Contents

1 Configuring MMU

1.1 Introduction

1.1.1 Overview

The memory management unit (MMU) is a technology for allocating the chip buffers reasonably. This technology makes devices better deal with all kinds of burst flows.

Flows are not always steady and various burst flows exist in the network. When the network flow is steady and the bandwidth is sufficient, all the data flows are processed properly. When burst flows exist in the network, data flows may be discarded even if the average flow rate does not exceed the bandwidth. Data packets are stored in the buffers of a device before being forwarded. Normally, data packets in the buffers will be forwarded in microseconds. When a burst flow arises, if the instantaneous rate of the burst flow is out of the processing capacity of the device, the data packets that cannot be processed in time are piled up in the buffers and packet loss takes place once the buffers are insufficient. In this case, the MMU technology can used to reasonably configure buffers and allocate different buffer sizes to services, with a view to optimizing the network.

1.1.2 Principles

1. Basic Concepts

- Port group

  All the ports that physically belong to the same switching chip and the same core (some chips are multi-core chips) are called a port group. All devices implement buffer management in port groups.

- Buffer slice

  The buffer of a port group may be divided into multiple buffer blocks, and each buffer block is called a buffer slice. Slice is the minimum unit for independent buffer management. Each buffer slice is managed independently, and a port queue may use one or more buffer slices.

- Guaranteed buffer

  A guaranteed buffer is also called an exclusive buffer. This part of buffer is distributed based on each port queue. The guaranteed buffer of a queue can be used by this queue only. A fixed number of guaranteed buffers are allocated to each queue by default. This part of buffer enables a queue to normally forward packets at the line rate under the stable flow.

- Shared buffer

  In the total buffer of a port group, the remaining part obtained after the guaranteed buffer of each port queue is deducted is the total shared buffer. The shared buffer is available to all the queues. A shared buffer threshold can be set for each queue. This threshold restricts the maximum number of shared buffers that can be used by a queue. When the sum of shared buffers configured for queues in a port group exceeds the total shared buffers of the port group, the "first-come first-served" buffer occupancy mechanism is adopted.

- Service pool

  In addition to the guaranteed buffer, each port queue can use shared buffers. The shared buffers are uniformly stored in the service pool for management. One buffer slice may have multiple service pools. After using its own guaranteed buffer, each queue can request a shared buffer from the service pool within the shared buffer threshold, and the mechanism of "first-request, first-served" is adopted for the service pool. One queue can use the shared buffer of only one service pool according to the mapping configuration.

- Priority group (PG)

  PG is the abbreviation of priority group. For the sake of better and flexible buffer management, a chip offers eight PG objects in the ingress direction to implement independent buffer management. Packets are configured with different priorities (for example, 0 to 7) and can be mapped to different PG objects or one PG so that different buffer allocation schemes are provided for packets with different priorities.

- Egress queues

  Egress queues of a port are classified into unicast queues and multicast queues (the number of queues depends on the product). Logically, a switching chip is divided into the ingress direction and egress direction. The egress queues are in the egress direction of a switching chip. Before packets are sent out of a device, they need to be queued at the egress queues, as shown in Figure 1-1.

Figure 1-1Egress Queues

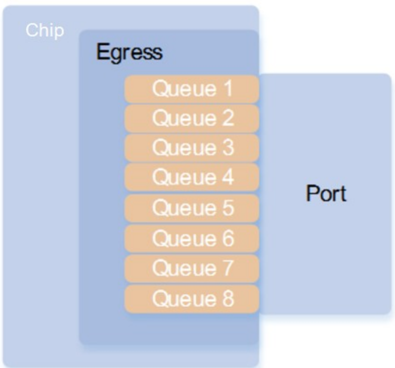

Currently, there are three types of egress queue models:

- ○ There are eight unicast queues and eight multicast queues at the egress. The known unicast packets enter the unicast queues, and all the other packets enter the multicast queues.
- ○ There are eight unicast queues and four multicast queues at the egress. The known unicast packets enter the unicast queues, and all the other packets enter the multicast queues.
- ○ There are only eight queues at the egress, with unicast and multicast queues not differentiated. The queues are managed as unicast queues in the device.

- Global multicast queue

  A global multicast queue refers to the multicast queue shared by all the ports in a port group, and is different from the multicast queue owned by a single physical port.

2. Buffer Adjustment

Users can adjust the buffer allocation of a chip so that the queue of each service has a different buffer size available, services are treated differently, and services of different priorities are provided.

