



# Release Notes for the Cisco 10000 Series Router for Cisco IOS Release 12.3(7)XI1

---

**First Published: August 2, 2004**

**Revised: September 7, 2006**

These release notes provide information about Cisco IOS Release 12.3(7)XI1, which provides broadband aggregation and leased-line features for the Cisco 10000 series router.

These release notes are updated as needed to describe new features, memory requirements, hardware support, software platform deferrals, and changes to the microcode and related documents.

Cisco IOS Release 12.3(7)XI1 is based on the following releases:

- Cisco IOS Release 12.2(16)BX
- Cisco IOS Release 12.3T

To review the release notes for Cisco IOS Release 12.2(16)BX, go to the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/aggr/10000/10krn/122bx/index.htm>

To review the release notes for Cisco IOS Release 12.3, go to the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123reInt/xprn123/index.htm>

## Contents

This document contains the following sections:

- [System Requirements, page 2](#)
- [New Features—Cisco IOS Release 12.3\(7\)XI1, page 3](#)
- [Limitations and Restrictions, page 12](#)
- [Important Notes, page 19](#)
- [Open Caveats—Cisco IOS Release 12.3\(7\)XI1, page 23](#)
- [Resolved Caveats—Cisco IOS Release 12.3\(7\)XI1, page 42](#)
- [Obtaining Documentation, page 43](#)
- [Documentation Feedback, page 44](#)



---

Corporate Headquarters:  
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2006 Cisco Systems, Inc. All rights reserved.

- [Obtaining Technical Assistance, page 44](#)
- [Obtaining Additional Publications and Information, page 45](#)

## System Requirements

Cisco IOS Release 12.3(7)XI1 requires that you have the performance routing engine (PRE), Part Number ESR-PRE2 installed in the Cisco 10000 series router chassis. To verify which PRE is installed in the router, use the **show version** command.

## Route Processor Redundancy Mode

The Cisco 10000 series router supports route processor redundancy (RPR) mode or RPR+ mode to provide fault resistance and to ensure high availability. In RPR mode, one supervisor engine is active and operational while the second supervisor engine is in standby mode waiting for the active supervisor to fail so that it can take over and maintain the operation of the router. In RPR+ mode, the standby supervisor engine is fully initialized and configured, which shortens the time needed to switch over to the standby supervisor.

When upgrading or downgrading the Cisco IOS software, the RPR mode used on the Cisco 10000 series router depends upon the Cisco IOS software currently running on the Cisco 10000 series router and the Cisco IOS software to which you want to upgrade or downgrade.

[Table 1](#) lists the RPR modes used when upgrading or downgrading Cisco IOS software. For example, when upgrading to Cisco IOS Release 12.3(7)XI1 from Release 12.2(16)BX, the router uses RPR mode instead of RPR+ mode. When downgrading to Cisco IOS Release 12.2(16)BX from Release 12.3(7)XI1, the router uses RPR mode.

**Table 1** RPR Modes for Cisco IOS Software Releases

Releases	12.2(16)BX	12.3(7)XI1
12.2(16)BX	RPR+	RPR
12.3(7)XI1	RPR	RPR+

## Before You Upgrade the Cisco IOS Software

Before you upgrade (or downgrade) the Cisco IOS software running on the Cisco 10000 series router, save the running configuration file. In RPR mode, the router synchronizes only the startup configuration.

## Upgrading to a New Software Release

For specific information about upgrading your Cisco 10000 series router to a new software release, refer to the [Cisco 10000 Series Router Software Configuration Guide](#).

For general information about upgrading to a new software release, refer to the product bulletin [Cisco IOS Upgrade Ordering Instructions](#).

For additional information about ordering Cisco IOS software, refer to the [Cisco IOS Software Releases](#).

# New Features—Cisco IOS Release 12.3(7)XI1

The following new features and improvements are supported on the Cisco 10000 series router in Cisco IOS Release 12.3(7)XI1. While some of the following features are supported on other releases on the Cisco 10000 series router, these features are newly supported in Cisco IOS Release 12.3(7)XI1:

- [3-Color Policer, page 4](#)
- [3-Level Hierarchical QoS Policies, page 4](#)
- [BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN, page 4](#)
- [Extended NAS-Port-Type and NAS-Port Support, page 4](#)
- [Half-Duplex VRF, page 5](#)
- [Hierarchical Shaping, page 6](#)
- [IEEE 802.1Q-in-Q VLAN Tag Termination, page 6](#)
- [Interface Oversubscription, page 6](#)
- [IP Receive ACLs, page 6](#)
- [IP Unnumbered on VLAN, page 7](#)
- [Lawful Intercept, page 7](#)
- [Local AAA Server, User Database—Domain to VRF, page 7](#)
- [MIB Enhancements, page 8](#)
- [MPLS QoS, page 8](#)
- [MPLS Traffic Engineering—DiffServ Aware, page 8](#)
- [Multirouter APS, page 9](#)
- [Percent-Based Policing, page 9](#)
- [Per DSCP WRED, page 9](#)
- [Per Precedence WRED Statistics, page 10](#)
- [RADIUS Packet of Disconnect, page 10](#)
- [Scaling Enhancements, page 10](#)
- [Strict Priority Queuing, page 11](#)
- [Time-Based ACLs, page 11](#)
- [VBR-nrt Oversubscription, page 11](#)
- [VC Weighting, page 11](#)
- [WRED with Queue Limit, page 12](#)

For more information about the new features in Cisco IOS Release 12.3(7)XI1, refer to the following documentation:

- [Cisco 10000 Series Broadband Aggregation and Leased-Line Configuration Guide](#)
- [Cisco 10000 Series Router Feature Map](#)

For information about new features supported on the Cisco 10000 series router in other releases, see the appropriate Release Notes at the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/aggr/10000/10krn/index.htm>

## 3-Color Policer

The 3-color policer feature provides a single-rate, 3-color marker. A 2-color marker as supported in earlier releases, meters a traffic stream classifying it into two groups (or colors): the traffic conforming to the specified committed information rate (CIR) and the burst parameters, and the traffic exceeding either the CIR or the burst parameters. A 3-color marker classifies the metered traffic into three groups, adding an additional color for the nonconforming traffic.

The 3-color marker distinguishes between the nonconforming traffic that occasionally bursts a certain number of bytes more than the CIR allowance and the traffic that continually violates the CIR allowance. A 3-color marker meets the requirements of applications that require three service levels: guaranteed, best effort, and deny. A three-color policer enables the Cisco 10000 series router to comply with RFC 2597.

## 3-Level Hierarchical QoS Policies

The 3-Level Hierarchical QoS Policies feature enables you to apply a service policy inside a policy map to define hierarchical policies. This feature increases the hierarchical levels of a nested QoS policy from two to three levels.

A hierarchical policy extends QoS by enabling you to combine one or more classes and applying specific actions on the aggregate traffic as well as executing class-specific actions. For example, a hierarchical policy can define a minimum bandwidth for two classes and specify a combined maximum bandwidth for the two classes. Similarly, a 3-level policy can define a minimum bandwidth for each type of traffic on a virtual circuit and a maximum bandwidth for the virtual circuit's total traffic. A 3-level hierarchical policy can also selectively police a subclass of each guaranteed class and place a maximum transmission limit on the aggregate traffic.

A 3-level policy is typically used to define the transmission capacity of a virtual circuit in the top level, class-based queuing at the middle level, and marking or metering in the bottom level.

## BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN

The BGP Multipath Load Sharing for eBGP and iBGP in an MPLS VPN feature allows you to configure multipath load balancing with both external Border Gateway Protocol (eBGP) and internal BGP (iBGP) paths in BGP networks that are configured to use Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). BGP Multipath Load Sharing provides improved load-balancing deployment and service offering capabilities and is useful for multihomed autonomous systems and provider edge (PE) routers that import both eBGP and iBGP paths from multihomed and stub networks.

## Extended NAS-Port-Type and NAS-Port Support

Cisco support for NAS-Port-Type (RADIUS attribute 61), NAS-Port (RADIUS attribute 5), and NAS-Port-ID (RADIUS attribute 87) has been changed as discussed in the following sections.

## NAS-Port-Type (RADIUS Attribute 61)

Remote Authentication Dial-In User Service (RADIUS) attributes are used to define specific Authentication, Authorization, and Accounting (AAA) elements in a user profile, which is stored on the RADIUS daemon. Currently the Internet Engineering Task Force (IETF) RADIUS attributes that are supported include an attribute 61, NAS-Port-Type. NAS-Port-Type indicates the type of physical port the network access server (NAS) is using to authenticate the user.

However there was no method to identify NAS-Port-Type based on a specific broadband service type because the RADIUS RFC does not support extended types that defines these types of ports. Basically all PPPoA, PPPoEoE, and PPPoEoA sessions were identified as being VIRTUAL and all PPPoEoVLAN and PPPoEoQinQ as ETHERNET.

The Extended NAS-Port-Type Attribute Support feature expands NAS-Port-Type, attribute 61, in order that the client can better identify what type of service is taking place on the different types of ports.

## NAS-Port (RADIUS Attribute 5)

The NAS-Port (RADIUS attribute 5) is a 32 bit value that uniquely represents the physical or logical port the user is attempting to authenticate on. A logical port can be represented by the virtual path identifier (VPI) and virtual channel identifier (VCI) for an ATM interface, or by the VLAN ID or Q-in-Q ID for an Ethernet interface.

Because each platform and service may have different port information which are relevant to their environment, there is no one unique way to populate this attribute. Currently Cisco has 4 hard wired formats (a-d) which are service specific and 1 configurable format (e) which can be tailored to customer and platform-specific needs.

Previously format e only allowed customizing 1 global format for all call types on a device, which limited its usefulness on devices that contained multiple services. With the extended NAS-port support, you can now configure a custom format e string for any and all service types based on the value of the NAS-Port-Type (RADIUS attribute 61). That is, when building the RADIUS Access or Accounting request, the encoding routine will pick the specific format e string defined for the session's NAS-Port-Type value and use that first instead of using the default global format e string.

## NAS-Port-ID (RADIUS Attribute 87)

The NAS-Port-ID (RADIUS attribute 87) contains the character text string identifier of the NAS port that is authenticating the user. This text string typically matches the interface description found under the CLI configuration. This attribute was previously available under Cisco Vendor Specific Attribute (VSA) "cisco-nas-port". But it is now sent by default under the IETF attribute 87 as per customer demand.

## Half-Duplex VRF

The Half-Duplex VRF (HDVRF) feature provides scalable hub and spoke connectivity for subscribers of a multiprotocol label switching-based virtual private network (MPLS VPN) service. These subscribers connect to the provider edge (PE) router of the wholesale service provider, and they use the same or different services (for example, the same or different VRFs). The HDVRF feature prevents local connectivity between subscribers at the spoke PE router and ensures that a hub site provides subscriber connectivity. Any sites that connect to the same PE router must forward intersite traffic using the hub site. This ensures that the routing done at the spoke site is always access side interface to network side interface, or network side interface to access side interface, and never access side to access side.

## Hierarchical Shaping

The Hierarchical Shaping feature provides two levels of shaping—per VC ATM level shaping and per VC packet level shaping—and provides per-VC and per-VP traffic shaping to control or modify the flow of traffic on an interface. Traffic shaping limits throughput by buffering excess traffic instead of dropping packets. The shaping function also ensures that traffic from one VC does not adversely impact another VC, resulting in loss of data.

The Cisco 10000 series router supports the Hierarchical Shaping feature for the following ATM line cards:

- OC-12
- 4-port OC-3
- 8-port DS3/E3

## IEEE 802.1Q-in-Q VLAN Tag Termination

For the emerging broadband Ethernet-based DSLAM market, the Cisco 10000 series router supports Q-in-Q encapsulation. With an Ethernet-based DSLAM model, customers typically get their own VLAN and all these VLANs are aggregated on a DSLAM.

