CHAPTER 10

Overhead Accounting

This chapter describes overhead accounting on the Cisco 10000 series router and contains the following topics:

Overhead Accounting Features

Overhead accounting enables the router to account for packet overhead when shaping traffic to a specific rate. This accounting ensures that the router executes quality of service (QoS) features on the actual bandwidth used by subscriber traffic.

The Cisco 10000 series router supports the following overhead accounting features:

- ATM Overhead Accounting, page 10-2
- MLP on LNS with HQoS and ATM Overhead Accounting, page 10-3
- Traffic Shaping Overhead Accounting for ATM, page 10-4
- Ethernet Overhead Accounting, page 10-4

Feature History for Overhead Accounting

Cisco IOS Release	Description	Required PRE[1]
Release 12.3(7)XI7	The ATM Overhead Accounting feature was introduced on the router to enable the router to account for various encapsulation types when applying QoS to packets.	PRE2
Release 12.2(28)SB	The ATM Overhead Accounting feature was introduced in Cisco IOS Release 12.2(28)SB.	PRE2
Release 12.2(31)SB2	The Traffic Shaping Overhead Accounting for ATM feature was introduced on the PRE3.	PRE3
Release 12.2(33)SB	The Ethernet Overhead Accounting feature was introduced on the PRE2, PRE3, and PRE4. ATM overhead accounting was enhanced on the PRE3 to allow a user-defined number of overhead bytes and introduced on the PRE4. The PRE4 inherits all overhead accounting features from the PRE3.	PRE2 PRE3 PRE4
Release 12.2(33)SB2	The MLP on LNS with HQoS and ATM Overhead Accounting feature was introduced on the PRE3 and PRE4.	PRE3 PRE4

1. Performance Routing Engine (PRE)

ATM Overhead Accounting

ATM overhead accounting enables the router to account for various encapsulation types when applying QoS to packets. Typically, in Ethernet Digital Subscriber Line (DSL) environments, the encapsulation from the router to the DSLAM is Gigabit Ethernet and the encapsulation from DSLAM to customer premise equipment (CPE) is ATM. ATM overhead accounting enables the router to account for ATM encapsulation on the subscriber line and for the overhead added by cell segmentation. This accounting enables the service provider to prevent overruns at the subscriber line and ensures that the router executes QoS features on the actual bandwidth used by ATM packets.

The router uses the encapsulation type you configure to calculate the ATM overhead per packet. When calculating ATM overhead at the subscriber line, the router considers the encapsulation used between the router and DSLAM and between the DSLAM and CPE (as described in the following list). You configure the encapsulation type and the router calculates the overhead associated with ATM cell segmentation.

- IEEE 802.1Q and qinq encapsulation are typically used between the router and DSLAM. Because the DSLAM removes the encapsulation, the router does not account for this encapsulation in the calculation.
- The encapsulation used between the DSLAM and the CPE is based on the Subnetwork Access Protocol (SNAP) and multiplexer (MUX) formats of ATM Adaptation Layer 5 (AAL5) and AAL3. These encapsulation types can be routed bridge encapsulation (RBE), PPP over Ethernet (PPPoE), or PPP over ATM (PPPoA), and IP. Because the DSLAM treats IP and PPPoE packets as payload, the router does not account for IP and PPPoE encapsulation in the calculation.

AAL5 segmentation processing adds the additional overhead of the 5-byte cell headers, the AAL5 Common Part Convergence Sublayer (CPCS) padding, and the AAL5 trailer. For more information, see the “Overhead Calculation on the Router” section on page 10-5.

MLP on LNS with HQoS and ATM Overhead Accounting

MLP on LNS (MLPoLNS) sessions are shaped at a negotiated bandwidth that reflects the downstream link bandwidth on the CPE. The link bandwidth on the L2TP network server (LNS) is received on the line either through the connect speed attribute-value pairs (AVP) or PPPoE tags from the L2TP access concentrator (LAC).