In terms of hardware, the buffer is managed in the input direction and output direction. Figure 1-1 shows the processing mechanism.

Figure 1-1Buffer Processing Mechanism

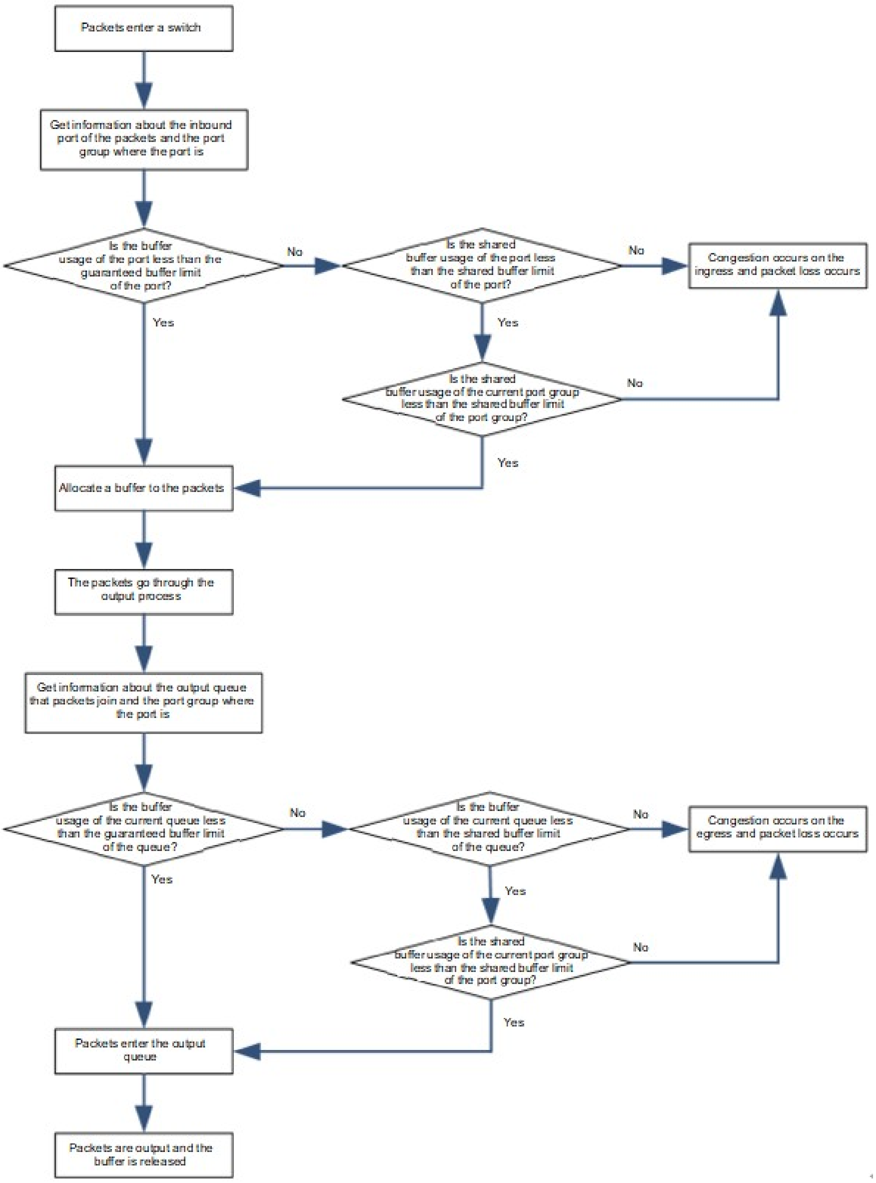

During buffer management, the buffer size is adjusted to the maximum value in the input direction to prevent packet loss in the input direction, and packet loss takes place in the output direction. Therefore, buffer adjustment is not allowed in the input direction, and a CLI command is available for buffer adjustment in the output direction only, including the guaranteed buffer and shared buffer of queues. Buffer adjustment allocates different buffer sizes to queues based on the guaranteed buffer threshold and shared buffer threshold configured for the queues.

3. Status Monitoring

Status monitoring monitors the buffer usage and the queue packet forwarding and discarding by using statistical methods, with a view to providing data support for network optimization and reasonable buffer configuration.

Status monitoring adopts the polling mode to periodically read, from a chip, the buffer usage information and information about packets forwarded/discarded by queues, and analyzes the information so as to display the device operating status.

