CHAPTER 2

# Scalability and Performance

The infrastructure of the service provider must be capable of supporting the services the enterprise customer or Internet service provider (ISP) wants to offer its subscribers. It must also be able to scale to an expanding subscriber base. You can configure the Cisco 10000 series router for high scalability.

This chapter discusses the following topics:



## Line Card VC Limitations

The Cisco 10000 series router supports four ATM service categories for virtual circuits (VCs):

- Constant Bit Rate (CBR)
- Variable Bit Rate-nonreal-time (VBR-nrt)
- Unspecified Bit Rate (UBR) with a peak cell rate (PCR), referred to as shaped UBR
- UBR without a PCR, referred to as unshaped UBR

The segmentation and reassembly (SAR) mechanism configures priority and additional traffic management parameters for the various ATM service categories. Table 2-1 lists the priority levels the SAR sets for the service categories.

*Table 2-1 ATM Service Categories*

| Parameter | CBR | VBR-rt | VBR-nrt | Shaped UBR | Unshaped UBR |
|---|---|---|---|---|---|
| Priority | 0 | 1 | 2 | 3 | None |

The number of SAR priority levels and the service categories supported at each priority level vary from line card to line card. For example, the 1-port OC-12/STM-1 line card supports the four levels of priority and the service categories listed in Table 2-2, but the 4-port OC-3 line card supports only two levels of priority and the service categories listed in the table.

The ATM line cards support a maximum number of VCs per priority. That VC limit depends on the VC limit of the SAR (SAR limit) and the number of priority levels configured. Table 2-2 describes how to determine the VC limit per priority level per port for the specified line cards.

*Table 2-2 Maximum Number of VCs per Priority*

| ATM Line Card | SAR Priority Levels | VC Rate | Maximum Number of VCs per Priority |
|---|---|---|---|
| 1-Port OC-12/ STM-1 | 0 = CBR VCs<br>1 = VBR-rt VCs<br>2 = VBR-nrt VCs<br>3 = UBR VCs | Full line rate | SAR limit / 2 / number of priority levels<br>4 priority system:<br>65,536 / 2 / 4 = 8192 VCs per priority level |
| | | Half line rate and below | SAR limit / number of priority levels<br>4 priority system:<br>65,536 / 4 = 16,384 VCs per priority level |
| 4-Port OC-3 | 0 = CBR, VBR-nrt VCs<br>1 = UBR VCs | Half line rate and below | SAR limit / number of PHYs / number of priority levels<br>2 priority system:<br>65,536 / 4 / 2 = 8192 VCs per priority level per port |
| 8-Port E3/DS3 | 0 = CBR VCs<br>0 = VBR-nrt VCs<br>1 = UBR VCs | Half line rate and below | SAR limit / number of PHYs / number of priority levels<br>2 priority system:<br>65,536 / 8 / 2 = 4096 VCs per priority level per port |

Configuring more channels or VCs than there are available priority locations can cause random channels or VCs to get stuck in the SAR. This occurs when an active channel tries to reschedule itself, but no priority locations are available. Therefore, the channel cannot find a place to reschedule itself, which results in a lost event for the channel, and the channel becomes stuck in the SAR.

On the PRE2, when a VC becomes stuck in the SAR, the PRE2 scheduler stops forwarding traffic on only the VC that is stuck in the SAR; the other VCs still carry traffic. On the PRE3, the PRE3 scheduler stops forwarding traffic on all the VCs configured on that ATM line card.

For example, suppose a 1-port OC-12 line card at full line rate is configured for four levels of priority and a 4-port OC-3 line card at half line rate is configured for two levels of priority. By calculating the maximum number of VCs as described in Table 2-2, you can configure 8192 VCs per priority level for

the 1-port OC-12 and 8192 VCs per priority level per port for the 4-port OC-3—a total of 16,384 VCs per priority level per port. If the number of VCs you configure exceeds the VC limit, the VCs get stuck in the SAR.

# Limitations and Restrictions

The Cisco 10000 series router has the following limitations and restrictions for scalability and performance:

- When Layer 4 Redirect (L4R) service is applied without Port Bundle Host Key (PBHK) service, the translations are all done in the PXF, except for those translations that encounter a collision condition. A collision occurs when a subscriber has two simultaneous TCP connections whose source ports have the same Modulo 64 result.

  For example, the subscriber has an active TCP connection on source port 1026, and while this connection is still alive the subscriber starts another TCP connection on source port 1090. A collision is created because the Modulo 64 result for both the source ports (1024 and 1090) is 2. In this example, L4R translation for the first traffic stream is done in the PXF and for the second TCP stream the packets are sent to the route processor (RP) where the L4R translation is done. This seperation prevents collisions.

- When the PBHK service is applied with L4R service, certain restrictions apply:
  - When the destination IP in any one of the access control entries of the PBHK ACL matches the redirected server IP address, then both L4R and PBHK translations are done in the RP.
  - When the destination IP address in the access control entries of the PBHK ACL does not match the redirect server IP address, then L4R translations are done in the PXF, and the packets that match the PBHK ACL are translated in the RP.

  For configuration examples, see the "Layer 4 Redirect Scaling" section on page 2-5.

- Certain restrictions apply on L4R translations for IP subnet sessions. If two subscribers send TCP traffic using the same source port, then L4R translation for the common port is done in the RP. However, if a group of IP subscribers in an IP subnet session send traffic on different source ports then L4R translations for all the subscribers are done in the PXF.

- For permanent L4R service, you can scale up to the number of sessions listed in Table 2-3. Scaling beyond these sessions can lead to an increase in CPU usage that is beyond the recommended limits.