In releases earlier than Cisco IOS 12.2(33), hierarchical policies were not supported on MLPoLNS single-member bundles. With MLPoLNS, if LNS sessions do not have a defined bandwidth associated with the session, users must specify the bandwidth of the MLP bundle in the downstream direction (LNS toward the CPE) to avoid drops downstream. The bandwidth of the bundle is the aggregate of all the member links. For single-member MLPoLNS bundles, the bundle bandwidth is the same as the member's link bandwidth. With hierarchical policies, the parent shape value overrides the bandwidth received through the AVP or PPPoE tag.

For more information on the MLP at LNS feature, see the "Configuring Multilink Point-to-Point Protocol Connections" chapter in the *Cisco 10000 Series Router Software Configuration Guide* at the following URL:

http://www.cisco.com/en/US/docs/routers/10000/10008/configuration/guides/broadband/mlp.html

The MLP on LNS with HQoS and ATM Overhead Accounting feature is described in the following sections:

- HQoS, page 10-3
- Overhead Accounting, page 10-3
- Restrictions and Limitations for Overhead Accounting, page 10-7
- Enabling ATM Overhead Accounting on the PRE3 and PRE4 for MLPoLNS, page 10-11

HQoS

The HQoS bandwidth from the parent policy overrides the default bandwidth (based on the rate received on the line) of the bundle. When the parent policy is removed, the default value is restored.

The user's RADIUS environment is responsible for providing the HQoS policy after determining the rate from the connect speed AVP and PPPoE downstream rate tag. The changes assume that platform will be presented the service policies (after the algorithm has run) through the existing API.

Overhead Accounting

With overhead accounting, the downstream transmission rate from the LNS is adjusted to meet the LAC-to-CPE bandwidth. This adjustment accounts account for the difference between the LNS-to-LAC overhead versus the LAC-to-CPE overhead to achieve optimal link utilization for the LAC-to-CPE interface. The overhead differences include IP/UDP/L2TP headers over the L2TP tunnel, as well as the header size and segmentation overhead when the physical interface from the LNS is Gigabit Ethernet and the LAC-to-CPE interface is ATM. Proper accounting can also avoid loss of data from any overruns between the LAC and the CPE.

Traffic Shaping Overhead Accounting for ATM

The Traffic Shaping Overhead Accounting for ATM feature enables the broadband aggregation system (BRAS) to account for various encapsulation types when applying QoS to packets. Typically, in Ethernet Digital Subscriber Line (DSL) environments, the encapsulation from the BRAS to the DSLAM is Gigabit Ethernet and the encapsulation from the DSLAM to the CPE is ATM. ATM overhead accounting enables the BRAS to account for ATM encapsulation on the subscriber line and for the overhead added by cell segmentation. This accounting enables the service provider to prevent overruns at the subscriber line and ensures that the router executes QoS features on the actual bandwidth used by ATM subscriber traffic.

The BRAS uses the encapsulation type you configure for the DSLAM-CPE side to calculate the ATM overhead per packet, except for IP and PPPoE packets. DSLAM-CPE encapsulation types are based on SNAP and MUX formats of ATM Adaptation Layer 5 (AAL5), followed by routed bridge encapsulation (RBE), IP, PPP over Ethernet (PPPoE), or PPP over ATM (PPPoA). Because the DSLAM treats IP and PPPoE packets as payload, the BRAS does not account for these encapsulations.

On the BRAS-DSLAM side, encapsulation is IEEE 802.1Q VLAN or qinq. However, because the DSLAM removes the BRAS-DSLAM encapsulation, the BRAS does not account for 802.1Q or qinq encapsulation.

AAL5 segmentation processing adds the additional overhead of the 5-byte cell headers, the AAL5 Common Part Convergence Sublayer (CPCS) padding, and the AAL5 trailer. For more information, see the "Overhead Calculation on the Router" section on page 10-5.

If the parent policy has overhead accounting enabled, you are not required to explicitly enable accounting on the child policy because, by default, child priority queues that do not contain the **shape** or **bandwidth** command have ATM overhead accounting enabled implicitly.