- Monitoring buffer occupancy in the egress direction

  The information about the guaranteed buffer, shared buffer, and service pool buffer occupancy of queues in the egress direction is regularly collected. The data can be displayed via a CLI command, in MIB mode, or other methods.

- Monitoring packet forwarding and discarding of the egress queues

  The information about the quantities and bytes of packets forwarded and discarded by queues in the egress direction is regularly collected, and the forwarding rate is also calculated. The data can be displayed via a CLI command, in MIB mode, or other methods.

- Sampling interval for status monitoring data

  Status monitoring data is sampled at a fixed interval, which can be configured based on user requirements.

4. Exception Detection

The exception detection function monitors the packet forwarding and discarding of each queue, and displays alarms when the buffer occupancy exceeds the threshold or queue packets are lost.

Exception detection adopts the polling mode to periodically read device buffer occupancy information and the number of packets forwarded by queues, and displays alarms through system logs based on the current configuration. When the buffer utilization of a queue exceeds the threshold, a system log alarm is printed to alert users.

1.2 Configuration Task Summary

MMU configuration includes the following tasks:

- Configuring Buffer Adjustment
- Configuring Status Monitoring
- Configuring Exception Detection

1.3 Configuring Buffer Adjustment

1.3.1 Overview

The buffer adjustment function aims to adjust the objects (such as queues) managed based on chip buffers.

1.3.2 Restrictions and Guidelines

- Configuration on an interface can be made on a physical port only.
- You are advised to configure buffer adjustment when there is no traffic on the device; otherwise the

configuration may cause temporary interruption of service traffic.

1.3.3 Configuration Tasks

Buffer adjustment configuration includes the following tasks:

- Configuring the Global Buffer Mode
- Configuring a Guaranteed Buffer for a Queue
- Configuring a Shared Buffer for a Queue
- Configuring a Flow Control Threshold for a Port

1.3.4 Configuring the Global Buffer Mode

1. Overview

This section describes how to configure the global buffer mode. After configuration, the device configures the default thresholds according to the mode.

2. Restrictions and Guidelines

N/A

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure the global buffer mode.

mmu buffer-mode { **burst-enhance** | **flowctrl-enhance** | **normal** | **qos-enhance** }

The default global buffer mode is **flowctrl-enhance**.

1.3.5 Configuring a Guaranteed Buffer for a Queue

1. Overview

This section describes how to configure a guaranteed buffer for a queue so that the queue uses this part of buffer exclusively.

2. Restrictions and Guidelines

- The **set** *value* parameter indicates the value of guaranteed buffer for a queue. Range: 1-50, unit: cell.
- The *queue-id* <1-8> parameter means that up to 8 queue numbers can be configured. Whether this parameter is supported and the value range of this parameter depend on the actual product version.
- The configuration is displayed when the **show running-config** command is executed even if the default value is used.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure a guaranteed buffer for queues.

mmu queue-guarantee { **output unicast** } [*queue-id*&<1-n>] **set** *value*

No guaranteed buffer is configured for a queue by default.

1.3.6 Configuring a Shared Buffer for a Queue

1. Overview

This section describes how to configure a shared buffer for a queue so as to control the size of the shared buffer that can be used by the queue.

2. Restrictions and Guidelines

- The *queue-id*&<1-8> parameter means that up to 8 queue numbers can be configured.
- The configuration is displayed when the **show running-config** command is executed even if the default value is used.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure a shared buffer for queues.

mmu queue-threshold { **output unicast** } [*queue-id*&<1-n>] **set** *threshold* |

No shared buffer is configured for a queue by default.

1.3.7 Configuring a Flow Control Threshold for a Port

1. Overview

This section describes how to configure a flow control threshold for a port and adjust the flow control threshold of the ingress.

2. Restrictions and Guidelines

- The configuration takes effect only when flow control/priority-based flow control (PFC) is enabled on a port.
- The configuration is displayed when the **show running-config** command is executed even if the default value is used.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure a flow control threshold for the port.

mmu fc-threshold set *threshold*

No flow control threshold is configured for a port by default.

1.4 Configuring Status Monitoring

1.4.1 Overview

This section describes how to collect statistics on and monitor the buffer usage and queue packet forwarding and discarding so as to provide guidance for buffer adjustment configuration.

1.4.2 Restrictions and Guidelines

- The configuration is displayed when the **show running-config** command is executed even if the default value is used.

1.4.3 Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure a sampling period.

mmu sample-period set *period*

No monitoring data sampling period is configured by default.