VLAN aggregation on a DSLAM results in a lot of aggregate VLANs that at some point need to be terminated on the broadband remote access servers (BRAS). Although the model could connect the DSLAMs directly to the BRAS, a more common model uses the existing Ethernet-switched network where each DSLAM VLAN ID is tagged with a second tag (QinQ) as it connects into the Ethernet-switched network.

The only model that is supported is PPPoE over Q-in-Q (PPPoEoQinQ). This can either be a PPP terminated session or as a L2TP LAC session. No IP over Q-in-Q is supported.

The Cisco 10000 series router already supports plain PPPoE and PPP over 802.1Q encapsulation; support for PPP over Q-in-Q encapsulation is new. PPP over Q-in-Q encapsulation processing is an extension to 802.1q encapsulation processing.

## Interface Oversubscription

The interface oversubscription feature offers providers the choice to improve network utilization of otherwise underutilized shared networks by leveraging statistical multiplexing on Frame Relay and IEEE 802.1Q networks.

## IP Receive ACLs

The IP Receive ACLs feature provides basic filtering capability for traffic that is destined for the router and protects the router from remote intrusions.

To restrict access to the router, you apply a numbered ACL to the ingress interface of the router. You can restrict access to the router to known and trusted sources, and to expected traffic profiles. The IP Receive ACLs feature supports both standard and extended ACLs. The rules for numbered ACLs also apply to the access control entries (ACEs) of the IP receive ACL.

The IP receive ACL filters traffic on the parallel express forwarding engine (PXF) before filtering the packets received by the route processor (RP). This feature protects the router from denial of service (DoS) floods, thereby preventing the flood from degrading the performance of the route processor (RP).

## IP Unnumbered on VLAN

The IP Unnumbered on VLAN feature helps to conserve IP address space for service provider configurations that include Ethernet VLAN subinterfaces.

Prior to Cisco IOS Release 12.3(7)X11, IP support for VLAN subinterfaces required that you configure separate IP subnets for each of the subinterfaces that terminate the VLAN. This resulted in inefficient use of the IP address space because an entire IP subnet is often not needed for the hosts assigned to a VLAN. The IP Unnumbered on VLANs feature helps to conserve IP address space for service provider configurations that include Ethernet VLAN subinterfaces.

VLAN subinterfaces with IP unnumbered configured support DHCP for IP address allocation. The DHCP server uses the information in DHCP Option 82 to assign IP addresses to the hosts on a VLAN. The routing table is dynamically updated to insert an IP route for the IP address assigned on each of the subinterfaces. These IP host routes exist until the DHCP lease time expires or the host releases the leased address.

## Lawful Intercept

Lawful intercept is a process that enables a Law Enforcement Agency (LEA) to perform electronic surveillance on an individual as authorized by a court order. To assist in the surveillance, the service provider intercepts the target's traffic as it passes through one of their routers, and sends a copy of the intercepted traffic to a third party mediation device (also in the service provider network). This third party mediation device formats and delivers the data to the LEA without the target's knowledge. The Lawful Intercept feature is available in the c10k2-k9p11u2-mz image.

## Local AAA Server, User Database—Domain to VRF

The Local AAA Server, User Database—Domain to VRF feature extends the Cisco IOS AAA Authorization to local AAA profiles on the router without using an AAA Server. The local user database acts as a local AAA server, and is fully compatible with any external AAA Server. If you want to maintain your user database locally or provide a failover local mechanism, you no longer have to sacrifice policy options when defining local users.

This flexibility allows you to provide complete user authentication and authorization locally within Cisco IOS without using an AAA Server, provided the local username list is relatively small. While authentication can be done on the router for a limited number of user names, it might make more sense and be much more scalable to use an AAA Server. Note that accounting is still be done on an AAA server and is not be supported on the router.

The key function this feature provides is a mapping of user domain names to local AAA profiles. This allows AAA attributes to be applied to the PPP session as part of the PPP session establishment. These local AAA attributes are RADIUS attributes that would normally be defined on a Radius Server but now are defined locally on the router.

Subscriber profiles are used to match user domain names, and on a match to use a defined AAA attribute list. The AAA attribute list contains a list of valid Cisco IOS format AAA attributes.

## MIB Enhancements

The MIB Enhancements feature includes the following additional MIBs and MIB support:

- MPLS-LSR-MIB
- MPLS-TE-MIB
- MPLS-VPN-MIB
- CISCO-TAP-MIB
- CISCO-IP-LOCAL-POOL-MIB
- Addition of the per precedence/DSCP/discard class statistics in the QoS MIB
- MPLS-LDP-MIB (Version 8)
- MPLS enhancements to the IF-MIB

For more information about MIBs supported on the Cisco 10000 series router, refer to:

- *Cisco 10000 Series Broadband MIB Specifications Guide*
- *Cisco 10000 Series Leased-Line MIB Specifications Guide*

## MPLS QoS

When a customer transmits IP packets from one site to another, the IP precedence field (the first three bits of the DSCP field in the header of an IP packet) specifies the class of service. Based on the IP precedence marking, the packet can be given a change in treatment such as the latency or the percent of bandwidth allowed. If the service provider network is an MPLS network, then the IP precedence bits are copied into the MPLS EXP field at the edge of the network. However, the service provider might want to set an MPLS packet's QoS to a different value determined by the service offering.

MPLS can be used to "tunnel" the QoS of a packet. The MPLS EXP field can be marked independent of the PHB. The service provider can choose from a variety of criteria (including those based on IP PHB) to classify a packet and set the MPLS EXP field. This allows the service provider to set the MPLS EXP field instead of overwriting the value in the customer's IP precedence field. The IP header remains available for the customer's use; the marking of an IP packet is not changed as the packet travels through the MPLS network. In some instances, it is desirable to extend the MPLS PHB to the egress interface between the provider edge (PE) router and customer edge (CE) router. This has the effect of extending the MPLS QoS tunnel, which allows the MPLS network owner to classify scheduling and discarding behavior on that final interface.

## MPLS Traffic Engineering—DiffServ Aware

The MPLS Traffic Engineering—DiffServ Aware (DS-TE) feature extends MPLS traffic engineering capabilities to provide stricter quality of service (QoS) guarantees. TE tunnels provide differentiated services (DiffServ) to satisfy bandwidth requirements of regular traffic. However, the bandwidth currently advertised for TE tunnels and the tunnel traffic do not correspond to any queue. Instead, the MPLS class of service (CoS) provides DiffServ service, which is adequate for most customer services. Special services such as voice, however, require stricter QoS guarantees. The DS-TE feature addresses this need, providing strict bandwidth guarantees for TE tunnels.

The DS-TE feature introduces awareness of a particular class of traffic referred to as the guaranteed bandwidth traffic. DS-TE enables service providers to perform separate admission control and separate route computation of the guaranteed bandwidth traffic. The service provider can, therefore, develop QoS services for end customers that rely on signaled QoS rather than provisioned QoS, which enables the service provider to build QoS services with hard commitments and without overprovisioning.

## Multirouter APS

The Multirouter APS (MR-APS) feature enables ATM connections to switch from one ATM circuit to another ATM circuit if a circuit failure occurs. ATM interfaces can be switched in response to a router failure, degradation or loss of channel signal, or manual intervention.

The protection mechanism used for this feature has a linear 1+1 architecture as described in the Bellcore publication TR-TSY-000253, SONET Transport Systems; Common Generic Criteria, Section 5.3. The connection may be bidirectional or unidirectional and revertive or nonrevertive. The default is bidirectional. The switching mode must be the same on the far end of the connection.

In Cisco IOS Release 12.3(7)X11, MR-APS is supported for the following line cards:

- OC-3 ATM
- OC-12 ATM
- 4-port Channelized STM-1

## Percent-Based Policing

The Percent-Based Policing feature enables you to configure traffic policing in bits-per-second or as a percentage of bandwidth of the network interface on which policing is applied. Configuring traffic policing based on bandwidth percentage enables you to use the same policy map for multiple interfaces with differing amounts of bandwidth.

## Per DSCP WRED

The per differentiated services code point weighted random early detection (DSCP WRED) feature enables the Cisco 10000 series router to randomly drop packets with a specific DSCP value, according to the DSCP thresholds you configure.

Differentiated Services (DiffServ) is a QoS model that increases the number of definable priority levels by reallocating bits of an IP packet for priority marking. The six most significant bits of the type of service (ToS) field are the DiffServ field. The last two bits in the DiffServ field are used as Early Congestion Notification (ECN) bits.

The per DSCP WRED feature enables you to configure eight unique drop precedence levels for one queue. Each of the 64 DSCP levels correspond to one of the eight levels. Previously, when you configured the eight unique drop precedence levels, all of the queues configured on an interface shared the different levels. The per DSCP WRED feature enhances support to provide eight unique levels per queue.

## Per Precedence WRED Statistics

The Enhanced Weighted Random Early Detection (WRED) Statistics feature maintains separate WRED drop statistics for each IP precedence, discard-class, and differentiated services code point (DSCP) value. The **show policy-map** command has been enhanced to show WRED drop counts for each profile. In earlier releases, RED drop counts were maintained only for each class.

## RADIUS Packet of Disconnect

In Cisco IOS Release 12.3(7)XI1, the RADIUS Packet of Disconnect feature consists of a method for terminating a session that has already been connected. This packet of disconnect (POD) is a RADIUS access\_request packet and is intended to be used in situations where the authenticating agent server wants to disconnect the user after the session has been accepted by the RADIUS access\_accept packet.

## Scaling Enhancements

The Scaling Enhancements feature 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 4,096 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

### Queue Scaling

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

## Strict Priority Queuing

The Priority Queuing feature guarantees latency for any packet that enters the priority queue regardless of the current congestion level on the link. Strict priority queue mode is supported as the *only* mode of operation for a priority queue in Cisco IOS Release 12.3(7)XI1.

## Time-Based ACLs

Time-based ACLs allow the network administrator to define a time range when certain resources may be accessed, thus providing greater control over resource usage. Time-based ACLs are functionally similar to extended ACLs and control access to the router for a specific time period.

A time range defines the specific times of the day and week that the ACL is active. A time range name identifies the time range. The access control entries (ACEs) reference the time range name, which causes the router to impose the time restriction on the ACEs. The time range relies on the router system clock to activate or deactivate an ACE.

Previously, access list statements were always in effect after they were applied to an interface. However, using the **time-range** command, network administrators can now define when the permit and deny statements in the ACL are in effect. Both named and numbered access lists can reference a time range.

## VBR-nrt Oversubscription

The Variable Bit Rate Non-Real Time (VBR-nrt) Oversubscription feature enables service providers to improve network utilization of otherwise under utilized shared networks by leveraging statistical multiplexing on ATM networks. Instead of supporting only unconditional reservation of network bandwidth to VCs, the router offers VC oversubscription to statistically guarantee bandwidth to VCs.

In releases prior to Cisco IOS Release 12.3(7)XI1, a call admission check (CAC) prevented you from assigning more bandwidth to virtual circuits (VCs) than a port's total bandwidth. The VBR-nrt Oversubscription feature enables you to specify the amount of oversubscription (oversubscription factor) you want to allow. The CAC check is based on the oversubscription factor you specify and evaluated separately for both VCs and VP tunnels into the port, and VCs into VP tunnels. When the total assigned bandwidth exceeds the physical capacity, the router provides each VC's bandwidth reservation, as long as a limited number of VCs activate at one time. By doing so, the router takes advantage of statistical multiplexing to provide better network utilization at the expense of degraded service under congestion.

The oversubscription factor is also used to evaluate the amount of bandwidth allocated for unspecified bit rate (UBR) VCs. Prior to Cisco IOS Release 12.3(7)XI1, UBR VCs received the bandwidth remaining after other VCs had been allocated bandwidth. The CAC check now adjusts the bandwidth for UBR VCs based on the oversubscription factor.