By default, child priority queues that do not contain the **shape** or **bandwidth** command have ATM overhead accounting enabled implicitly if the parent policy has overhead accounting enabled; you are not required to explicitly enable accounting on the child policy.

If a child traffic class contains the **shape** or **bandwidth** command, you must explicitly enable ATM overhead accounting on the class.

Ethernet Overhead Accounting

The Ethernet Overhead Accounting feature enables the router to account for downstream Ethernet frame headers when applying shaping to packets. A user-defined offset specifies the number of overhead bytes the router is to use when calculating the overhead per packet. Valid offset values are from +63 bytes to -63 bytes of overhead. Before applying shaping, the router calculates the overhead.

Ethernet interfaces and subinterfaces support overhead accounting. Using the **shape** or **bandwidth** command, you can configure accounting per VLAN and per port.

Configuration Commands for Overhead Accounting

To enable overhead accounting, use the **shape** and **bandwidth** commands. These commands allow you to specify the encapsulation type and user-defined offset that the router uses in calculating overhead. The commands have the following syntax:

shape *rate* [**account** {{**qinq** | **dot1q**} {**aal5** | **aal3**} {*subscriber-encap*}} | {**user-defined** *offset* [**atm**]}]

bandwidth {*bandwidth-kbps* | **percent** *percentage* | **remaining percent** *percentage*} {{**qinq** | **dot1q**} {**aal5** | **aal3**} {*subscriber-encap*}} | {**user-defined** *offset* [**atm**]}]

Note The options {{**qinq** | **dot1q**} {**aal5** | **aal3**} {*subscriber-encap*}} and {**user-defined** *offset* [**atm**]} are mutually exclusive.

Subscriber Line Encapsulation Types

The *subscriber-encap* option of the **shape** and **bandwidth** commands specifies the encapsulation type at the subscriber line. The router supports the following subscriber line encapsulation types:

- snap-1483routed
- mux-1483routed
- snap-dot1q-rbe
- mux-dot1q-rbe
- snap-pppoa
- mux-pppoa
- snap-rbe
- mux-rbe

Overhead Calculation on the Router

When calculating overhead for traffic shaping, the router considers the encapsulation type used between the BRAS and the DSLAM, and between the DSLAM and CPE.

Table 10-1 describes the fields the router uses for the various encapsulation types when calculating ATM overhead.

Table 10-1 Overhead Calculation

Encapsulation Type	Number of Bytes	Description
802.1Q	18	6-byte destination MAC address + 6-byte source MAC address + 2-byte protocol ID (0x8100) + 2-byte VID/CFI/PRIORITY + 2-byte length/type
802.3	14	6-byte destination MAC address + 6-byte source MAC address + 2-byte protocol ID (0x8000)
AAL5 MUX plus 1483	8	8-byte AAL5 trailer

Table 10-1 Overhead Calculation (continued)

Encapsulation Type	Number of Bytes	Description
AAL5 MUX plus PPPoA	10	8-byte AAL5 trailer + 2-byte protocol ID (0x0021)
AAL5 SNAP plus 1483	18	8-byte AAL5 trailer + 3-byte LLC header (0xAAAA03) + 3-byte OUI (0x0080c2) + 2-byte protocol ID (0x0007) + 2-byte PAD (0x0000)
AAL5 SNAP plus PPPoA	12	8-byte AAL5 trailer + 3-byte LLC header (0xFEFE03) + 1-byte protocol ID (0xCF)
PPPoE	6	1-byte version/type (0x11) + 1-byte code (0x00) + 2-byte session ID + 2-byte length
qinq	22	6-byte destination MAC address + 6-byte source MAC address + 2-byte protocol ID (0x8100) + 2-byte VID/CFI/PRIORITY + 2-byte protocol ID + 2-byte inner tag + 2-byte length or type

Overhead Accounting and Hierarchical Policies

In hierarchical policies, you can enable overhead accounting for shaping and bandwidth on top-level parent policies, middle-level child policies, and bottom-level child policies. If you enable overhead accounting on a:

- Parent class-default class, you are not required to enable accounting on a child traffic class that does not contain the **bandwidth** or **shape** command.
- Child policy, then you must enable overhead accounting on the parent policy.