1.5 Configuring Exception Detection

1.5.1 Overview

The exception detection function is used to analyze the data about device buffer and forwarding status, and generate an alarm when an exception occurs.

1.5.2 Configuration Tasks

Exception detection configuration includes the following tasks:

- Configuring an Alarm Threshold for the Buffer Utilization of a Port Group or Slice
- Configuring an Alarm Threshold for the Buffer Utilization of a Queue
- Configuring a Buffer Utilization Alarm Threshold for a Port

1.5.3 Configuring an Alarm Threshold for the Buffer Utilization of a Port Group or Slice

1. Overview

When the buffer utilization of a port group or slice exceeds the configured value, a system log alarm is printed.

2. Restrictions and Guidelines

- To prevent frequent log refreshing, the alarm logs of the same port group or slice are printed once in 30s at most. The maximum printing interval depends on the configured sampling period.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Configure a buffer utilization alarm threshold.

mmu usage-warn-limit set *threshold*

The default buffer utilization alarm threshold is **0**, indicating that no alarm is reported.

This configuration takes effect for all the port groups or slices.

1.5.4 Configuring an Alarm Threshold for the Buffer Utilization of a Queue

1. Overview

When the buffer utilization of a queue exceeds the configured value, a system log alarm is printed.

2. Restrictions and Guidelines

- The **set** *threshold* parameter indicates the alarm threshold for the buffer utilization of a queue. Range: 1-100, unit: percentage.
- The *queue-id* &<1-8> parameter means that up to 8 queue numbers can be configured. Whether this parameter is supported and the value range of this parameter depend on the actual product version.
- To prevent frequent log refreshing, the alarm logs of the same queue are printed once in 30s at most. The maximum printing interval depends on the configured sampling period.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configure a buffer utilization alarm threshold for a queue.

mmu usage-warn-limit { **unicast** } [*queue-id*&<1-n>] **set** *threshold*

The default buffer utilization alarm threshold is **0**, indicating that no alarm is reported.

1.5.5 Configuring a Buffer Utilization Alarm Threshold for a Port

1. Overview

When the buffer utilization of all the queues on a port exceeds the configured value, a system log alarm is printed.

2. Restrictions and Guidelines

- To prevent frequent log refreshing, the alarm logs of the same port are printed once in 30s at most. The maximum printing interval depends on the configured sampling period.

3. Procedure

(1) Enter the privileged EXEC mode.

enable

(2) Enter the global configuration mode.

configure terminal

(3) Enter the interface configuration mode.

interface *interface-type interface-number*

(4) Configuring a buffer utilization alarm threshold for a port.

mmu usage-warn-limit set *threshold*

The default buffer utilization alarm threshold is **0**, indicating that no alarm is reported.

1.6 Monitoring

Run the **show** commands to check the running status of a configured function to verify the configuration effect.

Run the **clear** commands to clear information.

Caution

Running the **clear** commands may lose vital information and thus interrupt services.

Table 1-1MMU Monitoring

Command	Purpose
show [**mmu**] **queue-counter** [**interface** [*interface-type interface-number*]]	Displays the packet statistics of a queue.
Show queue-buffer [**interface** [*interface-type interface-number*]]	Displays the buffer utilization information of a queue.
show mmu buffer-mode	Displays the current buffer mode.
clear queue-counter [{ **interface** *interface-type interface-number*]	Clears the packet statistics of a queue.
clear [**mmu**] **queue-buffer peaked** [{ **interface** *interface-*	Clears the historical buffer utilization peak

Command	Purpose
type interface-number]	values of a queue.
clear queue-buffer usage-warn-count	Clears the count of buffer alarms of a queue.

1.7 Configuration Examples

1.7.1 Configuring the Application of a Large Buffer Based on Egress Queues

1. Requirements

An enterprise needs a buffer large enough for the net disk service to avoid packet loss occurring in the service flow.

As shown in the following figure, device A is connected to 5 clients and 35 service servers, and 15 service servers are virtualized as 15 front-end servers.

The main service flow is as follows:

- A client sends a request packet to a front-end server.
- The front-end server sends the received request packet to a service server.
- After receiving the request packet, the service server sends a response packet to the front-end server.
- After receiving the response packet, the front-end server sends it to the client.
- The client receives the response packet, indicating that a session is created successfully.
- A many-to-one flow transmission mode exists in this service model:
- The request flows of multiple clients are sent to one front-end server.
- The request flows of multiple front-end servers are sent to one service server.
- The response flows of multiple service servers are sent to one front-end server.
- The response flows of multiple front-end servers are sent to one client.