## VC Weighting

In earlier releases, the weight of a particular VC was proportional to the VC speed and was not directly controllable by the user (other than by changing the VC rate). In Cisco IOS Release 12.3(7)XI1, the VC Weighting feature adds the ability to configure the VC weight directly.

## WRED with Queue Limit

The Weighted Random Early Detection (WRED) with Queue Limit feature is a congestion avoidance mechanism that expands your ability to customize the size of a WRED queue. Using this feature, you can configure a packet drop policy for a traffic class that includes a bandwidth guarantee and simultaneously limit the maximum number of packets allowed to accumulate in a traffic class queue.

In Cisco IOS Release 12.3(7)X11 or later, you can specify the **random-detect** and **queue-limit** commands in the same class of a policy. Earlier releases allowed you to specify either the **random-detect** command or the **queue-limit** command, but not both commands at the same time.

## Limitations and Restrictions

This section describes limitations and restrictions for the following areas. Be sure to review the following limitations and restrictions before using the features in the Cisco IOS Release 12.3(7)X11:

- [3-Level Hierarchical QoS Policies](#), page 13
- [BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN](#), page 13
- [Controlling the Rate of Logging Messages](#), page 14
- [Frame Relay](#), page 14
- [Half-Duplex VRF](#), page 14
- [Hierarchical Shaping](#), page 14
- [IEEE 802.1Q-in-Q VLAN Tag Termination](#), page 15
- [IP Receive ACLs](#), page 15
- [IP Unnumbered on VLAN](#), page 15
- [MPLS QoS](#), page 16
- [MPLS Traffic Engineering—Diffserv Aware](#), page 16
- [Multirouter Automatic Protection Switching](#), page 16
- [Per Domain VRF With Local Templates](#), page 17
- [Per DSCP WRED](#), page 17
- [Per Precedence WRED Statistics](#), page 17
- [PRE Network Management Ethernet Port](#), page 17
- [RADIUS Packet of Disconnect](#), page 17
- [Strict Priority Queuing](#), page 18
- [Testing Performance of High-Speed Interfaces](#), page 18
- [Time-Based ACLs](#), page 18
- [Variable Bit Rate Non-Real Time Oversubscription](#), page 18
- [WRED with Queue Limit](#), page 19

For more information about the restrictions for a specific feature, refer to the *Cisco 10000 Series Broadband Aggregation and Leased-Line Configuration Guide*.

## 3-Level Hierarchical QoS Policies

The 3-Level Hierarchical QoS Policies feature has the following restrictions:

- You can configure only the class-default class in the top-level policy. Configure the **shape** command for the class-default class and then configure the **service-policy** command to attach an inner policy. You must configure the **shape** command before the **service-policy** command.
- In an inner policy, you cannot configure the **police** and **set** commands for a class if you attach a **service-policy** command to the class. This restriction does not apply to classes that do not have a **service-policy** command configured.
- In a bottommost policy, you can configure only the **police** and **set** commands for a class.
- You cannot have default classes in the bottom most class.
- You cannot attach a **service-policy** command to a bottommost policy.



**Note** The actual shape rate applied to nested-policy traffic might differ from that specified in the policy. For example, a specified shape rate of 10.5 Mbps might be mapped to 11 Mbps. Use the **show policy-map interface** command to determine the actual shape rate.

## BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN

The BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature has the following restrictions:

- The Cisco 10000 series router supports recursive loadsharing, but with the following restriction.  
In recursive load sharing, the information required to forward a packet requires at least 2 lookups. The first lookup determines which provider edge (PE) router is used to reach the final destination. The second lookup determines how to reach the PE router (from first lookup).  
When you configure MPLS VPN, CEF uses recursive load sharing. The first lookup provides the VPN label, the second lookup provides the IGP label. When PXF forwards a packet, it does only 1 lookup which provides both a VPN and an IGP label; 2 lookups in CEF are combined into 1. The restriction for recursive load sharing when PXF forwards a packet is as follows.  
When there are multiple IGP paths between a Cisco 10000 Series PE router to a provider router (P), only per-tag load balancing is supported. That is, PXF is programmed with only one of the paths and this one path is chosen in a round-robin fashion. Because the path is chosen at prefix setup time, it is not possible to predict which path will be selected for which prefix. The path selected depends on the order in which the prefixes are configured in the routing table. The bandwidths of the IGP paths are not considered in the path selection.
- When the routing table contains multiple iBGP paths, a route reflector advertises only one of the paths (one next hop). If a router is behind a route reflector, all routers that are connected to multihomed sites are not advertised unless separate VRFs with different route distinguishers (RDs) are configured for each VRF.
- Each IP routing table entry for a BGP prefix that has multiple iBGP paths uses additional memory. We recommend not using this feature on a router with a low amount of available memory and especially when the router is carrying a full Internet routing table.

## Controlling the Rate of Logging Messages

It is important that you limit the rate that system messages are logged by the Cisco 10000 series router. This helps to avoid a situation in which the router becomes unstable and the CPU is overloaded. To control the output of messages from the system, use the **logging rate-limit** command.

Cisco recommends that you configure the **logging rate-limit** command as follows. This limits the rate of all messages to the console to 10 per second, except for messages with critical priority (level 3) or greater.

```
Router(config)# logging rate-limit console all 10 except critical
```

For more information, refer to the **logging rate-limit command** in the [Cisco IOS Configuration Fundamentals and Network Management Command Reference, Release 12.3](#).

## Frame Relay

The following limitations apply to the Cisco 10000 series router implementation of Frame Relay:

- The **ip rtp reserve** command is not supported.
- Only one priority queue per VC is allowed.

## Half-Duplex VRF

The Half-Duplex VRF feature has the following restrictions:

- In both the upstream and downstream VRFs, routing protocols are not supported on interfaces configured for half-duplex VRFs.
- Half-duplex VRFs apply only to virtual access interfaces (VAIs) and virtual template interfaces. Only IP unnumbered interfaces are supported.
- It is not supported with Routing with Bridged Encapsulation (RBE)

## Hierarchical Shaping

The Hierarchical Shaping feature has the following restrictions:

- The Cisco 10000 series router supports a maximum of 31,500 VCs when the Hierarchical Shaping feature is enabled.
- You can configure a maximum of 127 VP tunnels for each ATM line card. You can configure these 127 VP tunnels across the ports in any fashion.
- The OC-3 and OC-12 line cards support a maximum of 14,336 VCs when configured for hierarchical shaping. The DS3/E3 line card supports a maximum of 8,192 VCs when configured for shaping. You can configure the maximum number of VCs across the ports in any fashion, provided that you do not exceed the per-port maximum. The OC-3 line card is limited to 8,192 VCs per port and the DS3 is limited to 4,096 VCs per port.
- You must have the **atm pxf queuing** command configured on the port. If not, the SAR still does VP shaping and the VCs are sent to the tunnel based on a weighted round robin format; however, the PXF does not shape the VCs. The default queuing mode for a port is **atm pxf queuing**.

- Only variable bit rate (VBR) VCs are allowed in the VP tunnel. You cannot configure unspecified bit rate (UBR) VCs or constant bit rate (CBR) VCs in the tunnels.
- Congestion is not handled at the VP tunnel or at the port. During congestion, shaping is degraded.
- During congestion at the port-level, shaping degrades to simple round robin for all VPs contending for the port's capacity; shaping is not weighted based on the rate of the VPs.

## IEEE 802.1Q-in-Q VLAN Tag Termination

The IEEE 802.1Q-in-Q VLAN Tag Termination feature has the following restrictions:

- Supported on Ethernet, FastEthernet, or Gigabit Ethernet interfaces.
- Supports only Point-to-Point Protocol over Ethernet (PPPoE) packets that are double-tagged for Q-in-Q VLAN tag termination.
- IP and Multiprotocol Label Switching (MPLS) packets are not supported.
- Modular QoS services can be applied to unambiguous subinterfaces only.
- Limited ACL support.

## IP Receive ACLs

The IP receive ACLs feature has the following restrictions:

- A receive ACL must be a numbered ACL. You cannot use a named ACL as the receive ACL.
- The rules for numbered ACLs also apply to the access control entries (ACEs) of receive ACLs.
- Time-based and reflexive ACLs are not supported as receive ACLs.
- Only traffic processed by the RP is filtered. Traffic that is processed exclusively by the Forwarding Processor (FP) is not filtered. For example, GRE tunneled packets, L2TP tunneled packets, and some ICMP packets are not filtered.

## IP Unnumbered on VLAN

The IP Unnumbered on VLANs feature has the following restrictions:

- You can configure IP unnumbered on only Ethernet VLAN subinterfaces and point-to-point interfaces.
- If you configure more than 14,000 IP unnumbered subinterfaces and you have configured EIGRP on all interfaces on a router, the router can stop responding. To avoid this problem, use the **passive-interface default** command (which disables all router interfaces from sending routing updates) and then configure the **no passive-interface** command on selected interfaces you want to send routing updates.
- Service Selection Gateway (SSG) functionality is not supported.

## MPLS QoS

The following limitations apply to the Cisco 10000 series router implementation of MPLS QoS:

- The **match mpls experimental topmost** *exp-value* command (where *exp-value* is in the range 0-7) is supported on both input and output interfaces, on which MPLS is enabled.
- The **set mpls experimental imposition** *mpls-exp-value* command and the **set mpls experimental** *mpls-exp-value* command (where in both cases *mpls-exp-value* is in the range 0-7) are supported on the provider edge (PE) router input interface connecting to customer edge (CE) router. These commands can also be used on input interfaces on the CE, in pipe mode of MPLS QoS Diff Serv tunneling models.

These two commands have the same function, but because the **set mpls experimental** *mpls-exp-value* command is supported only for backward compatibility, Cisco recommends that you use the **set mpls experimental imposition** *mpls-exp-value* command.

- The **set-mpls-exp-imposition-transmit** option of the **police** command is only supported on the PE input interface that is connected to the CE.
- The **mpls ip encapsulate explicit-null** command is supported on the CE router interface that is connected to the PE. This command is only used in pipe mode of MPLS QoS Diff Serv tunneling models.
- When precedence-based weighted random early detection (WRED) is configured on an output policy map and outgoing packets are MPLS packets, the router drops the MPLS packets based on the 3 EXP bits in the MPLS label, instead of using the 3 bits of IP precedence in the underneath IP packets.
- When DSCP-based WRED is configured on an output policy map and outgoing packets are MPLS packets, the router drops the MPLS packets based on the 3 EXP bits in the MPLS label, instead of using the 6 bits of DSCP in the underneath IP packets. The router left shifts the 3 EXP bits and makes it 6 bits. For example, if the value of the EXP bits is 5 (binary 101), the router converts them to binary 101000 (makes it look like 6 DSCP bits), and drops packets based on this value.
- When configuring the **set** and **police** commands in a traffic class, regardless whether it is an input or output policy map, the **police** command is processed later than the **set** command. This means that whatever values implemented by the **police** command override values set by the **set** command. The value can be IP precedence, DSCP, qos-group, MPLS experimental imposition, Discard-class, or ATM CLP bit.
- Discard-class can be a number between 0 and 7; qos-group can be a number between 0 and 63.

## MPLS Traffic Engineering—Diffserv Aware

The DS-TE feature has the following restrictions:

- The total number of TE tunnels (regular TE tunnels and DS-TE tunnels) that can originate on a device is limited to 1013 tunnels.