The parent and child classes must specify the same encapsulation type when enabling overhead accounting and configuring an offset using the **user-defined** *offset* [**atm**] command option.

Table 10-2 summarizes the configuration requirements for overhead accounting. For example, if overhead accounting is currently enabled for a parent policy, then accounting can be disabled or enabled on a child policy.

Table 10-2 Overhead Accounting Configuration Requirements

Policy Map or Class	Current Configuration	Configuration Requirement
Parent	Enabled	Enabled on child policy
Child	Enabled	Enabled on parent policy
Child class	Enabled	Enabled on all classes in the child policy map, except priority classes with policing
Child class (nonpriority without policing)	Disabled	Disabled on all classes in the child policy map
Child class (priority with policing)	Disabled	Disabled or enabled on all nonpriority classes in the child policy map

Restrictions and Limitations for Overhead Accounting

- You can enable overhead accounting for shaping and bandwidth on top-level parent policies, middle-level child policies, and bottom-level child policies.
- If you enable overhead accounting on a parent policy, you are required to enable accounting on a child policy that is configured with the **shape** or **bandwidth** command. You are not required to enable accounting on a child policy that does not have the **shape** or **bandwidth** command configured.
- If you enable overhead accounting on a child policy, then you must enable overhead accounting on the parent policy.
- In a policy map, you must either enable overhead accounting for all classes in the policy or disable overhead accounting for all classes in the policy. You cannot enable overhead accounting for some classes and disable overhead accounting for other classes in the same policy.
- The router supports overhead accounting only for the **shape** and **bandwidth** commands.
- When you enter the **show policy-map interface** command, the resulting classification byte counts and the queuing feature byte counts do not match. This mismatch occurs because the classification byte count does not consider overhead, whereas the queuing features do consider overhead.
- This feature supports only Ethernet and ATM interfaces.
- Ethernet overhead accounting allows the automatic inclusion of downstream Ethernet frame headers in the shaped rate. However, policing is not supported for Ethernet overhead accounting.
- For the MLPoLNS feature, overhead accounting is supported only on HQoS.
- For MLPoLNS, HQoS with overhead accounting is supported only on single-member bundles and not on multimember bundles.
- QoS restriction on the main interface also apply to single-member MLPoLNS virtual-access bundles (for example, oversubscription of the bundle bandwidth with a parent shaper).
- For MLPoLNS single-member bundles with HQoS, 100 Mbps is the default bundle bandwidth. The bandwidth received on the line (Connect speed of AVP pairs or PPPoE tags) at the LNS overrides this bandwidth. If the connection speed of an AV pair of the MLP bundle is arbitrarily low, overriding with shaper is not possible.
- For the MLPoLNS feature, applying service policies on physical interfaces is not supported. Service policies must be applied on the virtual template of the MLP bundle or from the RADIUS server.

Configuring Overhead Accounting in a Hierarchical Policy

To configure overhead accounting in a hierarchical policy, enter the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# policy-map policy-map-name`	Creates or modifies the child policy. Enters policy-map configuration mode. *policy-map-name* is the name of the child policy map. The name can be a maximum of 40 alphanumeric characters.

