Router without pseudowire multi-path*

Applications	Traffic Type	Ingress Edge Router without pseudowire multi-path		
		Multiple Paths Legend: 1	Multiple Paths + LG Legend: 2	LG Legend: 3
Plain IP	IP	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
GRE with IP over IP	IP	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4
Soft GRE	IP	inner ip src + dst Legend: 4	inner ip src + dst Legend:4	inner ip src + dst Legend: 4
IP-based MPLS (L3VPN)	IP	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
VPLS	L2 Ethernet or other	inner label	inner label	inner label
L2VPN	L2 Ethernet or other	inner label	inner label	inner label
L2TP	any	tunnel + session id Legend: 5	tunnel + session id Legend: 5	tunnel + session id Legend: 5
L2 Bridge	L2 Ethernet	NA	NA	src + dst mac



Table 3 shows how the addition of pseudowire multi-path configuration changes the information that is used by the load balancing hash algorithm for VPLS and L2VPN applications:

- VPLS and L2VPN applications with IP payload traffic are balanced by a load balancing hash algorithm that uses information from a tuple of the source IP address and the destination IP address of a packet (ip src + dst) for all paths.
- VPLS and L2VPN applications with L2 Ethernet (non IP payload) traffic is balanced by a load balancing hash algorithm that uses information from the source and destination MAC address (src + dst mac) only if the `_next_header` is not IP.

Table 3 shows that the addition of pseudowire multi-path configuration does not change the information that is used by the load balancing hash algorithm for the other applications.

*Table 3 Ingress Edge Router with pseudowire multi-path*

Applications	Traffic Type	Ingress Edge Router with <code>pseudowire multi-path</code>		
		Multiple Paths Legend: 1	Multiple Paths + LG Legend: 2	LG Legend: 3
Plain IP	IP	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
GRE with IP over IP	IP	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4
Soft GRE	IP	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4
IP-based MPLS (L3VPN)	IP	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
VPLS	L2 Ethernet ( <code>_next_header</code> is IP)	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
	L2 Ethernet ( <code>_next_header</code> is not IP)	src + dst mac	src + dst mac	src + dst mac
	Other	inner label	inner label	inner label
L2VPN	L2 Ethernet ( <code>_next_header</code> is IP)	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
	L2 Ethernet ( <code>_next_header</code> is not IP)	src + dst mac	src + dst mac	src + dst mac
	Other	inner label	inner label	inner label



Table 3 Ingress Edge Router with pseudowire multi-path

Applications	Traffic Type	Ingress Edge Router with pseudowire multi-path		
		Multiple Paths Legend: 1	Multiple Paths + LG Legend: 2	LG Legend: 3
L2TP	any	tunnel + session id Legend: 5	tunnel + session id Legend: 5	tunnel + session id Legend: 5
L2 Bridge	L2 Ethernet	NA	NA	src + dst mac

Table 4 shows how P (transit) load balancing is the same for all applications with or without pseudowire multi-path, but the hash algorithm is different for non IP traffic:

- All application IP payload traffic is balanced by a load balancing hash algorithm that uses information from a tuple of the source IP address and the destination IP address of a packet (ip src + dst) for all paths by default. If the first nibble following the MPLS label stack is 0x4 or 0x6, it is treated respectively as an IPv4 or IPv6 payload.
- All application non IP payload traffic is balanced by a load balancing hash algorithm that uses information from the entire label stack and the incoming slot/port for all paths. If the first nibble following the MPLS label stack is 0x4 or 0x6, it is treated respectively as an IPv4 or IPv6 payload.
- L2TP traffic is the same for Ingress Edge Router and P (transit).

Table 4 P (transit) with or without pseudowire multi-path

Applications	Traffic Type	P (transit) with or without pseudowire multi-path		
		Multiple Paths Legend: 1	Multiple Paths + LG Legend: 2	LG Legend: 3
All	IP Legend: 6	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
	non IP Legend: 6	entire label stack + incoming slot/port	entire label stack + incoming slot/port	entire label stack + incoming slot/port

Table 5 is a detailed version of Table 4.

Table 5 P (transit) with or without pseudowire multi-path Detailed

Applications	Traffic Type	P (transit) with or without pseudowire multi-path Detailed		
		Multiple Paths Legend: 1	Multiple Paths + LG Legend: 2	LG Legend: 3
Plain IP	IP	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
GRE with IP over IP	IP	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4	inner ip src + dst Legend: 4



Table 5 P (transit) with or without pseudowire multi-path Detailed

Applications	Traffic Type	P (transit) with or without pseudowire multi-path Detailed		
		Multiple Paths Legend: 1	Multiple Paths + LG Legend: 2	LG Legend: 3
Soft GRE	IP	ip src + dst	ip src + dst	ip src + dst
IP-based MPLS (L3VPN)	IP	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
VPLS	L2 Ethernet (_next_header is IP)	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
	L2 Ethernet (_next_header is not IP)	entire label stack + incoming slot/port	entire label stack + incoming slot/port	entire label stack + incoming slot/port
	Other	entire label stack + incoming slot/port	entire label stack + incoming slot/port	entire label stack + incoming slot/port
L2VPN	L2 Ethernet (_next_header is IP)	ip src + dst Legend: 4	ip src + dst Legend: 4	ip src + dst Legend: 4
	L2 Ethernet (_next_header is not IP)	entire label stack + incoming slot/port	entire label stack + incoming slot/port	entire label stack + incoming slot/port
	Other	entire label stack + incoming slot/port	entire label stack + incoming slot/port	entire label stack + incoming slot/port
L2TP	any	tunnel + session id Legend: 5	tunnel + session id Legend: 5	tunnel + session id Legend: 5
L2 Bridge	L2 Ethernet	NA	NA	src + dst mac





## 2 Load-Balancing Commands

This following list provides references to the documentation that describes the syntax and usage guidelines for the commands used to enable load balancing and to show whether load balancing is enabled:

- *ecmp-transit*
- *pseudowire multi-path*
- *service load-balance ip layer 3*
- *service load-balance ip layer 4*

For more information on pseudowire multi-path load balancing, see the *Enabling Load Balancing on Pseudowires* section in *Configuring L2VPN* and the `pseudowire multi-path` command. For more information on Layer 3 and Layer 4 load balancing, see the *Layer 3 and Layer 4 Load Balancing* section in *Configuring Basic IP Routing* and the `service load-balance ip` command. For more information on ECMP load balancing, see the *Configuring an LDP Routing Instance* section in *Configuring LDP* and the `ecmp-transit` command.