## Multirouter Automatic Protection Switching

In Cisco IOS Release 12.3(7)XI1, MR-APS is supported for the following line cards:

- OC-3 ATM
- OC-12 ATM

- 4-port Channelized STM-1

## Per Domain VRF With Local Templates

Local templates can be used to forward users to a RADIUS Server for remote AAA. The **ip vrf forwarding** command is not supported under local templates. Therefore, you can only specify a virtual routing and forwarding instance (VRF) by using the **ip:vrf-id VSA** attribute on the RADIUS Server. Do not use Local templates with Subscriber Profiles; they are mutually exclusive.

## Per DSCP WRED

The per DSCP WRED feature has the following restrictions:

- Because Cisco IOS software applies the **random-detect** command on a per interface-basis, you cannot simultaneously configure precedence-based WRED and DSCP-based WRED on a particular interface.
- You cannot use this feature with Multiprotocol Label Switching (MPLS) encapsulated packets. The Cisco 10000 series router supports this feature for use with IP packets only.

## Per Precedence WRED Statistics

In the output of the **show policy-map interface** command, the Tail Drops counter indicates the number of packets dropped because the average queue length exceeds the maximum threshold for the given precedence. However, under burst conditions it is possible that packets can be dropped because the queue is full. These packets are not counted as Tail Drops. The number of packets that are dropped under burst conditions when the queue is full are counted as Output Queue Drops.

## PRE Network Management Ethernet Port

Ensure that the Fast Ethernet NME port on the PRE is configured for auto-negotiation mode, which is the system default. Duplex mode can cause problems, such as flapping. If the port is experiencing such problems and has been configured for duplex mode, use the **no half-duplex** or **no full-duplex** command to disable duplex mode.

## RADIUS Packet of Disconnect

Proper matching identification information must be communicated by the:

- Billing server and router configuration
- Router's original accounting start request
- Server's POD request

## Strict Priority Queuing

If you do not enter a **police** command with the **priority** command, other queues on the link can be starved for bandwidth.

After you use the **priority** command without a **police** command in a policy map, you cannot use the **bandwidth** command in other classes in the same policy map.

## Testing Performance of High-Speed Interfaces

Cisco IOS software running on the Cisco 10000 series router has multiple queues for all classes of traffic over high-speed interfaces. The software selects a queue based on the source and destination address for the packet. This ensures that a traffic flow always uses the same queue and the packets are transmitted in proper order.

When the Cisco 10000 series router is installed in a real network, the high-speed interfaces work efficiently to spread traffic flow equally over the queues. However, using single traffic streams in a laboratory environment may result in less-than-expected performance.

Therefore, to ensure accurate test results, you should test the throughput of the gigabit Ethernet, Packet over SONET (POS), or ATM uplink with multiple source or destination addresses.



Tip

---

To determine if traffic is being properly distributed, use the **show hardware pxf cpu queue** command.

---

## Time-Based ACLs

The Time-Based ACLs feature has the following restrictions:

- You can specify a time range for only IP extended access lists. Standard access lists are not supported.
- An ACE that refers to a non-existent time-range entry is considered active.
- You define time-based ACLs based on hours and minutes. You cannot specify seconds.

## Variable Bit Rate Non-Real Time Oversubscription

The VBR-nrt Oversubscription feature has the following restrictions:

### Congestion

- Due to congestion on the physical interface, the accuracy of priority queuing (PQ) and class-based weighted fair queuing (CBWFQ) on individual VCs degrades. For example, if you configure each of three queues at a distribution of 50, 30, and 20 percent, the actual distribution might be 45, 40, and 15 percent.
- The distribution of bandwidth for each VC might be less than expected based on the speed of the VC. Typically, low speed VCs are allocated the expected bandwidth while high speed VCs share the remaining bandwidth equally.
- The amount of bandwidth allocated for the PQ or latency might be less than expected.

### Oversubscription Feature

- Oversubscription of the ATM interfaces is off by default. Oversubscription of the tunnels (the number and bandwidth of VCs that can be in a tunnel) is on by default and is not subject to any oversubscription factor. Oversubscription of the tunnels cannot be adjusted or turned off.
- Use the **atm over-subscription-factor** command to enable the oversubscription feature for a particular interface or tunnel. Do *not* use the **atm oversubscribe** command to enable oversubscription, as this can cause undesirable results.
- It is recommended that the **atm over-subscription-factor** command be applied to all ports of an ATM line card. This command controls the allocation of resources that are managed on a line card. Enabling oversubscription on one port alone could result in other ports taking up more resources than they were supposed to use. This could result in starving other ports for resources, which could cause VC creation to fail.

## WRED with Queue Limit

The WRED with Queue Limit has the following restrictions:

- The Cisco 10000 series router supports the configuration of 131,072 queues. The router reserves 255 queues for high speed interfaces. Any link that has a speed greater than 622 Mbps is classified as a high speed interface.
- You can configure a maximum of 29 queues per link.
- The queue limits that you can configure on a high speed interface range from 128 to 65,536 packets and on a low speed interface the queue limits range from 8 to 4,096 packets.

## Important Notes

This section provides important information about the following topics:

- [Configuring the aaa new-model Command, page 19](#)
- [Provisioning for Scaling, page 20](#)
- [Enhancing Scalability of Per-User Configurations, page 21](#)
- [Inserting a New Line Card, page 22](#)
- [Multilink PPP, page 22](#)

## Configuring the aaa new-model Command

The **aaa new-model** command is disabled by default on the Cisco 10000 series router. In previous releases, the default configuration did not appear in the running configuration file. However, in Cisco IOS Release 12.3(7)XI1 or later releases, the running configuration file now includes the **no aaa new-model** command. This is an intentional change in behavior for this command and is the first step in a three-step process to change the default configuration to **aaa new-model**.



Note

This change in behavior differs from Cisco IOS software, which typically does not include default configurations in the running configuration file.

For example, when you enter the **show running-config** command, **no aaa new-model** appears in the configuration if either of the following conditions previously occurred:

- You did not configure the **aaa new-model** command on the router and instead accepted the default configuration of the file: **no aaa new-model**.
- You entered the **no aaa new-model** command to remove the previously configured **aaa new-model** command.

## Provisioning for Scaling

The following configuration parameters enhance scalability on the Cisco 10000 series router:

- [PPPoA Sessions with IP QoS Static Routes, page 20](#)
- [AAA Authentication on the NME Port, page 20](#)
- [Call Admission Control, page 20](#)

To configure the Cisco 10000 series router for high scalability, be sure to configure the configuration parameters as described in the sections that follow.

For more information, refer to the [Cisco 10000 Series Broadband Aggregation and Leased-Line Configuration Guide](#).

### PPPoA Sessions with IP QoS Static Routes

To scale to 32,000 PPPoA sessions with IP QoS enabled, you must limit the number of IP QoS static routes to 4,000 unidirectional QoS static routes.

### AAA Authentication on the NME Port

If you use AAA authentication on the NME port, set both the in and out interface hold queues to 4096. For example:

```
Router(config)# int fa 0/0/0
Router(config-if)# hold-queue 4096 in
Router(config-if)# hold-queue 4096 out
```

### Call Admission Control

We recommend that you set the Call Admission Control (CAC) to a maximum of 95. For example:

```
Router(config)# call admission limit 95
```