	Command	Purpose
Step 2	`Router(config-pmap)# class class-map-name`	Assigns the traffic class you specify to the policy map. Enters policy-map class configuration mode. *class-map-name* is the name of a previously configured class map.
Step 3	`Router(config-pmap-c)# bandwidth {bandwidth-kbps | percent percentage | remaining percent percentage} [account {{qinq | dot1q} {aal5} {subscriber-encap}} | {user-defined offset [atm]}]`	Enables class-based fair queuing and overhead accounting. *bandwidth-kbps* specifies or modifies the minimum bandwidth allocated for a class belonging to a policy map. Valid values are from 8 to 2,488,320, which represents from 1 to 99 percent of the link bandwidth. *percentage* specifies or modifies the maximum percentage of the link bandwidth allocated for a class belonging to a policy map. Valid values are from 1 to 99. **remaining** *percentage* specifies or modifies the minimum percentage of unused link bandwidth allocated for a class belonging to a policy map. Valid values are from 1 to 99. **account** enables ATM overhead accounting. **qinq** specifies queue-in-queue encapsulation as the BRAS-DSLAM encapsulation type. **dot1q** specifies IEEE 802.1Q VLAN encapsulation as the BRAS-DSLAM encapsulation type. **aal5** specifies the ATM Adaptation Layer 5 that supports connection-oriented variable bit rate (VBR) services. *subscriber-encap* specifies the encapsulation type at the subscriber line. For more information, see the "Overhead Accounting and Hierarchical Policies" section on page 10-6. **user-defined** indicates that the router is to use the offset value you specify when calculating ATM overhead. *offset* specifies the number of bytes the router is to use when calculating overhead. Valid values are from -63 to 63 bytes. **atm** applies the ATM cell tax in the ATM overhead calculation. **Note** Configuring both the *offset* and **atm** options adjusts the packet size to the offset size and then adds the ATM cell tax.
Step 4	`Router(config-pmap-c)# exit`	Exits policy-map class configuration mode.
Step 5	`Router(config-pmap)# policy-map policy-map-name`	Creates or modifies the top-level parent policy. *policy-map-name* is the name of the parent policy map. The name can be a maximum of 40 alphanumeric characters.
Step 6	`Router(config-pmap)# class class-default`	Configures or modifies the parent class-default class.

	Command	Purpose
Step 7	`Router(config-pmap-c)# shape [average] rate [account {{qinq | dot1q} {aal5} {subscriber-encap}} | {user-defined offset [atm]}]`	Shapes traffic to the indicated bit rate and enables overhead accounting. (Optional) **average** is the committed burst (Bc) that specifies the maximum number of bits sent out in each interval. This option is only supported on the PRE3. *rate* indicates the bit rate used to shape the traffic, in bits per second. When this command is used with backward explicit congestion notification (BECN) approximation, the bit rate is the upper bound of the range of bit rates that are permitted. **account** enables ATM overhead accounting. **qinq** specifies queue-in-queue encapsulation as the BRAS-DSLAM encapsulation type. **dot1q** specifies IEEE 802.1Q VLAN encapsulation as the BRAS-DSLAM encapsulation type. **aal5** specifies the ATM Adaptation Layer 5 that supports connection-oriented variable bit rate (VBR) services. *subscriber-encap* specifies the encapsulation type at the subscriber line. For more information, see the "Overhead Accounting and Hierarchical Policies" section on page 10-6. **user-defined** indicates that the router is to use the offset value you specify when calculating ATM overhead. *offset* specifies the number of bytes the router is to use when calculating overhead. Valid values are from -63 to +63 bytes. The router configures the offset size if you do not specify the *offset* option. **atm** applies the ATM cell tax in the ATM overhead calculation. **Note** Configuring both the *offset* and **atm** options adjusts the packet size to the offset size and then adds the ATM cell tax.
Step 8	`Router(config-pmap-c)# service-policy policy-map-name`	Applies a child policy to the parent class-default class. *policy-map-name* is the name of a previously configured child policy map. **Note** Do not specify the input or output keywords when applying a child policy to a parent class-default class.

Configuration Examples for Overhead Accounting

This section provides the following configuration examples:

- Enabling ATM Overhead Accounting, page 10-10
- Enabling ATM Overhead Accounting on the PRE3 and PRE4 for MLPoLNS, page 10-11
- Enabling Ethernet Overhead Accounting on the PRE2, page 10-11
- Enabling Ethernet Overhead Accounting on the PRE3 and PRE4, page 10-12

Enabling ATM Overhead Accounting

The following configuration example shows how to enable ATM overhead accounting using a hierarchical policy. The Child policy map has two classes: Business and Nonbusiness. The Business class has priority and is policed at 128,000 kbps. The Nonbusiness class has ATM overhead accounting enabled and has a bandwidth of 20 percent of the available bandwidth. The Parent policy map shapes the aggregate traffic to 256,000 kbps and enables ATM overhead accounting.

