

QoS in MC-MLPPP

In This Section

This section provides information to configure MC-MLPPP using the command line interface.

Topics in this section include:

- [Overview on page 912](#)
- [Basic Configurations on page 917](#)
- [Configuring MC-MLPPP on page 918](#)

Overview

If the user enables the multiclass option under an MLPPP bundle, the MDA egress data path provides a queue for each of the 4 classes of MLPPP. The user configures the required number of MLPPP classes to use on a bundle. The forwarding class of the packet, as determined by the ingress QoS classification, is used to determine the MLPPP class for the packet and hence which of the four egress MDA queues to store the packet. The mapping of forwarding class to MLPPP class is a function of the user configurable number of MLPPP classes. The default mapping for a 4-class, 3-class, and 2-class MLPPP bundle is shown in [Table 64](#).

Table 64: Default Packet Forwarding Class to MLPPP Class Mapping

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class 4-class bundle	MLPPP Class 3-class bundle	MLPPP Class 2-class bundle
7	NC	Expedited	0	0	0
6	H1	Expedited	0	0	0
5	EF	Expedited	1	1	1
4	H2	Expedited	1	1	1
3	L1	Non-Expedited	2	2	1
2	AF	Non-Expedited	2	2	1
1	L2	Non-Expedited	3	2	1
0	BE	Non-Expedited	3	2	1

[Table 65](#) shows a different mapping enabled when the user applies one of three pre-defined egress QoS profiles in the 4-class bundle configuration only.

Table 65: Packet Forwarding Class to MLPPP Class Mapping

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class (MLPPP Egress QoS profile 1, 2, and 3)
7	NC	Expedited	0
6	H1	Expedited	0
5	EF	Expedited	1
4	H2	Expedited	2
3	L1	Non-Expedited	2
2	AF	Non-Expedited	2
1	L2	Non-Expedited	2
0	BE	Non-Expedited	3

The MLPPP class queue parameters and its scheduling parameters are also configured by applying one of the three pre-defined egress QoS profiles to an MLPPP bundle.

Table 66 and Figure 50 provide the details of the class queue threshold parameters. Packets marked with a high drop precedence, such as out-of-profile, by the service or network ingress QoS policy will be discarded when any class queue reaches the OOP threshold. Packet with a low drop precedence marking, such as in-profile, will be discarded when any class queue reaches the max threshold.

Table 66: MLPPP Class Queue Threshold Parameters

	Class 0		Class 1		Class 2		Class 3	
Queue Threshold (in ms @ Available bundle rate)	Max	Oop	Max	Oop	Max	Oop	Max	Oop
2-Class Bundle Default Egress QoS Profile	250	125	750	375	N/A	N/A	N/A	N/A
3-Class Bundle Default Egress QoS Profile	50	25	200	100	750	375	N/A	N/A
4-Class Bundle Default Egress QoS Profile	10	5	50	25	150	75	750	375
4-Class Bundle Egress QoS Profile 1	25	12	5	3	200	100	1000	500
4-Class Bundle Egress QoS Profile 2	25	12	5	3	200	100	1000	500
4-Class Bundle Egress QoS Profile 3	25	12	5	3	200	100	1000	500





Figure 50: MLPPP Class Queue Thresholds for In-Profile and Out-of-Profile Packets

Table 67 and Table 51 provide the details of the class queue scheduling parameters.

Table 67: MLPPP Class Queue Scheduling Parameters

4-class MLPPP Egress QoS Profile	WRR Parameters			
	MIR	W1	W2	W3
Profile 1	85%	<1%	66%	33%
Profile 2	90%	<1%	89%	10%
Profile 3	85%	<1%	87%	12%

**Figure 51: MLPPP Class Queue Scheduling Scheme**

Note that all queue threshold and queue scheduling parameters are adjusted to the available bundle rate. If a member link goes down or a new member link is added to the bundle, the scheduling parameters MIR, W1, W2, W3, as well as the per class queue thresholds OOP and max are automatically adjusted to maintain the same values.