***Table 2-3 Scaling Limit of L4R Sessions***

| Cisco IOS Release | PRE2 | PRE3 | PRE4 |
|---|---|---|---|
| 12.2(31)SB | 4000 | 4000 | — |
| 12.2(33)SB | 4000 | 16000 | 16000 |

- You can apply access control lists (ACLs) to virtual access interfaces (VAIs) by configuring them under virtual template interfaces. You can also configure ACLs by using RADIUS attribute 11 or 242. Prior to Cisco IOS Release 12.2(28)SB, when you used attribute 242, a maximum of 30,000 sessions could have ACLs; this restriction was removed in release 12.2(28)SB and subsequent releases.

- For PRE2, the Cisco 10000 series router supports mini-ACLs (eight or fewer access control entries) and turbo ACLs (more than eight access control entries) for non-SSG interfaces. The limit for mini-ACLs is 32,000. The limit for turbo ACLs depends on the complexity of the defined ACLs. For PRE3, the Cisco 10000 series router does not use mini-ACLs.

- For SSG (RADIUS) configurations on PRE2, the following limitations apply:
  - For Cisco IOS Release 12.3(7)XI, ACLs defined through SSG configuration (RADIUS) are restricted to mini-ACLs only. Turbo ACLs cannot be used in combination with SSG and RADIUS. If you apply a Turbo ACL to an SSG session, the following syslog error is generated: "%C10K_ACLS-3-SSG_TURBO_ACL: acl is a Turbo ACL and cannot be used for SSG."

**Note** If a mini-ACL is on the verge of becoming a turbo ACL (that is, the ACL contains eight access control entries), SSG redirection can cause the mini-ACL to become a turbo ACL. For Cisco IOS Release 12.3(7)XI, this change would also cause a syslog error to be generated as follows: "%C10K_ACLS-3-SSG_ACL_ERR: acl is miniACL but cannot have another punt rule added."

- The Cisco 10000 series router supports a maximum of 2,000 authentication, authorization, and accounting (AAA) method lists. If you configure more than 2,000 AAA method lists by using the **aaa authentication ppp** or **aaa authorization network** command, traceback messages appear on the console.
- To avoid CPU overload and router instability, use the **logging rate-limit** command to limit the rate that the Cisco 10000 series router logs system messages. For more information, see the **logging rate-limit** command in the *Cisco IOS Configuration Fundamentals and Network Management Command Reference, Release 12.3*, located at the following URL:

  http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_book09186a008017d0a2.html
- The Cisco 10000 series router high-speed interfaces work efficiently to spread traffic flows equally over the queues. However, using single traffic streams in a laboratory environment might result in less-than-expected performance. To ensure accurate test results, test the throughput of the Gigabit Ethernet, OC-48 POS, or ATM uplink with multiple source or destination addresses. To determine if traffic is being properly distributed, use the **show pxf cpu queue** command.
- The Cisco 10000 series router supports a configuration file of up to 16 megabytes.
- If you configure create on demand PVCs (individual and within a range) and PPP sessions, RP CPU utilization can be extremely high when bringing up and tearing down sessions and PVCs. This usage is a concern only when the configuration contains approximately 30,000 PPP sessions, and additional services are enabled (such as DBS, ACLs, and service policies).

  To reduce the RP CPU usage for PPPoA sessions, reduce the number of configured PVCs in a single subinterface. To reduce the RP CPU usage for PPPoEoA sessions, use call admission control (**call admission limit** command).

# Scaling Enhancements in Cisco IOS Release 12.2(33)XNE

Starting from Cisco IOS Release 12.2(33)XNE, the **microcode reload pxf** command has been made for general availability. When this command is executed in a scalable scenario, CPUHOG messages may appear as the IOS software populates the parallel express forwarding (PXF) plane with the required information to resume forwarding of traffic as soon as possible. If there is lot of information to be populated, especially when the configuration is scaled up, CPUHOG messages may not appear till all the information is populated.

# Scaling Enhancements in Cisco IOS Release 12.2(33)SB

Cisco IOS Release 12.2(33)SB provides increased scalability for the Layer 4 Redirect feature.

## Layer 4 Redirect Scaling

The Layer 4 Redirect feature allows redirection of users' TCP or UDP traffic to a server to control and increase performance. In Cisco IOS Release12.2(33)SB, the ISG L4R feature is implemented in the PXF. This design increases the number of redirects to provide higher scalability and performance. This enhancement is a scalable solution for portals and self-provisioning and is supported on PRE3 and PRE4 only. On a PRE2 L4R translations are done in the RP.

PBHK translations are always done in the RP. The L4R feature is scalable when applied alone; however, certain scalability restrictions apply when it is used with PBHK. See also the "Limitations and Restrictions" section on page 2-3.

In Example 2-1, when the destination IP used in the PBHK ACL (162) matches the redirected server IP address, L4R translations are done in the RP.

***Example 2-1 L4R Translations in the Route Processor***

```
class-map type traffic match-any class-l4r
match access-group input 152

policy-map type service ser-l4r
class type traffic class-l4r
redirect to ip 200.0.0.2

ip portbundle
match access-list 162
source loopback 1

access-list 152 deny tcp any host 200.0.0.2
access-list 152 permit tcp any any

access-list 162 permit tcp any host 200.0.0.2
```

In Example 2-2, when the destination IP used in the PBHK ACL (162) is not the same as the redirected server IP address, L4R translations are done in the PXF.