Notice that Layer 2 overhead accounting is not explicitly configured for the Business traffic class. If the class-default class of a parent policy has ATM overhead accounting enabled, you are not required to enable ATM overhead accounting on a child traffic class that does not contain the **bandwidth** or **shape** command. Therefore, in this example, the Business priority queue implicitly has ATM overhead accounting enabled because its parent class-default class has overhead accounting enabled.

```
policy-map Child
    class Business
        priority
        police 128000
    class Nonbusiness
        bandwidth percent 20 account dot1q aal5 snap-rbe-dot1q
        exit
    exit
policy-map Parent
    class class-default
        shape 256000 account dot1q aal5 snap-rbe-dot1q
        service-policy Child
```

In the following configuration example, overhead accounting is enabled for bandwidth on the gaming and class-default classes of the child policy map named subscriber_classes, and on the class-default class of the parent policy map named subscriber_line. The voip and video classes do not have accounting explicitly enabled; these classes have ATM overhead accounting implicitly enabled because the parent policy has overhead accounting enabled. Notice that the features in the parent and child policies use the same encapsulation type.

```
policy-map subscriber_classes
    class voip
        priority level 1
        police 8000
    class video
        priority level 2
        police 20
    class gaming
        bandwidth remaining percent 80 account dot1q aal5 snap-rbe-dot1q
    class class-default
        bandwidth remaining percent 20 account dot1q aal5 snap-rbe-dot1q

policy-map subscriber_line
    class class-default
```

```
bandwidth remaining ratio 10 account dot1q aal5 snap-rbe-dot1q
shape average 512 account dot1q aal5 snap-rbe-dot1q
service policy subscriber_classes
```

Note The **shape average** *rate* command is available only on the PRE3 and PRE4. The PRE2 supports the **shape** *rate* command.

Example 10-1 shows that the Child policy map has two classes: Business and Nonbusiness. The Business class has priority and is policed at 128,000 kbps. The Nonbusiness class has ATM overhead accounting enabled and has a bandwidth of 20 percent of the available bandwidth. The Parent policy map shapes the aggregate traffic to 256,000 kbps and enables ATM overhead accounting. Notice that Layer 2 overhead accounting does not occur for the Business traffic class.

Example 10-1 Enabling ATM Overhead Accounting

```
Router(config)# policy-map Child
Router(config-pmap)# class Business
Router(config-pmap-c)# priority
Router(config-pmap-c)# police 128000 //*No Layer 2 overhead accounted*/
Router(config-pmap-c)# class Nonbusiness
Router(config-pmap-c)# bandwidth percent 20 account dot1q aal5 snap-rbe-dot1q
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# policy-map Parent
Router(config-pmap)# class class-default
Router(config-pmap-c)# shape 256000 account dot1q snap-rbe-dot1q
Router(config-pmap-c)# service-policy Child
```

Enabling ATM Overhead Accounting on the PRE3 and PRE4 for MLPoLNS

Example 10-2 shows how to enable ATM overhead accounting using the hierarchical service policy where there is a parent with a child policy. In the example, the child policy map has a predefined class prec2 that shapes the aggregate traffic to 10,000 kbps and enables ATM overhead accounting.