## 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** 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"
```

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.

Whenever you specify a VRF in a user profile, but you do not configure the VRF on the Cisco 10000 series router, in Cisco IOS Release 12.2(15)BX, the router accepted 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 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 1](#) shows how to redefine the VRF named *newyork* using the `ip:vrf-id` VSA.

### *Example 1*     *Redefining VRF Configurations*

Change:

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

To:

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

[Example 2](#) shows how to redefine the Loopback 0 interface using the `ip:ip-unnumbered` VSA.

### *Example 2*     *Redefining IP Unnumbered Interfaces*

Change:

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

To:

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

## Inserting a New Line Card

Unlike other Cisco routers, if you insert a new or different line card into a Cisco 10000 series router chassis slot that previously had a line card installed, the line card initially reports that it is administratively up.

## Multilink PPP

Multilink PPP (MLPPP) is not supported on Cisco IOS Release 12.3(7)XI1.

# Open Caveats—Cisco IOS Release 12.3(7)XI1

Table 2 describes Open Caveats in Cisco IOS Release 12.3(7)XI1.

**Table 2** Open Caveats in Cisco IOS Release 12.3(7)XI1

Caveat	Description
CSCdt94857	<p>High impact commands or commands used in high scaling environments impact scaling by increasing CPU cycles, increasing boot time, and decreasing control plane run-time efficiency.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCdy19642	<p>Performance counters under the VT1.5, T3, and VT2 controllers for DS1/E1 are not getting updated/displayed correctly. On inserting the CRC errors in different ways to generate various events which can be used to count errors at the T1/E1 levels under either VT1.5 T3 or VT2 controllers, the counters are not getting updated correctly.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCdy45049	<p>When scaling over 3000 serial interfaces, line rate traffic may not be achieved. This problem occurs when thousands of serial interfaces (PPP or HDLC) are used on the port and line rate traffic is sent through all interfaces.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCdz40002	<p>When you remove APS and then re-activate it, traffic convergence after an APS switchover takes longer than 2 seconds.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCdz83304	<p>A T3 link on a 4-port channelized OC-3 line card may not come up under Synchronous Digital Hierarchy (SDH) framing. This problem occurs when the 4-port channelized OC-3 line card interoperates with third-party vendor test equipment.</p> <p><b>Workaround:</b> Enter the <b>shutdown</b> controller configuration command followed by the <b>no shutdown</b> controller configuration command on the AU-3 controller that contains the T3 link.</p>
CSCea63115	<p>When you enter the <b>redundancy force-failover main-cpu</b> privileged EXEC command on a router that is configured with two Performance Routing Engines (PREs), an automatic protection system (APS) switchover occurs on OC-12 Packet-over-SONET (POS) line cards, which is incorrect behavior.</p> <p>This problem occurs when APS is configured on OC-12 POS line cards in two different Cisco 10000 series routers that are connected back-to-back and you enter the following sequence of commands:</p> <ol style="list-style-type: none"> <li>1. Enter the <b>aps force pos slot/subslot/port from working</b> interface configuration command on both routers.</li> <li>2. Enter the <b>show aps</b> EXEC command. The output displays the active channels for both routers.</li> <li>3. Enter the <b>redundancy force-failover main-cpu</b> privileged EXEC on one of the routers, causing an APS switchover to occur on this router.</li> </ol> <p><b>Workaround:</b> There is no workaround for this problem. However, when problem occurs, there is no loss of data.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCea63638	<p>When Automatic Protection Switching (APS) is enabled, if you issue the <b>hw-module reset</b> command on the primary APS slot, no change is observed because the router does not switch to the secondary APS slot. This problem occurs when the <b>hw-module reset</b> command is issued.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCea68229	<p>The traffic flow over multirouter automatic protection switching (MR- APS) connections can stop. This problem occurs under the following conditions:</p> <ul style="list-style-type: none"> <li>• MR-APS is enabled on both a Cisco 10000 series router and a Cisco ONS15454 platform.</li> <li>• The protect interface is configured on an interface of a 6-port OC-3 Packet-over-SONET (POS) line card in the Cisco 10000 series router.</li> <li>• You enter the <b>shutdown</b> interface configuration command followed by the <b>no shutdown</b> interface configuration command on the interface that is configured as the Protect Group Protocol link while the protect interface is active.</li> </ul> <p>The working interface that is configured for MR-APS on the Cisco ONS15454 platform should become active but fails to do so, causing the traffic flow over MR-APS connections to stop.</p> <p><b>Workaround:</b> Enter the <b>shutdown interface</b> configuration command followed by the <b>no shutdown interface</b> configuration command on the protect interface on the Cisco 10000 series router.</p>
CSCec13372	<p>The router can generate wrong or misleading sub-pool or global pool flooding messages when up or down thresholds for MPLS TE resource availability (bandwidth) are crossed. The configured thresholds for MPLS TE resource availability are crossed when defining bandwidth on the MPLS tunnel interface reserved on the physical interface/subinterface.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCec16481	<p>A Cisco device running Internetwork Operating System (IOS) and enabled for the Open Shortest Path First (OSPF) Protocol is vulnerable to a Denial of Service (DoS) attack from a malformed OSPF packet. The OSPF protocol is not enabled by default.</p> <p>The vulnerability is only present in IOS release trains based on 12.0S, 12.2, and 12.3. Releases based on 12.0, 12.1 mainlines and all IOS images prior to 12.0 are not affected. Refer to the Security Advisory for a complete list of affected release trains.</p> <p><b>Workaround:</b> Further details and the workarounds to mitigate the effects are explained in the Security Advisory which is available at the following URL: <a href="http://www.cisco.com/warp/public/707/cisco-sa-20040818-ospf.shtml">http://www.cisco.com/warp/public/707/cisco-sa-20040818-ospf.shtml</a>.</p>
CSCec37207	<p>On Cisco 10000 series routers running in PTA mode, PPPoEoA sessions using bandwidth queues drop packets if a priority queue is also configured in the policy map. When there is traffic sent to priority queue, all other queues can drop packets below line rate if the traffic consists of small packets.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCec38308	<p>SSG only supports one class attribute rather than several of them, although a RADIUS client is supposed to put all class attributes that it receives in Access-Accept messages into Accounting-Request messages that it sends for a session. (See RFC2865/2866.) This problem occurs on a Cisco platform that is configured as an SSG.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCec42315	<p>When scaling to 12000 Frame Relay DLCI interfaces, line rate traffic may not be achieved. This problem occurs when thousands of Frame Relay DLCIs are used on the port and line rate traffic is sent through all interfaces.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCec42451	<p>The RIP routing protocol does not function properly over VLAN interfaces with IP unnumbered.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCec43937	<p>When you run hierarchical shaping, a small number of UPC violations can occur on a DS3 ATM interface on an LS1010 that is connected to a LAC. The traffic shaped by the LAC is slightly above the tunnel PCR. The number of violations does not depend on the tunnel oversubscription but seems to be proportional to the tunnel PCR.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCec48111	<p>When sending 64 byte packets through 300 serial interfaces or more, line rate traffic may not be achieved. This problem occurs with 64 byte packets and a large number of interfaces.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCec62993	<p>The following error message appears with <b>debug vpdn l2x-errors</b> enabled when users are trying to connect to an LNS and 'user@domain' is the username and domain being used to connect with:</p> <pre>vpn_set_ppp_remote_name: Error inserting username, user@domain, into String DB</pre> <p>This problem occurs in a DSL environment with L2TP. The message appears on the LNS. It does not affect the functioning of the router.</p> <p><b>Workaround:</b> None needed. This message is only seen when debugging vpdn l2x-errors.</p>
CSCec78662	<p>Time-based ACLs do not work when placed inside a policy map. The results of placing a time-based ACL inside a policy map is either that the time-based rules are always active or inactive.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCec80927	<p>Call setup rate slower is for a particular configuration running on a 12.3(6)TX image compared with 12.2(16)BX. If the <b>mtu</b> command is added to the vtemplate for sessions, the command processing for the command takes significantly longer on a 12.3(6)TX image as compared to a 12.3(16)BX image.</p> <p><b>Workaround:</b> Remove the <b>mtu</b> command from the vtemplate configuration.</p>

**Table 2** Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCec85628	<p>Outgoing traffic is above VP speed on an 8e3d3atm line card. For this problem to occur, the total SCR value of all VBR-nrt VCs in a VP is above 80 percent of the VPs PCR value but still smaller than the total VP bandwidth (PCR). All the VCs should be overdriven by outgoing traffic.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed03248	<p>The CLI error "IP address is already defined as an interface" appears when the address is not used anywhere in the running configuration. The error occurs when the IP address was used by a serial interface and the interface was removed or unconfigured from the system.</p> <p><b>Workaround:</b> Use the <b>no ip address</b> command before removing a serial interface or use a different IP address (if possible).</p>
CSCed17570	<p>When using thousands of QoS queues with WRED configured in each queue, a traceback message can appear when you execute the <b>microcode reload pxf</b> command. The traceback message appears only when thousands of PXF queues are configured with random-detect enabled and the <b>microcode reload pxf</b> command is issued.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed20626	<p>'Exec' process CPUHOG occurs. This is caused by the <b>dir all</b> command, probably due to the attempted accesses to the secondary's PCMCIA slots.</p> <p><b>Workaround:</b> None needed. The router continues to function, but the console is unusable for a short while(10-30 seconds). Alternatively, use the <b>dir &lt;device&gt;</b> commands only for known good device names.</p>
CSCed29494	<p>The maximum queue limit for a low -peed link is 4096, but the system allows you to set the queue limit to 8192.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed54867	<p>The input service policy does not match traffic as shown by the <b>show policy-map interface</b> command if there is no action associated for that class.</p> <p><b>Workaround:</b> The workaround is to set up an action such as "set" or "police".</p>
CSCed59185	<p>When you apply the following example configuration to an output interface that is MPLS enabled, and send traffic from the CPU of the local router (ping other routers or hosts), the traffic is not policed by the policy map.</p> <pre> Policy Map exp2cos Class exp0 set cos 1 Class exp1 police 104000 5000 150800 conform-action transmit exceed-action drop violate-action drop Class exp2                     </pre> <p>This problem only affects the traffic from the router CPU, and does not affect traffic passing through the router.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCed62503	<p>When you apply a policy map to a tunnel interface on a router configured with a PRE2 processor, a traceback message appears. This problem occurs when the policy map is applied to a tunnel interface.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed65349	<p>When you configure 2000 PPP interfaces, traffic does not reach 99percent of the line rate after performing 4 HA RPR switchovers. The traffic rates keep fluctuating.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed68868	<p>A traceback message appears when you unconfigure the spoke PE router configured for half-duplex VRF over PPPoE. This problem occurs with 32k PPPoE sessions and 40 spoke VRF, therefore scaling to high values.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed70202	<p>A traceback message appears when you unconfigure the hub PE router configured for half-duplex VRF. This problem occurs with 32k sessions, therefore scaling to high values.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed71107	<p>When 2 time-based ACLs are configured to deny traffic at the same time and are applied to different interfaces, one of the ACLs fails to work properly.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed72023	<p>Excessive CPU utilization is detected for 5 minutes after unconfiguring half-duplex VRF with a large number of PPPoE user sessions. This problem occurs with 32k PPPoE sessions, therefore scaling to high values.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed72338	<p>The system allows non-nested queuing policy maps to be applied via the <b>frame-relay map-class</b> command on Frame Relay main interfaces and subinterfaces; it should not allow such policy maps to be configured.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed86371	<p>Automation Protection Switching (APS) active state does not stay with the lowest active odd port after a PRE switchover.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCed88782	<p>The secondary port does not go to a working state during a signal degrade of the primary port using threshold SON ERR RAT 1e-6.</p> <p><b>Workaround:</b> Set the BIP threshold to 6; do not set the BIP threshold to 7.</p>