***Example 2-2 L4R Translations in PXF***

```
class-map type traffic match-any class-l4r
match access-group input 152

policy-map type service ser-l4r
class type traffic class-l4r
redirect to ip 210.0.0.2

ip portbundle
match access-list 162
source loopback 1

access-list 152 deny tcp any host 200.0.0.2
access-list 152 permit tcp any any

access-list 162 permit tcp any host 200.0.0.2
```

For more information on configuring L4R, see the "Redirecting Subscriber Traffic Using ISG Layer 4 Redirect" chapter in the *Cisco IOS Intelligent Service Gateway Configuration Guide, Release 12.2 SB* at the following URL:

http://www.cisco.com/en/US/products/ps6566/products_configuration_guide_chapter09186a0080630d65.html#wp1048970

For more information on configuring PBHK, see the "Configuring ISG Port-Bundle Host Key" chapter in the *Cisco IOS Intelligent Service Gateway Configuration Guide, Release 12.2 SB* at the following URL:

http://www.cisco.com/en/US/products/ps6566/products_configuration_guide_chapter09186a0080630d6c.html

# Scaling Enhancements in Cisco IOS Release 12.3(7)XI1

Cisco IOS Release 12.3(7)XI1 provides increased limits with FIB scaling, policy-map scaling, and queue scaling.

## FIB Scaling

The FIB is a routing table that is used to look up the next hop route for the destination IP address and the reverse path forwarding (RPF) route using the source IP address. The FIB Scaling feature implements the following changes:

- Up to 1 million routes in the global FIB table are supported without MPLS VPN configuration.
- Total number of virtual routing and forwarding instances (VRFs) supported is 4095.
  - Up to 100 routes per VRF with 4095 VRFs configured.
  - Up to 70 routes per VRF with 4095 VRFs configured, plus 200,000 global BGP routes.
  - Up to 600 routes per VRF with 1000 or fewer VRFs configured.

## Policy-Map Scaling

The Policy-Map Scaling feature increases the system-wide number of quality of service (QoS) policy maps that you can configure. Depending on the complexity of your configuration, the Cisco 10000 series router supports up to 4096 policy maps. In complex configurations the maximum number of policy maps can be as small as a few hundred. Additionally, when you use percent-based policing in a service policy, the system may convert a single customer-configured service to multiple service policies (which count against the 4096 limit). The system uses one such service policy for each different speed interface that uses a service policy with percent-based policing

Each **policy-map** command counts as one policy map and applying the same policy map on different speed interfaces also counts as an extra policy map. The **policy-map** command syntax is unchanged. The maximum number of classes that you can configure in a policy is 127.

## Queue Scaling

The Queue Scaling feature increases the total number of queues that VTMS supports to 131,072. Of the total number, 254 queues are available for high speed interfaces, and 130,816 queues are available for low speed interfaces. This increase allows the support of the 31,500 priority queues (of 131,072 total queues) on 31,500 sessions or interfaces.

Each interface includes a class-default queue and a system queue. If you attach an output policy map with 1 priority queue and 1 class-based weighted fair queue (PQ/CBWFQ) to each of the 31,500 interfaces, the number of priority queues is 31,500 and the total number of queues is 31,500 x 4, or 126,000 queues.

The maximum number of queues per link remains at 32, of which 29 are user-configurable because there is 1 class-default queue, 1 system queue, and 1 reserved queue.

To support 131,072 queues, the queue limits range has changed. For high-speed interfaces (an interface that has a speed greater than 622 Mbps), the queue limit range is 128 to 65,536. For low-speed interfaces the queue limit range is 8 to 4,096. Because the total number of packet buffers for queue limits is 4,194,304, the average queue depth is less than or equal to 32 per queue with 131,072 queues configured.

On low-speed interfaces, the default queue size is 8 for all QoS CBWFQ queues, with the exception of WRED queues. The default queue size for WRED queues is 32.

The class-default queue size on low-speed interfaces has changed from 32 to 8. If the traffic is too bursty and packets drop, you can use the **queue-limit** command to increase the class-default queue size.

If you change the queue size for 131,072 queues while traffic is running, the queue size for a few queues might not be changed if packets were in the queues. An “out of resource” message can also appear. Use the **queue-limit** command to modify the queue size for those queues that were not changed.

The queue limits packet buffers can become fragmented after the queue sizes on 131,072 queues has been changed a few times. The system might indicate that there are not enough resources to increase queue size, even though there are enough free packet buffers. Removing and reapplying the policy map on the interfaces solves this problem.

Use the **show pxf cpu queue summary** command to see the number of packet buffers, packet buffers being recycled, and free packet buffers.

# Scaling Enhancements in Cisco IOS Release 12.3(7)XI2

Cisco IOS Release 12.3(7)XI2 provides increased limits with queue scaling and VC scaling.

## Queue Scaling

At least two queues are allocated for every interface or subinterface for which separate queues are created. The first queue is the default queue for normal traffic, and the second queue, known as the system queue, is used for a small amount of router-generated traffic that bypasses the normal drop mechanisms. For 32,000 VCs, this setup would require the allocation of a minimum of 64,000 queues. While Cisco IOS Release 12.3(7)XI1 adds support for up to 128,000 queues, a more effective use of these limited resources is realized by having the subinterfaces on a given main interface share the single system queue of the main interface.

In Cisco IOS Release 12.3(7)XI2, the subinterfaces on a given main interface share the single system queue of the main interface, which allows for 32,000 subinterfaces with a three-queue model that supports assured forwarding (AF) queues and expedited forwarding (EF) queues, in addition to the default best effort (BE) queues. Because a system queue does not exist for every subinterface, this setup frees up queues for a 4-queue model.