Example 10-2 Enabling ATM Overhead Accounting—MLPoLNS

```
Router(config)# policy-map child
Router(config-pmap)# class prec2
Router(config-pmap-c)# shape average 10000 account user-defined 63 atm
Router(config-pmap-c)# exit
Router(config)# policy-map parent
Router(config-pmap)# class class-default
Router(config-pmap-c)# shape average 256000 account user-defined 63 atm
Router(config-pmap-c)# service-policy child
```

Enabling Ethernet Overhead Accounting on the PRE2

The following configuration example shows how to enable Ethernet overhead accounting on the PRE2 and specify the number of bytes the router should take into account when calculating the overhead. In the example, the policy map named ethernet_ovrh contains the class-default class, which has overhead accounting enabled for shaping and a user-defined offset of 18 bytes specified. The ethernet_ovrh policy map is attached to subinterface Gigabit Ethernet 1/0/0.100.

```
policy-map ethernet_ovrh
   class class-default
```

```
        shape 200 account user-defined 18
!
interface GigabitEthernet1/0/0.100
    encapsulation dot1q 100
    pppoe enable group global
    no snmp trap link-status
    service-policy output ethernet_ovrh
```

Enabling Ethernet Overhead Accounting on the PRE3 and PRE4

The following configuration example shows how to enable Ethernet overhead accounting on the PRE3 or PRE4. In the example, the configuration of the policy map named ethernet_ovrh shapes class-default traffic at a rate of 200,000 kbps and enables overhead accounting with a user-defined value of 18. The ethernet_ovrh policy is attached to subinterface Gigabit Ethernet 1/0/0.100, thereby enabling overhead accounting on the subinterface.

```
Router# configuration-terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# policy-map ethernet_ovrh
Router(config-pmap)# class class-default
Router(config-pmap-c)# shape average 200000 account user-defined 18
!
Router(config)# interface GigabitEthernet1/0/0.100
Router(config-subif)# service-policy output ethernet_ovrh
!
Router# show running-config | begin 1/0/0.100
interface GigabitEthernet1/0/0.100
encapsulation dot1Q 101
pppoe enable group group_pta
service-policy output ethernet_ovrh
```

Verifying Overhead Accounting

To verify overhead accounting, enter any of the following commands in privileged EXEC mode:

Command	Purpose
`Router# show policy-map [interface interface]`	Displays information about the policy map attached to the interface you specify, including ATM overhead accounting. If you do not specify an interface, the command displays information about all of the policy maps configured on the router. **interface** *interface* is the interface type and number (for example, atm 4/0/0). **Note** When you enter the **show policy-map interface** command, the resulting classification byte counts and the queuing feature byte counts do not match. This mismatch occurs because the classification byte count does not consider overhead, whereas the queuing features do consider overhead.

Command	Purpose
Router# **show running-config**	Displays the running configuration on the router. The output shows the AAA setup and the configuration of the policy map, ATM VC, PPPoA, dynamic bandwidth selection, virtual template, and RADIUS server.

Verification Examples for Overhead Accounting

This section provides the following verification examples:

- Verifying ATM Overhead Accounting Using show policy-map, page 10-13
- Verifying Overhead Accounting Using show running-config, page 10-14
- Verifying Ethernet Overhead Accounting with User-Defined Option, page 10-14

Verifying ATM Overhead Accounting Using show policy-map

The following sample output from the **show policy-map** command indicates that ATM overhead accounting is enabled for shaping and disabled for bandwidth:

```
Service-policy output:unit-test

    Class-map: class-default (match-any)
    100 packets, 1000 bytes
    30 second offered rate 800 bps, drop rate 0 bps
    Match: any
    shape (average) cir 154400, bc 7720, be 7720
    target shape rate 154400
    overhead accounting: enabled
    bandwidth 30% (463 kbps)
    overhead accounting: disabled

    queue limit 64 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (packets output/bytes output) 100/1000
```

The following sample output from the **show policy-map** command indicates that ATM overhead accounting is enabled for the class-default class for shaping. The BRAS-DSLAM encapsulation is dot1q and the subscriber line encapsulation is snap-rbe based on the AAL3 service.

```
Policy Map unit-test
    Class class-default
        Average Rate Traffic Shaping
        cir 10% account dot1q aal3 snap-rbe
```

