



Release Notes for the Cisco 10000 Series ESR for Cisco IOS Release 12.0(23)SX3

May 19, 2003

These release notes provide information about Cisco IOS software Release 12.0(23)SX3 for the Cisco 10000 Series Router. These release notes describe new features, memory requirements, hardware support, software platform deferrals, and changes to the microcode and related documents.

Cisco IOS Release 12.0(23)SX3 is based on previous 12.0(23)SX releases. Cisco IOS 12.0(23)SX releases include all of the new features introduced in Cisco IOS Release 12.0(23)S. You can review the release notes for Cisco IOS Release 12.0(23)S at the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/relnote/7000fam/rn120s.htm>

Contents

These release notes contain the following sections:

- [Upgrading to a New Software Release, page 2](#)
- [System Requirements, page 2](#)
- [New Features in Cisco IOS Release 12.0\(23\)SX3, page 3](#)
- [New Features in Cisco IOS Release 12.0\(23\)SX2 and 12.0\(23\)SX1, page 3](#)
- [New Features in Cisco IOS Release 12.0\(23\)SX, page 3](#)
- [Limitations and Restrictions, page 11](#)
- [Important Notes, page 12](#)
- [Caveats Resolved in Cisco IOS Release 12.0\(23\)SX3, page 13](#)
- [Caveats Resolved in Earlier Releases, page 14](#)
- [Open Caveats, page 17](#)
- [Obtaining Documentation, page 19](#)
- [Obtaining Technical Assistance, page 20](#)



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

Copyright © 2003. Cisco Systems, Inc. All rights reserved.

Upgrading to a New Software Release

For specific information about upgrading your Cisco 10000 series ESR to a new software release, see the *Cisco 10000 Series ESR Software Configuration Guide* at the following URL:

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

For general information about how to upgrade to a new software release, see the product bulletin *Cisco IOS Upgrade Ordering Instructions* located at:

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/957_pp.htm

For information about how to order Cisco IOS software, refer to the Cisco IOS Software Releases URL:

<http://www.cisco.com/warp/public/cc/pd/iosw/iore/index.shtml>

Upgrading Software from Earlier Cisco IOS Releases

Upgrading from Cisco IOS Release 12.0(21)SX or Release 12.0(21)SX1

If you are upgrading your software from Cisco IOS Release 12.0(21)SX or Release 12.0(21)SX1 to Cisco IOS Release 12.0(23)SX3, you must upgrade the eboot image on the Cisco 10000 series edge services router (ESR). If you fail to upgrade the eboot image, the ESR configuration may not load properly, and a checksum error appears on the console. If you are upgrading the software from an earlier Cisco IOS release, you do not need to upgrade the eboot image.

Upgrading from Cisco IOS Release 12.0(14)SL or from Earlier 12.0(x)SL-based Releases

If you are upgrading your software from Cisco IOS Release 12.0(14)SL or from earlier 12.0(x)SL-based releases to Cisco IOS Release 12.0(23)SX3, save your current configuration file. If you decide to reinstall Release 12.0(14)SL or an earlier release, you must also reinstall the configuration file associated with that release. This is because some Border Gateway Protocol (BGP) configuration-file entries in Release 12.0(23)SX3 are not compatible with Release 12.0(14)SL or earlier releases.

Upgrading Software on Redundant PREs

When you upgrade software on redundant Cisco 10000 series Performance Routing Engines (PREs), be sure to download the software to both the active PRE and the standby PRE before you reload both PREs. For more information, refer to the “Upgrading Software on Redundant PREs” section in the “System Startup and Basic Configuration Tasks” of the *Cisco 10000 Series ESR Software Configuration Guide* at the following URL:

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



Note

In the procedure at the above URL, specify c10k-p10-mz instead of c10k-p6-mz for image name.

System Requirements

This release requires that you have the PRE1 version (part number ESR-PRE1) of the Performance Routing Engine (PRE) installed in the Cisco 10000 series ESR chassis. To verify which PRE is installed in the ESR, use the **show version** command.

Memory Requirements

The following table lists memory requirements for the ESR.