## VC Scaling

When configured for hierarchical shaping, ATM line cards support the following number of VCs:

- E3/DS3 line card supports a maximum of 4,096 VCs
- OC-12 ATM line card supports a maximum of 16,384 VCs (previously 14,436)
- OC-3 ATM line card supports a maximum of 8,191 VCs

# Scaling Enhancements in Cisco IOS Release 12.2(28)SB

In Cisco IOS Release 12.2(28)SB, up to 16,384 L2TP tunnels are supported. Because of a limit on the number of VPDN groups supported, it is not possible to configure 16,384 tunnel definitions using the CLI. Configure the remaining tunnel definitions using RADIUS.

# Configuring the Cisco 10000 Series Router for High Scalability

To ensure high scalability on the Cisco 10000 series router, perform the following configuration tasks:

# Configuring Parameters for RADIUS Authentication

If your network uses a RADIUS server for authentication, set the small, middle, and big buffers by using the **buffers** command. Table 2-4 lists the buffer sizes to configure (and see Example 2-3).

*Table 2-4 Buffer Sizes for RADIUS Authentication*

| Buffer | Size |
|---|---|
| Small | 15000 |
| Middle | 12000 |
| Big | 8000 |

*Example 2-3 Configuring Buffer Sizes*

```
Router(config)# buffers small perm 15000
Router(config)# buffers mid perm 12000
Router(config)# buffers big perm 8000
```

Typically, if the RADIUS server is only a few hops away from the router, we recommend that you configure the RADIUS server retransmit and timeout rates by using the **radius-server** command. Table 2-5 lists the recommended settings (and see Example 2-4).

*Table 2-5 RADIUS Server Parameters*

| Parameter | Value |
|---|---|
| RADIUS Server Retransmit Rate | 5 |
| RADIUS Server Timeout Rate | 15 |

*Example 2-4 Configuring RADIUS Server Parameters*

```
Router(config)# radius-server retransmit 5
Router(config)# radius-server timeout 15
```

# Configuring L2TP Tunnel Settings

Configure an L2TP tunnel password using Cisco IOS Release 12.2(4)BZ1 or later. We recommend that you configure the L2TP tunnel parameters listed in Table 2-6 (and see Example 2-5, Example 2-6, and Example 2-7).

*Table 2-6 L2TP Tunnel Settings*

| Parameter | Setting |
|---|---|
| No Session Timeout | 30 |
| L2TP Tunnel Receive Window | 100 |
| L2TP Tunnel Retransmit Timeout | 2 (minimum)<br>8 (maximum) |

**Note** The No Session Timeout parameter indicates the length of time a tunnel persists when there are no sessions in the tunnel.

***Example 2-5 Configuring an L2TP Tunnel Password***

```
Router(config)# vpdn-group tunnel1
Router(config-if)# l2tp tunnel password 7
```

***Example 2-6 Configuring the No Session Timeout Parameter***

```
Router(config)# vpdn-group tunnel1
Router(config-if)# l2tp tunnel nosession-timeout 30
```

***Example 2-7 Configuring the L2TP Tunnel Receive-Window and Retransmit Timeout Parameters***

```
Router(config)# vpdn-group tunnel1
Router(config-if)# l2tp tunnel receive-window 100
Router(config-if)# l2tp tunnel retransmit timeout min 2
Router(config-if)# l2tp tunnel retransmit timeout max 8
```

# VPDN Group Session Limiting

Before the introduction of the VPDN Group Session Limiting feature introduced in Cisco IOS software release 12.2(1)DX, you could only globally limit the number of VPDN sessions on a router with limits applied equally to all VPDN groups. Using the VPDN Group Session Limiting feature, you can limit the number of VPDN sessions allowed per VPDN group. For more information, see the VPDN Group Session Limiting feature documentation, located at the following URL:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5012/products_feature_guide09186a0080087ef2.html

# Configuring the PPP Authentication Timeout

To keep the L2TP network server (LNS) from timing out a PPP authentication process, set the PPP Timeout parameter to 100, using the **ppp timeout authentication** command (Example 2-8).

***Example 2-8 Configuring the PPP Authentication Timeout***

```
Router(config)# interface Virtual-Template1
Router(config-if)# ppp timeout authentication 100
```

# Disabling Cisco Discovery Protocol

To maximize scalability, do not enable the Cisco Discovery Protocol (CDP).

**Note** CDP is disabled by default.

# Disabling Gratuitous ARP Requests

To maximize the performance of the router, disable gratuitous ARP requests, using the **no ip gratuitous-arp** command (Example 2-9).

***Example 2-9 Disabling Gratuitous ARP Requests***

```
Router(config)# no ip gratuitous-arp
```

# Configuring a Virtual Template Without Interface-Specific Commands

If you configure a virtual template interface with interface-specific commands, the Cisco 10000 series router does not achieve the highest possible scaling. To verify that the router does not have interface-specific commands within the virtual template interface configuration, use the **test virtual-template** <*number*> **subinterface** command.

Including interface-specific commands in a virtual template can limit PPP session scaling. Table 2-7 lists the interface-specific commands that prevent the Cisco 10000 series router from attaining the highest possible PPP session scaling.