CSCed94283	<p>When scaling to 128K PXF queues with a policer configured on all queues, the router's CPU utilization runs at approximately 48 percent of capacity. This problem occurs with 128K PXF queues configured on ATM interfaces.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee02536	<p>When configuring MPLS Layer 3 VPN, the PXF CEF/FIB table can hold up to 4085 VRFs, although it is designed to hold 4095 VRFs. If more than 4085 VRFs are configured, 10 of those VRFs do not have an entry in PXF CEF/FIB table, so traffic is not forwarded in those 10 VRFs.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

**Table 2** Open Caveats in Cisco IOS Release 12.3(7)XI1 (continued)

Caveat	Description
CSCee03801	<p>After you issue the <b>clear ip bgp *</b> command, a Cisco 10000 series router takes longer than 30 minutes to achieve convergence. eBGP sessions between PE and CE routers can go up and down multiple times, and the IGP routing protocol and LDP session can also go down and up again.</p> <p>These problems occur under the following conditions:</p> <ol style="list-style-type: none"> <li>1. 4095 VRFs are configured on a router</li> <li>2. 500 eBGP sessions are established between the router (PE) and CE routers</li> <li>3. 540 VRF routes per VRF in the 500 VRFs that are running eBGP between PE and CE routers</li> <li>4. 40 VRF routes per VRF in the rest of 3595 VRF</li> </ol> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee06089	<p>When you apply a nested policy map using the <b>bandwidth</b> command in the child policy map to a POS OC48 interface, PXF stops responding. This problem occurs when you allocate a small amount of bandwidth, and it only occurs on POS OC48 interfaces.</p> <p><b>Workaround:</b> Allocate more bandwidth in the child policy map.</p>
CSCee07016	<p>Output rate limited traffic on an L2TP VAI is less than expected or less than in previous releases. The output policing algorithm is different in the 12.3(7)XI image than in previous images. The changes to the algorithm increased the overhead included in the policed bps rate.</p> <p><b>Workaround:</b> Account for the additional overhead of the output policer. The configured police bps rate when applied to an L2TP Virtual-Access interface includes 40 bytes of per packet overhead: L2TP (8 bytes), PPP (4 bytes), outer IP (20 bytes), and UDP (8 bytes).</p>
CSCee08859	<p>Backward compatibility of RADIUS aaa tunnel authorization fails because the expected IP local address pool is not being used to assign addresses to the VAIs.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee14864	<p>Policing under a created queue, when attached at an MLP interface, accounts for only 2 bytes of the L2 header, so that policing is done at a higher rate than configured. This can cause a problem with priority queue CBWFQ functionality because the priority queue is configured with policing and its dequeue rate can be higher than intended.</p> <p><b>Workaround:</b> Do not configure policing under a created queue.</p>
CSCee15674	<p>When broadband PTA is configured with 114,000 queues, executing the <b>microcode reload pxf</b> command causes the ATM interface to display a big number of total output drops.</p> <p><b>Workaround:</b> Clear the counters.</p>
CSCee20418	<p>If the you change the amount of intercepted streams from 8 to 2 streams, the wrong amount of packets is intercepted. This occurs in Lawful Interception scenarios.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCee25615	<p>This problem occurs when almost all the system resources (vcci) are in use, after an OIR (slot reset) is issued, and in the OC3 ATM line card. The reason it occurs in the OC3 ATM line card is that it happens in an ATM line card with multiple ports. The symptom is that all the sessions in the same port stop passing traffic after OIR.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee27630	<p>A low-bandwidth class can be allocated more than its share of bandwidth at the expense of a high-bandwidth class. This problem occurs when the ratio of the configured bandwidths between two data classes is high (8:1 or higher) and when there is a priority class that receives traffic at (at least) 20 percent of the line rate. The traffic that is received by the data classes should be in the ratio of the configured bandwidths.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee36429	<p>If you configure or modify random detect parameters for a policy map that is already applied to an interface, the router can display an error message. This problem occurs if the modified policy map is part of a hierarchical policy map configuration.</p> <p><b>Workaround:</b> Remove the hierarchical policy map from the interface before proceeding with the random detect modifications.</p>
CSCee41961	<p>The following error messages appears:</p> <pre>%C10K-2-BADSRRCNUM: Invalid resource number from PXF (15). (PLEASE REPORT THIS!) PXF divert PAK (len 62, rsrc 15, col 0, cause 0) contents: 22 3E 00 01 03 3E 45 20 00</pre> <p>PXF also crashes. This problem occurs when a 2 level output policy map is applied on an OC48 POS interface, the child policy map has WRED configured, and traffic is being sent out the POS interface.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee42746	<p>When using multiple intercepts in Lawful Intercept mode, the MIB information is not completely cleared after intercepts are cleared from SNMP. This problem occurs when 35 or more streams are intercepted at the same time.</p> <p><b>Workaround:</b> Use Cisco IOS to delete the stream that was not deleted by SNMP.</p>
CSCee44273	<p>The <b>show activity</b> line card debug command shows the VC configuration from the perspective of the line card, but the autovc information is not shown. Also, after you delete or create an autovc, the counter is inaccurate.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee44988	<p>In 12.3(7)X1, when <b>vpdn authen-before-forward</b> and <b>aaa authorization network default local group radius</b> are configured, a second unwanted access request for authorization is sent.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee45306	<p>With 40 or more intercept streams in Lawful Intercept mode, the LI engine fails to intercept correctly for UDP traffic. This problem occurs when 40 or more streams are intercepted at the same time.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCee45378	<p>When intercepting streams at 5 Mbps or above in Lawful Intercept, the router CPU runs at about 78 percent of capacity. This problem occurs when 35 or more streams are intercepted at the same time.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee50060	<p>A Cisco 10000 series router with PPPoA VCs can, under abnormal conditions (such as a denial-of-service attack involving the sending of PPPoA data packets before the PPPoA session is up), experience heavy RP CPU use. The router with PPPoA VCs can forward PPPoA data packets for non-existent sessions.</p> <p>This problem occurs when PPPoA data traffic is sent before the session reaches the PTA forwarded state. A normal PPPoA client does not send traffic before the session is up.</p> <p><b>Workaround:</b> Configure RPF on all ATM subinterfaces containing PPPoA sessions. The subinterface should have an RPF check in addition to using an RPF check in the virtual template. Configuring RPF on the subinterface forces all PPPoA data traffic to be dropped by the PXF before the session reaches the PTA forward state.</p>
CSCee54408	<p>When the 1choc12 line card uses SDH framing, the Path Trace Buffer is unstable for au3 mode. This problem occurs only with SDH framing; the Path Trace Buffer is stable with SONET framing.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee54426	<p>When the 1choc12 line card uses SDH framing, the J1 Path trace message is not received. This problem occurs only with SDH framing; The J1 Path Trace message is received when SONET framing is used.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee54473	<p>A loss of frame (LOF) alarm appears for a T1 when framing SF is configured on both ends. This problem occurs when you configure <b>T1 1 framing sf</b> under AU-3 on a 1 port channelized OC12 line card.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee54971	<p>The <b>show policy-map interface</b> command output does not display the layer 2 frame size correctly. The actual output policing rate is 6.6 percent higher than the configured policing rate on gigabit Ethernet and POS OC48 interfaces. The problem occurs when a <b>police</b> command is configured in a policy map, and the policy map is applied to a gigabit Ethernet or POS OC48 interface as an output policy map.</p> <p><b>Workaround:</b> Use shaping instead of policing.</p>
CSCee55828	<p>You cannot configure <b>t1 1 framing esf</b> and <b>t1 loopback remote</b> at the same time on a 1-port channelized OC12 line card. This problem occurs when you configure <b>t1 1 framing esf</b> under an AU-4 on a 1-port channelized OC12 line card.</p> <p><b>Workaround:</b> Configure <b>t1 1 framing esf</b> without the loopback configured for the T1.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCee57219	<p>The <b>set cos</b> command in an output policy map applied to a VLAN subinterface does not work if the outgoing traffic is MPLS packets (with MPLS labels). The problem occurs when outgoing traffic is MPLS packets.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee57357	<p>When scaling Frame Relay DLCIs on routers running IOS version 12.3(7)XI, traceback messages can appear on the console when bringing up the high number of DLCIs. This problem occurs when there are more than 3000 DLCIs on the interface.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee58454	<p>On a router running 12.3(7)XI, if the LAC tries to redirect a call to the bid-winning LNS and fails after three attempts, a new RADIUS disconnect cause code with the value as 608 is not being sent to RADIUS by the LAC.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee60038	<p>When a proxy service profile defined with V &amp; X attributes is configured locally on the router, which is enabled to run SSG, an SSG host cannot activate the service it has been subscribed to.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee61067	<p>In 2-level policy map configurations using a parent shaper, the shaped traffic rate might not be within plus or minus 1 percent of the configured value. This problem occurs with certain parent shaper values and mostly small packet sizes.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee61353	<p>AAA server counters, private and global, are not cleared. The AAA server counters are not reset to 0 (zero) when the <b>clear aaa counters servers radius all</b> command or <b>clear aaa counters servers all</b> command is issued.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee61485	<p>Several PIM-related messages appear on the console when you remove, then re-apply a PIM configuration on the interface. This problem occurs when the removal and re-application of the configuration is done in a rapid manner.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee61502	<p>When configuring an MLPPP interface on a redundant system, the standby PRE adds the <b>no ip route-cache cef</b> interface command to multilink interfaces. This additional line causes the system to generate the following error when the new standby PRE is reloaded:</p> <pre>May 19 13:20:47.222 EDT: %REDUNDANCY-3-CONFIG_SYNC: Active and Standby bulk configuration out of sync</pre> <p><b>Workaround:</b> Remove the <b>no ip route-cache cef</b> command from each multilink interface.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCee62159	<p>Actual output and expected output for packet 1 does not match at nibble 8. This packet (packet_no 1, fragment_no : 1) is received in the wrong order. Other packets are also received in the wrong order. This problem occurs with the bootflash:c10k2-p11-mz.v123_7_xi_throttle.040510 image and the test is passed with Feb17 bba image.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee63156	<p>Traceback messages appear when running regression tests on the 1choc12 line card at T1 speed with sdh_au4_hdlc framing. These messages do not seem to affect the functionality of the line card.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee63564	<p>The output queue counter is wrong when executing the <b>show policy-map interface X/Y/Z</b> command. This problem occurs when policing is attached at the OC48 POS interface's output.</p> <p>Workaround: Use the <b>show int</b> command counter for packet output.</p>
CSCee63636	<p>MPLS:Traceroute does not show Labels being switched-propagate-ttl ON.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee66066	<p>BERT testing over a clear channel DS3 interface for the 1CHOC12 line card fails as a result of the DS3 interface, which remains in a DOWN state.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee66091	<p>During SNMP polling of the AAA Server MIB, the casDeadCount variable can cause a CPU hog on the router. This problem occurs with a large number of RBE interfaces (16K) and bi-directional traffic running.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee66183	<p>If you reload the peer router (containing 768CGs MLPPP configuration in startup-config) while the traffic is flowing over the configured bundles, upon reload the configured interface remains in an up/down state, and bundles in a down/down state. Additionally, the router log indicates traceback and ipc failure messages.</p> <p><b>Workaround:</b> Clean all the MLP-related configurations, stop the traffic, and copy the configuration back from a tftp server.</p>
CSCee66314	<p>In Lawful Intercept mode a traceback message might appear on the Intercept Access Point (IAP) router when the interface to the mediation router is shut down. This problem occurs when traffic is sent through the IAP and interception is turned on.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee68404	<p>If a PRE2 is in the early process of booting up, sometimes the SEND-BREAK character sequence can cause the router to crash instead of gracefully dropping back into ROMMON. This problem occurs when the PRE2 is in the early stages of the boot process and the SEND-BREAK is issued. If the PRE2 is already booted up, this is not an issue.</p> <p><b>Workaround:</b> To gracefully drop the PRE2 into ROMMON, if the configuration register is set to accept SEND-BREAK, wait until the PRE2 is fully booted.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCee68480	<p>Priority queue latency can exceed the threshold of 2MTU+6msec. This problem occurs when more than 3 queues are configured on a interface, in addition to the priority queue.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee69274	<p>The throughput is low as a percentage of the maximum available. The problem occurs only with PVCs configured within a PVP and for traffic consisting of packets that are less than ~128 bytes in size. This problem affects situations where all traffic consists of packets less than ~128 bytes and there are PVCs configured within a PVP.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee71816	<p>When traffic engineering tunnels are active, issuing the <b>show pxf cpu statistics drop tunnel</b> command causes the following traceback message to appear:</p> <pre>*May 31 20:12:04.947: %GENERAL-3-EREVENT: pxf_drop_interface: No c10k_tt_hwdb -Traceback= 60D8F458 60D8BD98 60D8D9BC 603B322C 601404D0 603CD1B4 6045DD88 6045DD6C PE-1#</pre> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee72919	<p>AAA accounting records for a PPPoA session terminated on a Cisco 10000 series router in a PTA fashion shows repeated entries for the Framed-Route attribute (attribute 22).</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee72931	<p>When a PPPoA session is cleared on the PTA router using the <b>clear pppatm interface ATM X/Y/Z.A</b> command or the <b>clear int virtual-access</b> command, the accounting stop record does not display the Octet and Packet counters. This problem occurs only when the session is cleared on the PTA router. If the user disconnects the session, the counters are displayed correctly.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee73535	<p>In Lawful Intercept mode, the intercept stream sometimes is not deleted after the configured time to live (TTL) has expired. This problem occurs if the TTL value is changed while the intercept is active.</p> <p><b>Workaround:</b> Do not change the intercept TTL while it is active.</p>