The following sample output from the **show policy-map interface** command indicates that ATM overhead accounting is enabled for shaping and disabled for bandwidth:

```
Service-policy output:unit-test

    Class-map: class-default (match-any)
    100 packets, 1000 bytes
    30 second offered rate 800 bps, drop rate 0 bps
    Match: any
    shape (average) cir 154400, bc 7720, be 7720
    target shape rate 154400
     overhead accounting: enabled
```

```
    bandwidth 30% (463 kbps)
     overhead accounting: disabled

    queue limit 64 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (packets output/bytes output) 100/1000
```

Verifying Overhead Accounting Using show running-config

The following sample output from the **show running-config** command indicates that ATM overhead accounting is enabled for shaping. The BRAS-DSLAM encapsulation is dot1q and the subscriber line encapsulation is snap-rbe based on the AAL5 service.

```
subscriber policy recording rules limit 64
no mpls traffic-eng auto-bw timers frequency 0
call rsvp-sync
!
controller T1 2/0
    framing sf
    linecode ami
!
controller T1 2/1
    framing sf
    linecode ami
!
!
policy-map unit-test
    class class-default
        shape average 10 account dot1q aal5 snap-rbe
```

Note The **shape average** *rate* command is available only on the PRE3 and PRE4. The PRE2 supports the **shape** *rate* command.

Verifying Ethernet Overhead Accounting with User-Defined Option

The following sample output for the policy map named ethernet_ovrh indicates that Ethernet overhead accounting is enabled for shaping and the user-defined offset is 18 bytes. The sample output from the **show policy-map interface** command indicates that the ethernet_ovrh policy map is attached to the subinterface Gigabit Ethernet 1/0/0.100, enabling overhead accounting on the subinterface.

```
Router# show policy-map ethernet_ovrh

Policy Map ethernet_ovrh
Class class-default
Average Rate Traffic Shaping
cir 200000 (bps) account user-defined 18

Router# show policy-map interface GigabitEthernet1/0/0.100

GigabitEthernet1/0/0.100

Service-policy output: ethernet_ovrh

Class-map: class-default (match-any)
0 packets, 0 bytes
30 second offered rate 0 bps, drop rate 0 bps
Match: any
```

```
0 packets, 0 bytes
30 second rate 0 bps
Queueing
queue limit 8 packets
(queue depth/total drops/no-buffer drops) 0/0/0
(pkts output/bytes output) 0/0
shape (average) cir 200000, bc 800, be 800
target shape rate 200000
Overhead Accounting Enabled
```

Related Documentation

This section lists additional Cisco documentation for the features discussed in this chapter. When appropriate, paths to applicable sections are listed below the documentation title.

Feature	Related Documentation
ATM overhead accounting	*Traffic Shaping Overhead Accounting for ATM, Release 12.2(31)SB2* feature module
Class-based shaping	*Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.3* Part 4: Policing and Shaping > Configuring class-Based Shaping
Class maps	*Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2* Part 8: Modular Quality of Service Command-Line Interface > Configuring the Modular Quality of Service Command-Line Interface > Modular QoS CLI Configuration Task List > Creating a Traffic Class
Percentage-based traffic shaping	*QoS: Percentage-Based Shaping, Release 12.2(31)SB2* feature module
Policing	*Comparing Traffic Shaping and Traffic Policing for Bandwidth Limiting*
Policy maps	*Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2* Part 8: Modular Quality of Service Command-Line Interface > Configuring the Modular Quality of Service Command-Line Interface > Modular QoS CLI Configuration Task List > Creating a Traffic Policy
QoS service policies	*QoS Configuration and Monitoring, Creating Time-of-Day QoS Service Policies* Tech Note *QoS Configuration and Monitoring, Monitoring Voice over IP Quality of Service* Tech Note *Site-to-Site MPLS VPN Solution for Service Providers, Service Provider Quality-of-Service Overview* Tech Note
Traffic shaping	*Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.3* Part 4: Policing and Shaping
MLP on LNS	*Cisco 10000 Series Router Software Configuration Guide* Chapter 19: Configuring Multilink Point-to-Point Protocol Connections