***Table 2-7 Interface-Specific Commands That Prevent PPP Scaling***

| Command | Function |
|---|---|
| access-expression | Builds a bridge Boolean access expression. |
| asp | Asynchronous Port (ASP) subcommands. |
| autodetect | Autodetects encapsulations on serial interfaces. |
| bridge-group | Transparent bridging interface parameters. |
| bsc | Binary Synchronous Communications (BSC) interface subcommands. |
| bstun | Block Serial Tunnel (BSTUN) interface subcommands. |
| carrier-delay | Specifies delay for interface transitions. |
| cdp | Cisco Discovery Protocol (CDP) interface subcommands. |
| clock | Configures the serial interface clock. |
| compress | Sets the serial interface for compression. |
| custom-queue-list | Assigns a custom queue list to an interface. |
| diffserv | Differentiated Services (diffserv) for provisioning. |
| down-when-looped | Forces a looped serial interface down. |
| encapsulation | Sets the encapsulation type for an interface. |
| fair-queue | Enables fair queuing on an interface. |
| full-duplex | Configures full-duplex operational mode. |
| h323-gateway | Configures the H.323 Gateway. |
| half-duplex | Configures half-duplex and related commands. |

*Table 2-7 Interface-Specific Commands That Prevent PPP Scaling (continued)*

| Command | Function |
|---|---|
| help | Provides a description of the interactive help system. |
| hold-queue | Sets the hold queue depth. |
| lan-name | Specifies a name for the LAN that is attached to the interface. |
| lapb | X.25 Level 2 parameters (Link Access Procedure, Balanced). |
| load-interval | Specifies the interval for load calculation for an interface. |
| locaddr-priority | Assigns a priority group. |
| logging | Configures logging for an interface. |
| loopback | Configures the internal loopback on an interface. |
| mac-address | Manually sets the MAC address for an interface. |
| max-reserved-bandwidth | Specifies the maximum reservable bandwidth on an interface. |
| mpoa | Multiprotocol over ATM (MPOA) interface configuration commands. |
| multilink | Configures multilink parameters. |
| multilink-group | Puts the interface in a multilink bundle. |
| netbios | Defines Network Basic Input/Output System (NetBIOS) access list or enables name-caching. |
| ntp | Configures the Network Time Protocol (NTP). |
| priority-group | Assigns a priority group to an interface. |
| qos pre-classify | Enables quality of service (QoS) preclassification. |
| random-detect | Enables weighted random early detection (WRED) on an interface. |
| roles | Specifies roles (by entering roles mode). |
| sap-priority | Assigns a priority group. |
| sdlc | Configures Synchronous Data Link Control (SDLC) to Logical Link Control type 2 (LLC2) translation. |
| serial | Serial interface commands. |

*Table 2-7 Interface-Specific Commands That Prevent PPP Scaling (continued)*

| Command | Function |
|---|---|
| snmp | Modifies Simple Network Management Protocol (SNMP) interface parameters. |
| source | Gets the configuration from another source. |
| stun | Serial Tunnel (STUN) interface subcommands. |
| transmit-interface | Assigns a transmit interface to a receive-only interface. |
| trunk-group | Configures an interface to be in a trunk group. |
| tx-ring-limit | Limits the number of particles or packets that can be used on a transmission ring on an interface. |

In Example 2-10, the output of the **test virtual-template** <*number*> **subinterface** command indicates that the interface-specific command **carrier-delay** is set.

***Example 2-10 Verifying Interface-Specific Commands in the Virtual Template***

```
Router(config)# test virtual-template 11 subinterface

Subinterfaces cannot be created using Virtual-Template11
Interface specific commands:
carrier-delay 45
```

# Monitoring PPP Sessions Using the SNMP Management Tools

To prevent the virtual-access subinterfaces from being registered with the SNMP functionality of the router and using memory, do not use the router's SNMP management tools to monitor PPP sessions. Use the **no virtual-template snmp** command to disable the SNMP management tools (Example 2-11).

***Example 2-11 Preventing SNMP Registration of Virtual-Access Subinterfaces***

```
Router(config)# no virtual-template snmp
```

# SNMP Process and High CPU Utilization

Network management applications retrieve information from devices by using SNMP. If a user application polls the SNMP MIBs while the router is updating its routing table, the SNMP engine process can cause CPU HOG messages to appear and sessions and tunnels to go down until the process releases the CPU.

For information about how to avoid high CPU utilization by an SNMP process, see the *IP Simple Network Management Protocol (SNMP) Causes High CPU Utilization* Tech Note, located at the following URL:

http://www.cisco.com/warp/public/477/SNMP/ipsnmphighcpu.shtml#polling

# CISCO-ATM-PVCTRAP-EXTN-MIB

The Cisco 10000 series router does not support the CISCO-ATM-PVCTRAP-EXTN-MIB for large numbers of permanent virtual circuits (for example, 32,000 PVCs). To exclude the Cisco-ATM-PVCTRAP-EXTN-MIB from the Simple Network Management Protocol (SNMP) view and enhance scalability, configure the following commands in global configuration mode:

| | Command | Purpose |
|---|---|---|
| **Step 1** | `Router(config)# snmp-server view` *view-name oid-tree* **included** | Creates or updates a view entry.<br>The *view-name* argument is a label for the view record that you are updating or creating. The name is used to reference the record.<br>The *oid-tree* argument is the object identifier of the ASN.1 subtree to be included from the view. Specify a valid oid-tree from where you want to poll the information.<br>The **included** argument configures the OID (and subtree OIDs) specified in the *oid-tree* argument to be included in the SNMP view. |
| **Step 2** | `Router(config)# snmp-server view` *view-name* **ciscoAtmPvcTrapExtnMIB excluded** | Configures the CISCO-ATM-PVCTAP-EXTN-MIB OID (and subtree OIDs) to be explicitly excluded from the SNMP view. You must specify the oid-tree as shown in the command line.<br>The *view-name* argument must match the *view-name* you specified in step 1. |
| **Step 3** | `Router(config)# snmp-server community` *string* [**view** *view-name*] [**ro** \| **rw**] [access-list-number] | Sets up the community access string to permit access to SNMP.<br>The *string* argument is a community string that acts like a password and permits access to the SNMP protocol.<br>The *view-name* argument must match the *view-name* you specified in step 1. |

Example 2-12 shows how to create or modify the SNMP view named *myview* to include the information polled from the Internet oid-tree and to exclude the CISCO-ATM-PVCTRAP-EXTN-MIB oid-tree. The community access string named private is set up and access to SNMP is read-only (**ro**) access.