Feature Set by Router	Image Name	Flash Memory	DRAM Memory	Runs From
Edge Services Router	c10k-p10-mz	40 MB	512 MB	RAM
Service Provider/ Secured Shell 3DES	c10k-k4p10-mz	40 MB	512 MB	RAM

New Features in Cisco IOS Release 12.0(23)SX3

Cisco IOS Release 12.0(23)SX3 contains no new features. Several known problems are fixed in this release (see the [“Caveats Resolved in Cisco IOS Release 12.0\(23\)SX3”](#) section on page 13).

Cisco IOS Release 12.0(23)SX3 includes all of the new features introduced in Release 12.0(23)SX, and all of the problems fixed in Release 12.0(23)SX2 and Release 12.0(23)SX1.

New Features in Cisco IOS Release 12.0(23)SX2 and 12.0(23)SX1

Cisco IOS Release 12.0(23)SX2 and Cisco IOS Release 12.0(23)SX1 contained no new features. Both releases fixed known problems (see the [“Caveats Resolved in Earlier Releases”](#) section on page 14).

Cisco IOS Release 12.0(23)SX2 and Release 12.0(23)SX1 include all of the new features introduced in Release 12.0(23)SX. For information about those features, see the [“New Features in Cisco IOS Release 12.0\(23\)SX”](#) section on page 3 below.

New Features in Cisco IOS Release 12.0(23)SX

The following sections describe the new features and improvements introduced in Cisco IOS Release 12.0(23)SX:

- [Multirouter–Automatic Protection Switching, page 4](#)
- [Multicast for MPLS VPN, page 4](#)
- [MPLS Traffic Engineering Fast Reroute, page 5](#)
- [DiffServe Aware Traffic Engineering for MPLS, page 5](#)
- [eIBGP Multipath Load Sharing, page 5](#)
- [Link Fragmentation and Interleaving, page 6](#)
- [Single Rate 3-Color Marker for Traffic Policing, page 10](#)
- [Cisco 10000 ESR MIB Enhancements, page 10](#)
- [SNMP Trap Filtering, page 11](#)

Multirouter–Automatic Protection Switching

The multirouter–automatic protection switching (MR–APS) feature allows switchover of SONET connections in the event of circuit failure. MR–APS is often required when connecting SONET equipment to telco equipment. APS refers to the mechanism of using a "protect" SONET interface in the SONET network as the backup for a working SONET interface. When the working interface fails, the protect interface quickly assumes its traffic load.

The protection mechanism provided by MR–APS 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 non-revertive.

MR–APS provides a protection mechanism in which the “working” and “protect” SONET interfaces are on separate routers. MR–APS also allows both interfaces to be on the same router. However, in this case, we recommend you use single router–APS instead.

Single router–APS (SR–APS) provides a protection mechanism in which the working and protect interfaces are both on the same router. This protection mechanism provides a linear 1+1 unidirectional, non-revertive architecture. Note that the ESR supports bridging in the SR–APS implementation.

For more information on this feature, see the following URL. Note that the documentation includes **aps** commands for POS interfaces. However, the same commands also apply for ATM interfaces.

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios112/ios112p/gsr/posapsgs.htm>

Recommendation for MR–APS Configurations with Limited Static Routes

If your MR–APS configuration has a limited number of static routes (for example, less than 10), we recommend that you use the following command for gigabit Ethernet, ATM, and POS interfaces that have static routes configured for them.

```
Router(config-if)# ip route static update immediate
```

The command adds static routes to the routing table immediately when an interface becomes active, rather than after a route dampening period ends.

Restrictions

- Currently, the ESR supports MR–APS only on OC-12ATM interfaces.
- The **aps reflector** command is not supported in this release.

Multicast for MPLS VPN

The Multicast for Multiprotocol Label-Switching (MPLS) Virtual Private Network (VPN) feature allows service providers to offer multicast services over their MPLS core network. Multicast for MPLS VPN allows end-user customers to improve productivity and communication flow for applications such as corporate communication, e-learning, data warehousing, content synchronization, trading stocks and commodities (stock quotes and ticker information), and emergency messaging services.

Live video and video on demand applications also require a multicast solution that offers scalability, efficiency, and performance. As demand grows for multicast and VPN, enterprises will require service providers to enable multicast across VPNs. Multicast for MPLS VPN provides that solution.