**Table 2** Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
<p><b>CSCee78997</b></p>	<p>If changing the MTU of a serial interface on a ch-oc12 or ch24E1T1 line card, the optional sizes provided is in the range 64-17940. This is not correct. The correct range should be 64-9108. If you change the MTU to a value greater than 9108, the following error appears:</p> <pre> %GENERAL-3-EREVENT: c10k_ttc_m_icb_update: attempt to set max_mtu to 9320 , overridden to 9216 -Traceback= 60D24E68 60D28294 6011B36C 60423510 604241A0 603B3110 601405A8 603CD098 6045DD30 6045DD14                     </pre> <p>example:</p> <pre> router(config-if)#int s3/0/23:0 router(config-if)#mtu ? &lt;64-17940&gt; MTU size in bytes  router(config-if)#mtu 9200 router(config-if)# *Jun  8 16:45:50.950 EDT: %GENERAL-3-EREVENT: c10k_ttc_m_icb_update: attempt to set max_mtu to 9320 , overridden to 9216 -Traceback= 60D24E68 60D28294 6011B36C 60423510 604241A0 603B3110 601405A8 603CD098 6045DD30 6045DD14 router(config-if)# router(config-if)#                     </pre> <p><b>Workaround:</b> Do not set the MTU to a value greater than 9108.</p>
<p><b>CSCee79228</b></p>	<p>On a Cisco 10000 series router configured as PTA device, some very small amount of memory is not released as PPP sessions are brought up and torn down. This problem can cause the router to run out of memory after a long period of time.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
<p><b>CSCee81270</b></p>	<p>When a source sends packets to a destination under the TCP protocol, the destination sends an echo response back to the sender. With the intercepting router configured to intercept "all", those echo packets should also be picked off. This does not occur.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCee82413	<p>The following errors appear when setting up 31,500 PPPoX sessions on a PRE2 as LAC with DBS enabled:</p> <pre>Jun 9 12:34:48.520: %C10KATM-3-DBS: C10K internal DBS error, DBS: modify() failure: validation of params unsuccessful(1) ATM3/0/0 2277 1/2377 -Traceback= 600878C0 60162C30 60C6C344 60C6C4D8 60C68348 60C6ACE8 60C6B2EC 60C63094 Jun 9 12:34:48.524: %ATM-3-FAILMODIFYVC: ATM failed to modify VC(VCD=2277, VPI=1, VCI=2377) on Interface ATM3/0/0, (Cause of the failure: Failed to have the driver to modify the VC) Jun 9 12:34:48.524: %C10KATM-3-DBS: C10K internal DBS error, DBS: modify() failure: validation of params unsuccessful(1) ATM3/0/0 2277 1/2377 -Traceback= 600878C0 60162C30 60C6C344 60C6C508 60C6989C 60C60F00 60C6B1A8 60C6B2EC 60C63 094 Jun 9 12:34:48.524: %C10KATM-3-DBS: C10K internal DBS error, DBS: modify() failure: validation of params unsuccessful(1) ATM3/0/0 2279 1/2379</pre> <p>The result is that QoS parameters that should be derived from RADIUS via DBS are not set for some ATM VCs.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCee84775	<p>MPLS routes are untagged if you configure virtual private network (VPN) parameters before configuring MPLS.</p> <p>If you configure MPLS before VPN, the router is correct:</p> <pre>Router#show mpls forward Local Outgoing Prefix Bytes tag Outgoing Next Hop tag tag or VC or Tunnel Id switched interface 17 Pop tag 192.168.28.11/32 0 Gi5/0/0.20 192.168.201.14</pre> <p>If you configure VPN before MPLS, the routes are UNTAGGED:</p> <pre>Router#show mpls forward 17 Untagged 192.168.28.11/32 0 Gi5/0/0.20 192.168.201.14</pre> <p><b>Workaround:</b> Configure MPLS before configuring the VPNs.</p>
CSCee85029	<p>The class attribute is not sent in prepaid authorization requests for PPP users. This problem occurs in all releases after 12.3(2)T.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCee90736	<p>The ATM line card can experience a crash when there is a lot of change activity going on (VCs being added and deleted). On the IOS console messages similar to the following appear:</p> <pre>#sh log Jun 16 15:19:00.423 BST: %SYS-5-CONFIG_I: Configured from console by provuser on vty2 (address deleted) Jun 16 15:56:56.961 BST: %IPCGRP-3-SYSCALL: System call for command 405 (slot3/0) : ipc_send_rpc_blocked failed (Cause: timeout) -Traceback= 6053D6BC 6053D994 6053DB50 60096EAC 60083294 60164CC8 60DB8FB4 60DB96F0 60DB9BE8 Jun 16 15:57:02.962 BST: %IPCGRP-3-SYSCALL: System call for command 401 (slot3/0) : ipc_send_rpc_blocked failed (Cause: timeout) -Traceback= 6053D6BC 6053D994 6053DB50 60096EAC 60083294 60164CC8 60DB8FB4 60DB96F0 60DB9BE8 Jun 16 15:57:03.962 BST: %IPCOIR-3-TIMEOUT: Timeout waiting for a response from slot 3/0. Jun 16 15:57:03.962 BST: %IPCOIR-2-CARD_UP_DOWN: Card in slot 3/0 is down. Notifying 4oc3atm-1 driver. Jun 16 15:57:03.962 BST: %IPCGRP-3-CMDOP: IPC command 401 (slot3/0): line card ipc is disabled - dropping non-blocking ipc command -Traceback= 6053D940 6053E410 Jun 16 15:57:05.970 BST: %LINK-3-UPDOWN: Interface ATM3/0/1, changed state to down Jun 16 15:57:08.362 BST: %LINK-3-UPDOWN: Interface ATM3/0/0, changed state to down Jun 16 15:57:08.786 BST: %LINEPROTO-5-UPDOWN: Line protocol on Interface ATM3/0/1, changed state to down Jun 16 15:57:09.362 BST: %LINEPROTO-5-UPDOWN: Line protocol on Interface ATM3/0/0, changed state to down Jun 16 15:57:14.123 BST: %IPCOIR-5-CARD_DETECTED: Card type 4oc3atm-1 (0x2D8) in slot 3/0 Jun 16 15:57:14.123 BST: %IPCOIR-5-CARD_LOADING: Loading card in slot 3/0 Jun 16 15:57:15.315 BST: %C10K-5-LC_NOTICE: Slot [3/0] 4oc3atm-1 Image Downloaded...Booting... Jun 16 15:57:37.124 BST: %IPCOIR-5-CARD_DETECTED: Card type 4oc3atm-1 (0x2D8) in slot 3/0 Jun 16 15:57:37.124 BST: %IPCOIR-2-CARD_UP_DOWN: Card in slot 3/0 is up. Notifying 4oc3atm-1 driver. Jun 16 15:57:50.161 BST: %LINK-3-UPDOWN: Interface ATM3/0/0, changed state to up Jun 16 15:57:50.421 BST: %LINK-3-UPDOWN: Interface ATM3/0/1, changed state to up Jun 16 15:57:51.201 BST: %LINEPROTO-5-UPDOWN: Line protocol on Interface ATM3/0/0, changed state to up Jun 16 15:57:51.421 BST: %LINEPROTO-5-UPDOWN: Line protocol on Interface ATM3/0/1, changed state to up Jun 16 16:19:04.759 BST: %IPCGRP-3-CMDOP: IPC command 405 (slot3/0): waiting for a keepalive -Traceback= 6053D6BC 6053D994 6053DB50 60096EAC 60083294 60166834 601637E8 60DB8E90 60DB9A14 60DB9BE8</pre> <p>(display text omitted)</p> <p>On the line card console the following appears:</p> <pre>#if-con 3/0 Connecting console for slot 3/0 Type "^C^C^C" or "if-quit" to end this session log dump ----- Start of console log ----- oc3atm-3/0&gt; FPGA: fatal FPGA interrupt encountered (0x00000010) ASSERT Failed: in ../src-c10k-atm/ocXatm_fpga.c::fpga_int_handler() L1407 backtrace: 8000CCA4 80008334 8003A754 80007880</pre> <p><b>Workaround:</b> There is no workaround for this problem. The line card will reboot and the system will recover without user intervention.</p>
CSCee90765	<p>On Cisco 10000 series routers used as L2TP managed LNS, a CPUHOG message can appear on the console if several thousand sessions are active. This problem occurs when you scale the managed LNS to high numbers.</p> <p><b>Workaround:</b> There is no workaround for this problem</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)XI1 (continued)

Caveat	Description
CSCee90904	In the presence of a large number of static routes (16k - 32k), line card flap/ router reload/OIR cause high CPU usage for a long period of time. <b>Workaround:</b> There is no workaround for this problem.
CSCee96582	With broadband multipoint 31,500 PVCs with 30k sessions up, 126k queues, and you add a class with the <b>set</b> command in an output policy map on the fly, the router hangs for a long time then crashes. This problem occurs with broadband multipoint PVCs with 30k sessions up, 120k queues, then you add a class with the <b>set</b> command in a policy map on the fly. <b>Workaround:</b> There is no workaround for this problem. With a large number of sessions and queue scaling, avoid changing policy map on the fly.
CSCef00499	With broadband queue scaling with an input police policy on virtual-template and output CBWFQ policy on ATM 31,500 subinterfaces is configured, and then the output policy is removed, an XCM access error message occurs continuously. This condition occurs with broadband queue scaling with input and output policies configured, and then the output policy is subsequently removed. <b>Workaround:</b> There is no workaround for this problem.
CSCef00808	The <b>show pxf cpu stat security</b> command shows incorrect statistics when Legal Intercept is configured along with time-based or regular access lists. This problem occurs only if Legal Intercept and access lists are configured and are interoperating. <b>Workaround:</b> There is no workaround for this problem.
CSCef01772	Cisco IOS can crash when receiving a malformed PPPoE packet. Without having a PPPoE configuration on the router, the router can crash if it receives a PPPoE session Packet (0x8864) with Session ID = 0. This problem can occur on a Cisco 10000 series router running the 12.3(07)XI image. <b>Workaround:</b> The workaround is to configure bba-group before sending any PPPoE session packet. This prevents a crash.
CSCef03281	Error messages that occur during bandwidth oversubscription with ATM PVPs do not contain information for which interfaces or features are oversubscribed. Messages indicate the peak rate exceeding the available bandwidth Link oversubscribed by 92240 kbps. The error messages occur when the sum of ATM PVP Peak Cell Rates exceed the interface bandwidth. <b>Workaround:</b> There is no workaround for this problem.
CSCef04501	In some cases assertion failures occur when modifying a VC in PVC range, such as adding a VC class under a VC that is a part of a range. <b>Workaround:</b> There is no workaround for this problem.
CSCef08967	The WRED sampling frequency is too slow, which can cause jitter for the overall algorithm. <b>Workaround:</b> There is no workaround for this problem.

**Table 2** Open Caveats in Cisco IOS Release 12.3(7)XI1 (continued)

Caveat	Description
CSCef09119	<p>With broadband PTA 128k queue with input and output policy map, removing the input policy from Virtual-Template causes a CPUHOG traceback message. This occurs when configuring 31.5k ATM subinterfaces with output CBWFQ policy, and input police policy in Virtual-Template, bringing up 30k PPPoE sessions, and removing the input policy map.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef09165	<p>With SSG configured, VPDN parameters are locally provisioned but VPDN tunnels are not established between the LAC and the LNS. SSG VPDN services are not working.</p> <p><b>Workaround:</b> Configure the <b>aaa authorization network default group radius</b> command.</p>
CSCef11074	<p>The following errors appear when starting up 31,500 PPPoX sessions at 5 CPS on a PRE2 as PTA with SSG auto-logon configured:</p> <pre>*Jul 7 10:05:38.602: SSG-CTL-ERR: Unable to add HostRoute in CEF table x.x.x.x *Jul 7 10:05:38.602: SSG-CTL-ERR: host route addition failed *Jul 7 10:05:41.770: SSG-CTL-ERR: Unable to add HostRoute in CEF table x.x.x.x *Jul 7 10:05:41.770: SSG-CTL-ERR: host route addition failed</pre> <p>The results is that the Active HostObject Count in <b>sho ssg host</b> output does not match the ConnectionCount in <b>sh ssg serv</b> output. There is no upstream traffic on these lost connections.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef14249	<p>When sending traffic with 1024 byte large size packets over 120k queues with 80 percent oc12atm line rate, traffic drops 10 percent due to buffer_low packet drop. This problem occurs when 120k queue scaling is configured with only large packet size traffic.</p> <p><b>Workaround:</b> There is no workaround for this problem. Send traffic with mixed size packets, tending to small packets.</p>
CSCef14961	<p>The "%ATM-3-FAILCREATEVC: ATM failed to create VC" and "Attempting to over-subscribe tunnel bandwidth" messages canbe erroneously logged upon deletion and addition of VCs and VPs (Hierarchical Traffic Shaping).</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef15141	<p>On Cisco 10000 routers running 12.3(7)XI, the Priority Queue latency values (in milliseconds) is higher than 2*MTU + 6ms on 4Mbps and 8Mbps subrates of the 8e3ds3 line card.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef17789	<p>A performance degradation can occur when VP and VC shaping are configured on the same interface. Output drops can occur at the SAR level and affect the established PPPoX sessions.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCef19196	<p>The PVC 'weight' configurable parameter is not visible under the <b>show atm pvc</b> or <b>show atm vc</b> command. As this parameter is made configurable, a 'show' command should display the default or non-default value.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef20523	<p>PPPoEoA sessions using CBWFQ experience BQ drops. In some cases, when aggregate traffic is near the VC rate, the BQ tail drops packets. This problem appears with low bandwidth VCs, in this case 196 kbps.</p> <p><b>Workaround:</b> Changing the queue-limit via the policy map and/or the VC queue depth will improve the result.</p>
CSCef20554	<p>The CPU stays at 100 percent utilization after a session has been cleared. This problem occurs with 1000 VLANs and 32k sessions and 32 sessions per VLAN, clearing sessions while traffic is being sent over the sessions.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef22815	<p>Some memory can be lost when you configure and remove several Multilink PPP interfaces. This problem can lead to buffer exhaustion over time and require a reload of the router.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef24008	<p>When using a 4choc3 line card and 300 or more VT T1 interfaces are configured with PPP encapsulation, some T1 links do not achieve full traffic line rate. This problem occurs when all 300+ interfaces are sending traffic at line rate concurrently.</p> <p><b>Workaround:</b> There is no workaround for this problem</p>