***Example 2-12 Excluding CISCO-ATM-PVCTRAP-EXTN-MIB from the SNMP View***

```
Router(config)# snmp-server view myview internet included
Router(config)# snmp-server view myview ciscoAtmPvcTrapExtnMIB excluded
Router(config)# snmp-server community private view myview ro
```

For more information about the **snmp-server view** and **snmp-server community** commands, see the *Cisco IOS Configuration Fundamentals and Network Management Command Reference, Release 12.3*, located at the following URL:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_book09186a008017d0a2.html

# Configuring the Trunk Interface Input Hold Queue

To ensure high scalability, set the trunk interface input hold queue to a high value (Example 2-13).

**Note** The default value for the OC-12 ATM line card trunk interface input hold queue is 27230. Cisco laboratory tests have shown this setting to result in the highest scalability for the OC-12 ATM line card. We recommend that you not change the default setting.

***Example 2-13 Setting the Trunk Interface Input Hold Queue***

```
Router(config)# interface gig1/0/0
Router(config-if)# hold-queue 4096 in
```

# Configuring no atm pxf queuing

**Note** We do not recommend using this mode for QoS-sensitive deployments.

Configuring the **no atm pxf queuing** command on each port of the Cisco 10000 series router enables the router to support a high number of VCs. PPPoA supports one session per VC and requires that you enable **no atm pxf queuing** to support 32,000 PPPoA sessions. Enabling **no atm pxf queuing** is not required for L2TP, and might not be required for PPPoE, because you can have 32,000 sessions on a single VC.

The Cisco 10000 series router supports three ATM traffic classes when you configure **no atm pxf queuing**: unshaped UBR (no PCR is specified), shaped UBR (PCR is specified), and VBR-nrt. To configure an unspecified bit rate (UBR) quality of service (QoS) and specify the output peak cell rate (PCR), use the **ubr** command in the appropriate configuration mode. In ATM VC configuration mode, the syntax is:

```
Router(config-if-atm-vc)# ubr output-pcr
```

If you do not specify a PCR, unshaped UBR is configured.

To configure the variable bit rate-nonreal-time (VBR-nrt) QoS, use the **vbr-nrt** command in the appropriate configuration mode and specify the output PCR, output sustainable cell rate (SCR), and the output maximum burst cell size (MBS) for a VC class. Note that if the PCR and SCR values are equal, the MBS value is 1.

```
output-pcr output-scr output-mbs
```

**Note** Before you configure VCs on an interface, configure the **atm pxf queuing** mode for the port (**atm pxf queuing** or **no atm pxf queuing**). After you configure the mode, then configure the VCs. Do not change the mode while VCs are configured on the interface. If you need to change the mode, delete the VCs first and then change the mode. Changing the mode while VCs are configured can produce undesired results, and the change will not take effect until the next router reload.

# Configuring atm pxf queuing

The Cisco 10000 series router supports two ATM traffic classes when you configure **atm pxf queuing**: unshaped UBR and VBR-nrt. When you specify an output PCR for an unshaped UBR class, the Cisco 10000 series router accepts the PCR. However, the router does not use the PCR value and it does not notify you of this omission.

For information about configuring the traffic classes, see the "Configuring no atm pxf queuing" section on page 2-15.

**Note** Before you configure VCs on an interface, configure the **atm pxf queuing** mode for the port (**atm pxf queuing** or **no atm pxf queuing**). After you configure the mode, then configure the VCs. Do not change the mode while VCs are configured on the interface. If you need to change the mode, delete the VCs first and then change the mode. Changing the mode while VCs are configured can produce undesired results.

Table 2-8 lists the number of active VCs the ATM line cards support in **atm pxf queuing** mode for Cisco IOS Release 12.3(7)XI2 or later releases.

*Table 2-8 Active VCs on ATM Line Cards*

| Line Card | Maximum VCs per Port | Maximum VCs per Module | No. VBR, CBR, Shaped UBR VCs |
|---|---|---|---|
| E3/DS3 | 4,096 | 32,768[1] | 28,672[2] |
| OC-3 | 8,191 | 32,764[3] | 28,672[4] |
| OC-12 | 16,384 (previously 14,436) | 16,384 | 16,384 |

1. For 32,768 VCs per module, 4096 of them must be unshaped UBR VCs.
2. For 28,672 VBR, CBR, and shaped UBR VCs, no VCs can be in shaped VP tunnels. If VCs are in shaped VPs, the number of VBR, CBR, and shaped UBR VCs is 22,204.
3. For 32,764 VCs per module, 4096 of them must be unshaped UBR VCs.
4. For 28,672 VBR, CBR, and shaped UBR VCs, no VCs can be in shaped VP tunnels. If VCs are in shaped VPs, the number of VBR, CBR, and shaped UBR VCs is 22,204.

You can configure the maximum number of VCs across the ports in any fashion, provided that you do not exceed the per-port maximum.

Although the maximum number of VBR, CBR, and shaped UBR VCs per E3/DS3 and OC-3 ATM line card is 28,672 VCs, the router supports a maximum of 22,204 VBR, CBR, and shaped UBR VCs per line card that you can place within virtual path (VP) tunnels. If you attempt to bring up more than 22,204 VCs in a configuration that includes VP tunnels and VCs (hierarchical traffic shaping configuration), the VCs might not assign traffic correctly or the VCs might not come up at all. Be sure to limit the number of configured VBR, CBR, and shaped UBR VCs on an ATM card to less than 22,204 VCs if you place the VCs in VP tunnels.