Class 0 queue is serviced at MLPPP at available bundle rate. Class 1 queue is guaranteed a minimum service rate but is allowed to share additional bandwidth with class 2 and 3 queues based on the configuration of WRR weight W1.

Class queues 2 and 3 can be given bandwidth guarantee by limiting MIR of class 1 queue to less than 100% and by setting the WRR weights W1, W2 and W3 to achieve the desired bandwidth distribution among all three class queues.

Note that there is one queue per bundle member link to carry link control packets, such as LCP: PPP, and which are serviced with strict priority over the 4 class queues (not shown).

In the default 2-class, 3-class, and 4-class egress QoS profile, the class queues are service with strict priority in ascending order of class number.

Ingress MLPPP Class Reassembly

For a MLPPP bundle with the multi-class option enabled, there is a default profile for setting the re-assembly timer value for each class. When the pre-defined MLPPP ingress QoS profile 1 is applied to a 4-class bundle, the values of the timers are modified as shown in [Table 68](#).

Table 68: MLPPP Ingress QoS Profile: Reassembly Timers (msec)

	Class 0	Class 1	Class 2	Class 4
MLPPP ingress QoS default profile (2-Class bundle)	25ms	25ms	NA	NA
MLPPP ingress QoS default profile (3-Class bundle)	25ms	25ms	25ms	NA
MLPPP ingress QoS default profile (4-Class bundle)	25ms	25ms	100ms	1000ms
MLPPP ingress QoS profile 1 (4-class bundle)	10	10	100	1000

Configuring MC-MLPPP QoS Parameters

A 4-class MLPPP bundle can be configured. This feature cannot be used with MC-MLPPP bundles with fewer than 4 classes.

The following describe the parameters and the configuration processes and rules:

1. The user creates an ingress QoS profile, `mlppp-profile-ingress`, to configure the desired value of the ingress per-class re-assembly timer. Ingress QoS profile #1 is reserved to the pre-defined profile with parameter values in [Table 68](#). The user is allowed to edit this profile and change parameter values. However, the default value of a parameter when a user creates a profile with a profile-id higher than 1, or performs the no option on the parameter, will always be the one in [Table 68](#) for the ingress QoS Profile #1 regardless what parameter value the edited Profile #1 has at that point in time.
2. The user creates an egress QoS profile, `mlppp-profile-egress`, to configure the desired values for the per-class queue and queue scheduling parameters. The user is also able to configure the mapping of the system forwarding classes to the MLPPP classes. Egress QoS profiles #1, 2, and 3, are reserved to the pre-defined profiles with parameter values shown in [Table 65](#), [Table 66](#), or [Table 67](#). The user is allowed to edit these profiles and change parameter values. However, the default value of a parameter when a user creates a profile with a profile-id

higher than 3, or when the user performs the no option on the parameter, will be the one shown in [Table 65](#), [Table 66](#), or [Table 67](#) for the egress QoS Profile #1. This is regardless of the parameter value the edited profiles have at that point in time.

