

QoS for 802.3ad Link Aggregation Interfaces Overview

You can configure QoS for 802.3ad link aggregation interfaces. To ensure that QoS is applied properly to the interface column, you configure the QoS profile using either a hashed loadbalancing scheme or a subscriber loadbalancing scheme.

Types of Load Balancing

For hashed load balancing, you configure the scheduler hierarchy with Ethernet queues, and the system replicates them on each link within the link aggregation group (LAG). The system demultiplexes each packet to one of the active links in the LAG using a random hash generated by fields in the packet header. For example, when an IP packet is routed to a LAG, the hash algorithm is based on the IP Source Address and Destination Address in the IP header.

For subscriber load balancing, you configure the scheduler hierarchy with IP, VLAN, and S-VLAN queues and the system allocates them to individual ports in the LAG. The system demultiplexes each packet to an active link based on the subinterface underlying the egress interface. For example, when an IP packet is routed to an IP interface over a LAG, the system binds the underlying VLAN, PPPoE, or MPLS subinterface to one of the active links in the LAG. The packet is transmitted over the interface.

Most network operators configure QoS over 802.3ad LAGs using subscriber load balancing to take advantage of subscriber class-based queueing (SCBQ) features. However, configuring hashed load balancing is useful for achieving fine-grained distribution of multicast VLAN traffic or for any high bandwidth VLAN that does not require shared shaping.

To ensure that QoS is symmetrically applied to all the links, the router periodically rebalances the traffic on the LAG. You can control the loadbalancing parameters.

Munged QoS Profiles and Load Balancing

To determine whether to use hashed load balancing or subscriber load balancing, the system munges a QoS profile for a subscriber.

In typical Ethernet configurations, the munged QoS profile for a given subscriber interface comprises the accumulated rules of the QoS profiles attached below the subscriber interface in the interface column. Rules in higher-attached QoS profiles override or eclipse rules in lower-attached QoS profiles. For example, rules from specific interface attachments such as a VLAN override those from attachments at S-VLANs or ports.

When applying QoS to LAGs, the system uses a modified algorithm to munge QoS profile attachments. The system automatically builds the munged QoS profile using the rules in the QoS profile attached at the LAG interface.

For example, the munged QoS profile for VLAN 0,0 consists of the munge of:

- Attachment 1—QoS profile attached to the VLAN

- Attachment 2—QoS profile attached to the S-VLAN
- Attachment 3—QoS profile attached to the LAG

If there is no QoS profile attached to the LAG, the system locates the lag-default QoS profile indicated in the **qos-port-type-profile** command.

If the resulting QoS profile specifies only Ethernet queues, the system uses the hash algorithm to balance the links. If the resulting QoS profile specifies any VLAN, IP, or L2TP-Session queues, then the system uses subscriber load balancing.

802.3ad Link Aggregation and QoS Parameters

You can create parameter instances for IEEE 803.ad LAG interfaces. A parameter instance for LAG can control an Ethernet port or a node, but you cannot create parameter instances for the Ethernet interfaces within the LAG.

For example, a LAG instance can specify a shaping rate of 100 Mbps on an Ethernet port or a group node. The system shapes all Ethernet ports or group nodes to the same rate within the LAG. Using load balancing, the system strives to balance the traffic each link equally.

QoS and Ethernet Link Redundancy

You can configure Ethernet link redundancy for LAG interfaces. When you configure QoS for those links, be sure to consider the following behaviors.

Active Link Failure and QoS

When an active link fails, traffic that is hashed-load balanced is redirected onto the remaining active links in the LAG. Traffic that is hashed-load balanced might be lost on the disabled link, but from the moment of switchover, traffic arriving from the fabric on the egress line module is directed towards one of the remaining hashed load-balanced queues.

Subscriber loadbalanced traffic takes more time to reestablish on active links because of the amount of computation (approximately 3 ms per subscriber). During this time period, traffic directed to the disabled link might be lost.

Administratively Disabling a Link and QoS

When a link is administratively disabled, the system immediately redirects traffic from the link to other links in the LAG.

Adding a New Link to the LAG and QoS

When you add a new link to the LAG, the system immediately sends traffic that is hashed-load balanced to the link. Traffic that is subscriber-load balanced moves to the new link as new subscribers log in. The system automatically rebalances traffic to the new link based on the load rebalance configuration for the LAG.

- Related Topics**
- Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview
 - Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview
 - For more information about configuring the lag-default QoS profile for default subscriber load balancing, see *Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups*
 - For more information about Ethernet link redundancy, see *JUNOS Link Layer Configuration Guide*
 - For more information about configuring QoS parameters, see *Parameter Definition Attributes for QoS Administrators Overview*
 - For more information about the munge algorithm, see *Munged QoS Profile Overview*
 - For a list of modules that support 802.3ad link aggregation, see the *ERX Module Guide* and the *E120 and E320 Module Guide*

Published: 2010-03-24