For the OC-12 ATM line card, the router supports 16,384 VCs in VP tunnels.

# Configuring keepalive

The **keepalive** command sets the keepalive timer for a specific interface. To ensure proper scaling and to minimize CPU utilization, set the timer for 30 seconds or longer (Example 2-14). The default value is 10 seconds.

***Example 2-14 Configuring keepalive for a Virtual Template Interface***

```
interface Virtual-Template1
 ip unnumbered Loopback1
 keepalive 30
 no peer default ip address
 ppp authentication pap
```

# Enhancing Scalability of Per-User Configurations

To enhance scalability of per-user configurations without changing the router configuration, use the **ip:vrf-id** and **ip:ip-unnumbered** RADIUS attributes. These per-user vendor specific attributes (VSAs) are used to map sessions to VRFs and IP unnumbered interfaces. The VSAs apply to virtual access subinterfaces and are processed during PPP authorization.

In releases earlier than Cisco IOS Release 12.2(16)BX1, the **lcp:interface-config** RADIUS attribute is used to map sessions to VRFs. This per-user VSA applies to any type of interface configuration, including virtual access interfaces. Valid values of this VSA are essentially any valid Cisco IOS interface command; however, not all Cisco IOS commands are supported on virtual access subinterfaces. To accommodate the requirements of the **lcp:interface-config** VSA, the per-user authorization process forces the Cisco 10000 series router to create full virtual access interfaces, which consume more memory and are less scalable.

In Cisco IOS Release 12.2(16)BX1 and later releases, the **ip:vrf-id** attribute is used to map sessions to VRFs. Any profile that uses the **ip:vrf-id** VSA must also use the **ip:ip-unnumbered** VSA to install IP configurations on the virtual access interface that is to be created. PPP that is used on a virtual access interface to be created requires the **ip:ip-unnumbered** VSA. An Internet Protocol Control Protocol (IPCP) session is not established if IP is not configured on the interface. You must configure either the **ip address** command or the **ip unnumbered** command on the interface so that these configurations are present on the virtual access interface that is to be created. However, specifying the **ip address** and **ip unnumbered** commands on a virtual template interface is not required because any pre-existing IP configurations are removed when the **ip:ip-vrf** VSA is installed on the virtual access interface. Therefore, any profile that uses the **ip:vrf-id** VSA must also use the **ip:ip-unnumbered** VSA to install IP configurations on the virtual access interface that is to be created.

These per-user VSAs can be applied to virtual access subinterfaces; therefore, the per-user authorization process does not require the creation of full virtual access interfaces, which improves scalability.

## Setting VRF and IP Unnumbered Interface Configurations in User Profiles

Although the Cisco 10000 series router continues to support the **lcp:interface-config** VSA, the **ip:vrf-id** and **ip:ip-unnumbered** VSAs provide another way to set the VRF and IP unnumbered interface configurations in user profiles. The **ip:vrf-id** and **ip:ip-unnumbered** VSAs have the following syntax:

```
Cisco:Cisco-AVpair = "ip:vrf-id=vrf-name"
Cisco:Cisco-AVpair = "ip:ip-unnumbered=interface-name"
```

You should specify only one **ip:vrf-id** and one **ip:ip-unnumbered** value in a user profile. However, if the profile configuration includes multiple values, the Cisco 10000 series router applies the value of the last VSA received, and creates a virtual access subinterface. If the profile includes the **lcp:interface-config** VSA, the router always applies the value of the **lcp:interface-config** VSA, and creates a full virtual access interface.

In Cisco IOS Release 12.2(15)BX, when you specify a VRF in a user profile, but do not configure the VRF on the Cisco 10000 series router, the router accepts the profile. However, in Cisco IOS Release 12.2(16)BX1 and later releases, the router rejects the profile.

## Setting VRF and IP Unnumbered Interface Configuration in a Virtual Interface Template

You can specify one VSA value in the user profile on RADIUS and another value locally in the virtual template interface. The Cisco 10000 series router clones the template and then applies the values configured in the profiles it receives from RADIUS, resulting in the removal of any IP configurations when the router applies the profile values.

## Redefining User Profiles to Use the ip:vrf-id and ip:ip-unnumbered VSAs

The requirement of a full virtual access interface when using the **lcp:interface-config** VSA in user profiles can result in scalability issues such as increased memory consumption. This situation is especially true when the Cisco 10000 series router attempts to apply a large number of per-user profiles that include the **lcp:interface-config** VSA. Therefore, when updating your user profiles, we recommend that you redefine the **lcp:interface-config** VSA to the scalable **ip:vrf-id** and **ip:ip-unnumbered** VSAs.

Example 2-15 shows how to redefine the VRF named *newyork* using the **ip:vrf-id** VSA.

***Example 2-15 Redefining VRF Configurations***

```
Change:
Cisco:Cisco-Avpair = "lcp:interface-config=ip vrf forwarding newyork"

To:
Cisco:Cisco-Avpair = "ip:vrf-id=newyork"
```

Example 2-16 shows how to redefine the Loopback 0 interface using the **ip:ip-unnumbered** VSA.

***Example 2-16 Redefining IP Unnumbered Interfaces***

```
Change:
Cisco:Cisco-Avpair = "lcp:interface-config=ip unnumbered Loopback 0"

To:
Cisco:Cisco-Avpair = "ip:ip-unnumbered=Loopback 0"
```

# Placing PPPoA Sessions in Listening Mode

For better scalability and faster convergence of PPPoA, PPPoEoA, or LAC sessions, set sessions to passive mode, using the **atm pppatm passive** command in ATM subinterface configuration mode. This command places PPP or L2TP sessions on an ATM subinterface into listening mode. For large-scale PPP terminated aggregation (PPPoA and PPPoEoA) and L2TP (LAC), the **atm pppatm passive** command is required.