These flows are basically transmitted through device A, which can easily lead to network congestion. Such a problem can be fixed by configuring a large buffer on the device.

2. Topology

Figure 1-1Application Scenario of a Large Buffer Based on Egress Queues

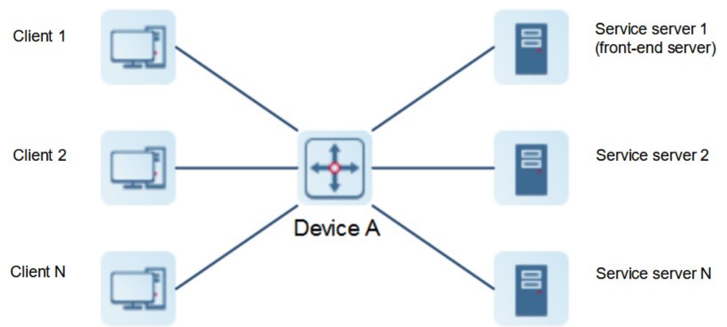

3. Notes

- On all the service ports (namely, the ports connecting to clients and servers), set the shared buffer to 100% for the queues where the service is.
- On all the service ports, set the guaranteed buffer to the minimum size for queues not in use.
- On all the ports not in use, set the guaranteed buffer to the minimum size for all the queues.

4. Procedure

(1) Configure a shared buffer.

On device A, set a shared buffer for queues on all the service ports.

```
DeviceA> enabel
DeviceA# configure terminal
DeviceA(config)# interface range gigabitethernet 0/1-40
DeviceA(config-if-range)# mmu queue-threshold output unicast 1 set 100
DeviceA(config-if-range)# exit
```

(2) Configure a guaranteed buffer.

On device A, set a guaranteed buffer for queues on all the service ports.

```
DeviceA# configure terminal
DeviceA(config)# interface range gigabitethernet 0/1-40
DeviceA(config-if-range)# mmu queue-guarantee output unicast 2 3 4 5 6 7 8 set
1
DeviceA(config-if-range)# exit
```

On device A, set a guaranteed buffer for queues on all the ports not in use.

```
DeviceA(config)# interface range gigabitethernet 0/41-48
DeviceA(config-if-range)# mmu queue-guarantee output unicast set 1
DeviceA(config-if-range)# exit
```

5. Verification

On device A, verify that the command configuration for interface GigabitEthernet 0/1 is correct.

```
DeviceA# configure terminal
DeviceA(config)# interface gigabitethernet 0/1
DeviceA(config-if-GigabitEthernet 0/1)# show this

Building configuration...
!
 mmu queue-threshold output unicast 1 set 100
 mmu queue-guarantee output unicast 2 set 1
 mmu queue-guarantee output unicast 3 set 1
 mmu queue-guarantee output unicast 4 set 1
 mmu queue-guarantee output unicast 5 set 1
 mmu queue-guarantee output unicast 6 set 1
 mmu queue-guarantee output unicast 7 set 1
 mmu queue-guarantee output unicast 8 set 1
```

6. Configuration Files

Device A configuration file

The configuration of interface GigabitEthernet 0/1 to interface GigabitEthernet 0/40 is exactly the same, and therefore only the configuration of GigabitEthernet 0/1 is provided below. The configuration of interface GigabitEthernet 0/41 to interface GigabitEthernet 0/48 is exactly the same, and therefore only the configuration of GigabitEthernet 0/41 is provided below.

```
hostname DeviceA
!
interface GigabitEthernet 0/1
 mmu queue-threshold output unicast 1 set 100
 mmu queue-guarantee output unicast 2 set 1
 mmu queue-guarantee output unicast 3 set 1
 mmu queue-guarantee output unicast 4 set 1
 mmu queue-guarantee output unicast 5 set 1
 mmu queue-guarantee output unicast 6 set 1
 mmu queue-guarantee output unicast 7 set 1
 mmu queue-guarantee output unicast 8 set 1
!
interface gigabitEthernet 0/41
 mmu queue-guarantee output unicast 1 set 1
 mmu queue-guarantee output unicast 2 set 1
 mmu queue-guarantee output unicast 3 set 1
 mmu queue-guarantee output unicast 4 set 1
 mmu queue-guarantee output unicast 5 set 1
 mmu queue-guarantee output unicast 6 set 1
 mmu queue-guarantee output unicast 7 set 1
 mmu queue-guarantee output unicast 8 set 1
!
```