3. A maximum of 128 ingress QoS profiles and 128 egress QoS profiles can be created on the system.
4. The values of the ingress per-class re-assembly timer are configured in the ingress QoS profile.
5. The mapping of the system forwarding classes to the MLPPP Classes are configured in the egress QoS profile. There is a many-to-one relationship between the system FC and an MLPPP class. See [Table 65](#) for the mapping when one of the three pre-defined 4-class egress QoS profiles is selected.
6. The maximum size for each MLPPP class queue in units of msec at the available bundle rate is configured in the egress QoS profile. This is referred to as max in [Figure 50](#) and as max-queue-size in CLI. The out-of-profile threshold for an MLPPP class queue, referred to as oop in [Figure 50](#), is not directly configurable and is set to 50% of the maximum queue size rounded up to the nearest higher integer value.
7. The MLPPP class queue scheduling parameters is configured in the egress QoS profile. The minimum information rate, referred to as MIR in [Figure 51](#) and mir in CLI, applies to Class 1 queue only. The MIR parameter value is entered as a percentage of the available bundle rate. The WRR weight, referred to as W1, W2, and W3 in [Figure 51](#) and weight in CLI, applies to class 1, class 2, and class 3 queues. Note that W1 in [Figure 51](#) is not configurable and is internally set to a value of 1 such that Class 1 queue shares 1% of the available bundle rate when the sum of W1, W2, and W3 equals 100. W2 and W3 weights are integer values and are user configurable such that Class 2 queue shares and Class 3 queue shares of the available bundle rate.
8. The user applies the ingress and egress QoS profiles to a 4-class MLPPP bundle for the configured QoS parameter values to take effect on the bundle.
9. The following operations require the bundles associated with a QoS profile to be shutdown to take effect.
 - A change of the numbered ingress or egress QoS profile associated with a bundle.
 - A change of the bundle associated ingress or egress QoS profile from default profile to a numbered profile and vice-versa.
10. The following operations can be performed without shutting down the associated bundles:
 - Changes to any parameters in the ingress and egress QoS profiles.

Basic Configurations

Configure an egress and ingress MLPPP profile.

The following displays the profile configuration examples:

```
A:ALA-12>config>qos# info
#-----
mlppp-profile-ingress 2 [create]
  description my-4-class-bundle-ingress-profile
  class class 0
    reassembly-timeout 10
  class class 1
    reassembly-timeout 100
  class class 2
    reassembly-timeout 500
  class class 3
    reassembly-timeout 1000
mlppp-profile-egress 4 [create]
  description my-4-class-bundle-egress-profile
  fc be mlppp-class 3
  fc l2 mlppp-class 2
  fc af mlppp-class 2
  fc l1 mlppp-class 2
  fc h2 mlppp-class 2
  fc ef mlppp-class 1
  fc h1 mlppp-class 1
  fc nc mlppp-class 0
  class 0
    max-queue-size 10
  class 1
    max-queue-size 50
    mir 25
    class 2
      max-queue-size 100
      weight 25
    class 3
      max-queue-size 1000
      weight 75

#-----
A:ALA-12>config>qos#

A:ALA-12>config>port>ml-bundle# info
#-----
multilink-bundle bundle-ppp-6/1.1
multilink-bundle
  fragment-threshold 384
  mlppp
  multiclass 4
    ingress
      qos-profile 2
    exit
    egress
      qos-profile 4
    exit
  exit
```

```
exit
member 1/1/1.1.1.1
member 1/1/1.1.2.1
member 1/1/1.1.3.1
member 1/1/1.1.4.1
minimum-links 2
```

Configuring MC-MLPPP

Use the following CLI syntax to create an MC-MLPPP:

```
CLI Syntax:config>qos# mlppp-profile-ingress 2 [create]
      description dmy-4-class-bundle-ingress-profile
      class class 0
            reassembly-timeout 10
      class class 1
            reassembly-timeout 100
      class class 2
            reassembly-timeout 500
      class class 3
            reassembly-timeout 1000
```

```
CLI Syntax:config>qos# mlppp-profile-egress 4 [create]
      description dmy-4-class-bundle-egress-profile
```

QoS in MLFR and FRF.12 Fragmentation

The following sections describe MLFR and FRF.12 feature descriptions and implementation.

QoS in MLFR

The MLFR feature introduces the following new QoS requirements on the MDA:

- Four MDA queues are provided per MLFR bundle to store the fragments of the FR SAP packets. One queue per FR scheduling class. Fragments of all FR SAPs of a given scheduling class will be queued in the same queue.
- The fragments of an FR SAP packet must be queued in the same MDA queue regardless of which forwarding class queue they use in the IOM.