Instead of sending out Link Control Protocol (LCP) packets to establish the sessions actively, the sessions listen to the incoming LCP packets and become active only after they receive their first LCP packet. When PPPoX is in passive mode, the LAC brings up the sessions only when the subscribers become active and does not waste processing power polling all the sessions.

The following example configures passive mode for the PPPoA sessions on an ATM multipoint subinterface:

```
Router(config)# interface atm 1/0.1 multipoint
Router(config-subif)# atm pppatm passive
Router(config-subif)# range range-pppoa-1 pvc 100 199
Router(config-subif-atm-range)# encapsulation aal5mux ppp virtual-template 1
```

# Scaling L2TP Tunnel Configurations

To prevent head-of-the-line blocking of the IP input process and save system resources, configure the following command in global configuration mode:

```
Router(config)# vpdn ip udp ignore checksum
```

When you configure this command, the router directly queues L2TP Hello packets and Hello acknowledgements to the L2TP control process. We recommend that you configure this command in all scaled LAC and LNS L2TP tunnel configurations.

If you do not configure the **vpdn ip udp ignore checksum** command, the L2TP software sends the packet to UDP to validate the checksum. When too many packets are queued to the IP input process, the router starts selective packet discard (SPD), which causes IP packets to be dropped.

**Note** Head-of-the-line blocking of the IP input process might occur in other non-L2TP configurations. A flush occurring on an input interface indicates that SPD is discarding packets.

# Using the RADIUS Attribute cisco-avpair="lcp:interface-config"

When you use the **lcp:interface-config** RADIUS attribute to reconfigure the virtual-access subscriber interface, scaling on the Cisco 10000 series router decreases for the following reasons:

- The **lcp:interface-config** command syntax includes an IOS interface configuration command. This command is any valid IOS command that can be applied to an interface. When the **lcp:interface-config** attribute is downloaded from the RADIUS server to the Cisco 10000 series router, the command parser is activated to configure the interface as per AV-pair, determining if the option is valid and then applying the configuration to the virtual access interface (VAI).
- The **lcp:interface-config** command forces the Cisco 10000 series router to create full VAIs instead of subinterface VAIs. Full VAIs consume more memory and are less scalable, and they follow a significantly slower and different path when sessions are established.
- The **lcp:interface-config** command degrades the call rate.

To enhance the scalability of per-user configurations, in many cases different Cisco AV-pairs are available to place the subscriber interface in a virtual routing and forwarding (VRF) instance or to apply a policy map to the session. For example, use the **ip:vrf-id** and **ip:ip-unnumbered** VSAs to reconfigure the user's VRF. For more information, see the "Enhancing Scalability of Per-User Configurations" section on page 2-17.

# Using Full Virtual Access Interfaces

A virtual access interface (VAI) is an interface that is dynamically created to terminate PPP subscribers. The Cisco router indicates full VAIs using a notation similar to **Virtual-Access6** (without a .number suffix).

**Note** For Cisco IOS Release 12.3(7)XI and later releases, the router does not support the use of full VAIs for broadband interfaces due to the scaling implications full VAIs have.

In general, the router creates full VAIs for one or more of the following reasons:

- Virtual template interface-specific configuration

  Some Cisco IOS configuration commands configured under the virtual template, such as the **carrier-delay** command, can force the router to create a full VAI. You can use the test command to determine the interface-specific configuration under the virtual template that triggered the full VAI.
- RADIUS attribute **lcp:interface-config**
- Global configuration **no virtual-template subinterface** command

# Preventing Full Virtual Access Interfaces

The **lcp:interface-config** RADIUS attribute is used to reconfigure the subscriber interface. To accommodate the requirements of this attribute, the per-user authorization process forces the router to create full VAIs.

Cisco IOS Release 12.2(31)SB2, Release 12.2(28)SB6, and later releases include an enhancement that allows you to use the **lcp:interface-config** attribute while preserving subvirtual access subinterfaces. You can achieve this behaviour in the following ways:

- Entering the following command in global configuration mode to preserve virtual access subinterfaces:

  ```
  Router(config)# aaa policy interface-config allow-subinterface
  ```

- Sending a Cisco attribute-value pair (AV-pair) in the user's profile on the RADIUS server:

  ```
  cisco-avpair="lcp:interface-config allow-subinterface=yes"
  ```

When you use the **aaa policy interface-config allow-subinterface** command, the router does not allow you to reconfigure the router using any commands that interact with the interface's hardware interface descriptor block (HWIDB), for example, the **compression** command.

When you use the **lcp:interface-config** attribute, sessions are not established if the sessions receive the attribute and the attribute reconfigures the HWIDB for the virtual access interface (VAI).

When the **allow-subinterface=yes** option is used in the Cisco AV-pair or the **aaa policy interface-config allow-subinterface** command is set, enter the following command to verify the condition for which a full VAI reconfiguration is required:

```
Router# debug sss feature-name interface-config {error | event}
```

In general, for interface reconfiguration, use the dedicated Cisco vendor specific attributes (VSAs). For example, use **Cisco-Policy-Up** or **Cisco-Policy-Down,** or **ip:vrf-id** instead of **lcp:interface-config**. Alternatively, when no dedicated Cisco AV-pair is present, use **lcp:interface-config** with the **allow-subinterface=yes** option, or the **aaa policy interface-config allow-subinterface** command to preserve VAI subinterfaces (for example, to enable multicast on the subscriber interface).