CSCef24338	<p>When a service policy configured with the <b>police percent</b> command is applied onto a Virtual-Access interface bound to a shaped PVC, the calculated police class bandwidth is not based on the PVC rate but on a 100 Mbps value.</p> <p><b>Workaround:</b> The workaround is to use absolute kbps bandwidth value with the <b>police</b> command.</p>
CSCef24551	<p>When running Automated Protection Switching (APS), the router can experience traffic loss after the <b>hw-module slot x reset</b> command is executed.</p> <p><b>Workaround:</b> Avoid executing <b>hw-module slot x reset</b>.</p>
CSCef24564	<p>Accounting of input packets/bytes is not happening correctly. When a client is connected to SSG and further on to a service linked via a gigabit Ethernet uplink subinterface, on an extended ping from client to service, the accounting of input packets/bytes is erroneous. The same result is reflected in output of the RADIUS accounting logs.</p> <p><b>Workaround:</b> The workaround is to use gig main interface as SSG uplink</p>
CSCef24716	<p>A traceback message appears when trying to log in an rfc1483 SSG host to a service over a gigabit Ethernet uplink. With a gigabit Ethernet uplink interface to the services, when an rfc1483 SSG host tries to log in to the service, a traceback message appears.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCef25091	<p>On Cisco 10000 series routers running PPPoEoQinQ, if you add and then remove policy maps several times on the virtual template, a spurious memory access message can appear on the console:</p> <pre>Defaulted buffer is too big</pre> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef26366	<p>An ACL has no effect even though it is configured. In <b>show pxf cpu statistics security</b>, packets are neither denied nor permitted. In <b>show pxf cpu context</b>, there are no feedback packets. The ACL has to be split. In <b>show pxf cpu access-list security</b>, the value of the table column has to be greater than 1 to have a split ACL.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef26525	<p>When a router has over 100 BGP peers, traceback messages can appear on the console after an RPR+ switchover. The system recovers and normal activities are resumed afterwards. This problem occurs only when the router has over 100 BGP peers and a switchover is performed.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef27158	<p>If you try to configure more than the supported 128K VTMS queues, the system does not recover and can become unstable. This problem occurs when the router is configured over its supported limit.</p> <p><b>Workaround:</b> Configure up to 128K queues.</p>
CSCef27202	<p>On Cisco 10000 series routers running in PTA mode, a CPU hog message appears if you execute the <b>show vpdn session</b> command when there are more than 30,000 sessions active. This problem occurs if the number of active sessions is large.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef27221	<p>When a router runs as a LAC and the rate at which PPPoA sessions are established is high, some sessions may not be established and the router can display an error message on the console. This problem occurs when 30,000 PPPoA sessions or more are established at high rate, such as when the ATM link to the DSLAM is restored after a link failure.</p> <p><b>Workaround:</b> Reduce the call admission rate for the PPPoA sessions.</p>
CSCef27417	<p>Output drops can be erroneously reported on the ATM OC12 interface upon reloading the router and without any traffic sent or received on the interface. The output drops interface counter may also report invalid non-zero values with a light traffic load on the interface (PPPoX session establishment). This problem occurs when a high number of VCs is configured on the interface.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef27539	<p>PPPoEoA sessions experience priority traffic drops when using an absolute priority configuration. This problem occurs during traffic congestion; with 8000 PPPoEoA sessions, priority traffic is dropped at the line card.</p> <p><b>Workaround:</b> Modifying the VC queue depth improves but does not alleviate the drops. Changing the configuration to a generic PQ configuration (without absolute priority) alleviates the drops.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)X11 (continued)

Caveat	Description
CSCef28767	<p>If Multilink PPP is configured, Quality of Service (QoS) may not function properly when using a strict priority queue with other bandwidth queues. This problem occurs when the traffic sent to the priority queue exceeds the configured policer bandwidth. All traffic is forwarded through the priority queue regardless of the policer configuration, which has a negative effect on the effective bandwidth of the other queues.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef28798	<p>PPPoX sessions can fail to connect on an ATM interface with the following error message logged on the console port:</p> <pre>XCM access error at ../toaster/c10k_rp/c10kds2_qos.c (4874)</pre> <p>This problem occurs when several thousand QoS service policies are applied on the ATM PVCs. This problem can potentially cause the active PRE2 to crash if the ATM PVCs are configured as create on-demand and the idle-timeout is enabled.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef29360	<p>When a router has over 8000 PPPoEoQinQ active sessions and input and output policing applied to all subinterfaces, the router can report a spurious memory access after you execute the <b>microcode reload pxf</b> command. This problem occurs only when the <b>microcode reload pxf</b> command is executed.</p> <p><b>Workaround:</b> Do not execute the <b>microcode reload pxf</b> command on production routers.</p>
CSCef29940	<p>On Cisco 10000 series routers running as PTA and terminating 31,500 PPPoA sessions, the router can run out of I/O memory when communications to the RADIUS server is lost and PPPoA sessions continue to be established. This problem occurs when the router cannot communicate with the RADIUS server.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef30497	<p>With Q-in-Q configured on the subinterfaces, policy maps applied to the main interface are not inherited by the subinterfaces. This problem only affects subinterfaces configured with Q-in-Q. Policy maps applied to the main interface are not inherited by the subinterfaces.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef30736	<p>When using WRED with 10,000 queues on 4,000 ATM subinterfaces after counters have been cleared, the total output drops on the ATM interface increases without any traffic.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>
CSCef30873	<p>The router can crash due to an "Unexpected Exception" when you flap several Multilink PPP interfaces several times. This problem occurs when over 50 MLPPP interfaces are concurrently brought up, then down, several times in a short period of time.</p> <p><b>Workaround:</b> There is no workaround for this problem.</p>

Table 2 Open Caveats in Cisco IOS Release 12.3(7)XI1 (continued)

Caveat	Description
CSCin67424	When ATM interfaces are connected through an ATM switch and ATM OAM Manage is configured on the router, the ILMI VC cannot be established, therefore traffic is not forwarded. When ATM OAM Manage is configured on an ATM VC passing through an ATM switch, the ILMI VC fails to establish and no traffic is forwarded. <b>Workaround:</b> There is no workaround for this problem.
CSCin74068	When <b>aaa authen login def enable</b> and <b>aaa author exec def gr radius</b> are configured for a new telnet connection, authentication succeeds (with getting a username) on entering the correct enable password, but an access-request is sent to the RADIUS with NULL username for authorization. Authorization should be suppressed when the username is not known and a RADIUS access-request should not be sent with a null username. <b>Workaround:</b> There is no workaround for this problem.
CSCin74698	Two accounting stop records are seen when "rsh" session is established to the router. This happens only when <b>aaa accounting send stop-recod authentication failure</b> command is configured. <b>Workaround:</b> Disable <b>aaa accounting send stop-record authentication failure</b> command if it's not needed.
CSCin76251	The LAC can crash when 8,000 or more PPPoA/LT2P sessions are retrying to connect all at the same time after the LNS and user have had all their sessions cleared. The LAC crashes only under low memory conditions. <b>Workaround:</b> The workaround is to not use the LAC at low memory conditions

## Resolved Caveats—Cisco IOS Release 12.3(7)XI1

This section describes caveats that were fixed in Cisco IOS Release 12.3(7)XI1.

For information about caveats fixed in other Cisco IOS releases, refer to the appropriate Release Note document at the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/aggr/10000/10krn/index.htm>

### CSCea22552

GRE implementation of Cisco IOS is compliant with RFC2784 and RFC2890 and backward compatible with RFC1701.

As an RFC compliancy this DDTS adds the check for bits 4-5 (0 being the most significant) of GRE header.

This issue does not cause any problem for router operation.

### CSCed40933

Cisco Internetwork Operating System (IOS) Software is vulnerable to a Denial of Service (DoS) attack from crafted IPv6 packets when the device has been configured to process IPv6 traffic. This vulnerability requires multiple crafted packets to be sent to the device which may result in a reload upon successful exploitation.

More details can be found in the security advisory, which is posted at <http://www.cisco.com/warp/public/707/cisco-sa-20050126-ipv6.shtml>.

#### CSCee08584

Cisco Internetwork Operating System (IOS) Software release trains 12.1YD, 12.2T, 12.3 and 12.3T, when configured for Cisco's IOS Telephony Service (ITS), Cisco CallManager Express (CME) or Survivable Remote Site Telephony (SRST) may contain a vulnerability in processing certain malformed control protocol messages.

A successful exploitation of this vulnerability may cause a reload of the device and could be exploited repeatedly to produce a Denial of Service (DoS). This advisory is available at <http://www.cisco.com/warp/public/707/cisco-sa-20050119-itscme.shtml>

Cisco has made free software upgrades available to address this vulnerability for all affected customers. This vulnerability is documented by Cisco bug ID CSCee08584.

## Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

### Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

International Cisco websites can be accessed from this URL:

[http://www.cisco.com/public/countries\\_languages.shtml](http://www.cisco.com/public/countries_languages.shtml)

## Ordering Documentation

You can find instructions for ordering documentation at this URL:

[http://www.cisco.com/univercd/cc/td/doc/es\\_inpk/pdi.htm](http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm)

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:  
<http://www.cisco.com/en/US/partner/ordering/index.shtml>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

# Documentation Feedback

You can submit e-mail comments about technical documentation to [bug-doc@cisco.com](mailto:bug-doc@cisco.com).

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems  
Attn: Customer Document Ordering  
170 West Tasman Drive  
San Jose, CA 95134-9883

We appreciate your comments.

## Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, the Cisco Technical Assistance Center (TAC) provides 24-hour-a-day, award-winning technical support services, online and over the phone. Cisco.com features the Cisco TAC website as an online starting point for technical assistance. If you do not hold a valid Cisco service contract, please contact your reseller.

### Cisco TAC Website

The Cisco TAC website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The Cisco TAC website is available 24 hours a day, 365 days a year. The Cisco TAC website is located at this URL:

<http://www.cisco.com/tac>

Accessing all the tools on the Cisco TAC website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a login ID or password, register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

### Opening a TAC Case

Using the online TAC Case Open Tool is the fastest way to open P3 and P4 cases. (P3 and P4 cases are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Case Open Tool automatically recommends resources for an immediate solution. If your issue is not resolved using the recommended resources, your case will be assigned to a Cisco TAC engineer. The online TAC Case Open Tool is located at this URL:

<http://www.cisco.com/tac/caseopen>

For P1 or P2 cases (P1 and P2 cases are those in which your production network is down or severely degraded) or if you do not have Internet access, contact Cisco TAC by telephone. Cisco TAC engineers are assigned immediately to P1 and P2 cases to help keep your business operations running smoothly.

To open a case by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete listing of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

## TAC Case Priority Definitions

To ensure that all cases are reported in a standard format, Cisco has established case priority definitions.

Priority 1 (P1)—Your network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Priority 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Priority 3 (P3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Priority 4 (P4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

## Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Go to this URL to visit the company store:  
<http://www.cisco.com/go/marketplace/>
- The Cisco *Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the Cisco Product Catalog at this URL:  
<http://cisco.com/univercd/cc/td/doc/pcat/>
- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press online at this URL:  
<http://www.ciscopress.com>
- *Packet* magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access Packet magazine at this URL:  
<http://www.cisco.com/packet>
- *iQ Magazine* is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:  
<http://www.cisco.com/go/iqmagazine>

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:  
<http://www.cisco.com/ipj>
- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:  
<http://www.cisco.com/en/US/learning/index.html>

CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0709R)

© 2005 Cisco Systems, Inc. All rights reserved.

♻️ Printed in the USA on recycled paper containing 10% postconsumer waste.