The FR class queue parameters and its scheduling parameters are configured by applying an Egress QoS profile to an MLFR bundle.

[Table 69](#) and [Figure 52](#) provide the details of the class queue threshold parameters. Packets that are marked with high drop precedence, for example, out-of-profile, by the service or network ingress QoS policy will be discarded when any class queue reaches the OOP threshold. Packets with a low drop precedence marking, for example, in profile, will be discarded when any class queue reaches the max threshold. Only the max threshold is user configurable and is referred to as max-queue-size in the CLI. The OOP threshold is always set to 50% of the max threshold.

Table 69: Default FR Class Queue Threshold Parameters

	Class 0		Class 1		Class 2		Class 3	
Queue threshold (in ms@available bundle rate)	Max	Oop	Max	Oop	Max	Oop	Max	Oop
	10	5	50	25	150	75	750	375




Figure 52: FR Class Queue Thresholds for In-Profile and Out-of-Profile Packets

Table 70 and Figure 53 provide the details for the class queue scheduling parameters for an MLFR bundle.

Table 70: Default FR Class Queue Scheduling Parameters

WRR Parameters			
MIR	W1	W2	W3
90%	<1%	89%	10%




Figure 53: FR Class Queue Scheduling for an MLFR Bundle

The minimum information rate, referred to as MIR in Figure 53 and MIR in CLI, applies to Class 1 queues only. The MIR parameter value is entered as a percentage of the available bundle rate. The WRR weight, referred to as W1, W2, and W3 in Table 70 and weight in CLI, applies to class 1, class 2, and class 3 queues. Note that W1 is not configurable and is internally set to a value of 1 such that Class 1 queue shares 1% of the available bundle rate when the sum of W1, W2, and W3 equals 100. W2 and W3 weights are integer values and are user configurable such that Class 2

queue shares $W2/(W1+W2+W3)$ and Class 3 queue shares $W3/(W1+W2+W3)$ of the available bundle rate.

Note that all queue threshold and queue scheduling parameters are adjusted to the available bundle rate. If a member link goes down or a new member link is added to the bundle, the scheduling parameters MIR, W1, W2, W3, as well as the per class queue thresholds OOP and max are automatically adjusted to maintain the same values.

In addition, operator user can configure the value of the FR scheduling class ingress re-assembly timeout for an MLFR bundle. The default values of the timers are shown in [Table 71](#).

Table 71: Default FR Ingress QoS Profile: Reassembly Timers (msec)

Class 0	Class 1	Class 2	Class 3
10	10	100	100

The following operations require the bundles or links associated with a QoS profile to be shutdown to take effect.

- A change of the numbered ingress or egress QoS profile.

The following operations can be performed without shutting down the associated bundles or links:

- Changes of parameters in the currently assigned ingress and egress QoS profiles.

QoS in FRF.12 End-to-End Fragmentation

When end-to-end fragmentation is enabled on an FR SAP, the queuing and scheduling on the MDA is reusing the existing FR SAP high and low priority queues are show in [Figure 54](#):




Figure 54: DLC Egress Channel Queue Scheduling

Some minor modifications are introduced to accommodate end-to-end FRF.12 fragments. The user configured FR scheduling class for this SAP will dictate which of the FR SAP MDA queues should be used to queue the fragments. Classes 0 and 1 map all fragments of the FR SAP to the Hi-Priority SAP queue. Classes 2 and 3 map all fragments of the FR SAP to the Low-Priority SAP queue.

The scheduling parameters of these queues have to be modified from existing ones as follows:

- Hi priority (HP) SAP queue: $SAP_HP_Q_PIR = \text{Sum}\{\text{all } SAP_FC_Q_PIR\}$
- Low priority (LP) SAP queue: $SAP_LP_Q_WRR_Weight = \text{Sum}\{\text{all } SAP_FC_Q_CIR\}$