MPLS Traffic Engineering Fast Reroute

The Fast Reroute (FRR) link protection feature of MPLS traffic engineering provides link protection to label-switched paths (LSPs). MPLS traffic engineering automatically establishes and maintains LSPs across the backbone using Resource ReSerVation Protocol (RSVP). Paths for LSPs are calculated at the headend, based on the LSP resource requirements and available network resources such as bandwidth. Under failure conditions, the headend determines a new route for the LSP. This provides for the optimal use of resources. However, due to messaging delays, recovery at the headend is not as quick as recovery at the point of failure.

FRR link protection enables the router to reroute traffic around a failed link without involving the headend in the rerouting decision. This provides quicker recovery time and prevents any further packet loss caused by the failed link. The headend of the tunnel is also notified of the link failure through the IGP or RSVP; the headend then attempts to establish a new LSP that bypasses the failure.

For more information about this feature, see the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120st/120st10/fastrout.htm>



Note

Local reroute prevents any further packet loss caused by a failed link, and gives the headend of the tunnel time to re-establish the tunnel along a new, optimal route.

Restrictions

The Fast Reroute link protection feature works only on:

- Packet over SONET (POS) links (SDH in the European standard)
- Links that use MPLS global label allocation (GLA)

DiffServe Aware Traffic Engineering for MPLS

The DiffServer Aware Traffic Engineering feature extends DiffServe quality of service (QoS) over an MPLS backbone that uses traffic engineering. Bandwidth pools assigned to tunnel interfaces ensure that critical data is associated with a tunnel that has enough bandwidth to transport data over the MPLS network. For information about this feature, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120st/120st14/ds_te.htm

eiBGP Multipath Load Sharing

This section describes the Border Gateway Protocol (BGP) Multipath Load Sharing for external BGP (eBGP) and internal BGP (iBGP) in a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) feature. This feature allows you to configure multihomed autonomous systems and provider edge (PE) routers to:

- Distribute traffic across eBGP and iBGP paths in BGP networks
- Import eBGP and iBGP paths from multihomed and stub networks

BGP installs up to the maximum number of paths allowed (configured by the **maximum-paths** command). BGP uses the best path algorithm to select one multipath as the best path, insert the best path into the routing information base (RIB), and advertise the best path to BGP peers. Other multipaths may be inserted into the RIB, but only one path will be selected as the best path.

The multipaths are used by Cisco Express Forwarding (CEF) to perform load balancing on a per-packet or per-source or destination pair basis. The eIBGP Multipath Load Sharing feature performs unequal cost load balancing by default by selecting BGP paths that do not have an equal cost of the Interior Gateway Protocol (IGP). To enable this feature, configure the router with MPLS VPNs that contain VPN routing and forwarding instances (VRFs) that import both eBGP and iBGP paths. The number of multipaths can be configured separately for each VRF.

**Note**

This feature operates within the configuration parameters of the existing outbound routing policy.

For more information about this feature, see the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t4/fteibmpl.htm>

Restrictions

Route Reflector Limitation—With multiple iBGP paths installed in a routing table, 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 will not be advertised unless separate VRFs with different route distinguishers are configured for each VRF.

Memory Consumption Restriction—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.

Link Fragmentation and Interleaving

Interactive traffic (Telnet, voice over IP) is susceptible to increased latency and jitter when the network processes large packets (for example, LAN-to-LAN FTP transfers traversing a WAN link), especially as they are queued on slower links. The Link Fragmentation and Interleaving (LFI) feature reduces delay and jitter on slower-speed links by breaking up large datagrams and interleaving low-delay traffic packets (such as voice) with the smaller packets resulting from the fragmented datagram. LFI was designed especially for lower-speed links where serialization delay is significant.

The Cisco 10000 series ESR implementation of LFI provides the following benefits:

- Supports the following encapsulation protocols:
 - Frame Relay (FRF.12) end-to-end
 - Multilink PPP (MLPPP)
- Supports LFI on up to 1000 ESR interfaces.
- Prioritized traffic is always sent intact, while all other traffic is subject to fragmentation.
- During packet reassembly, the ESR detects and discards any packets that are missing fragments.

Cisco 10000 Series ESR LFI Limitations and Restrictions

The following limitations and restrictions apply to the Cisco 10000 series ESR implementation of LFI:

- Multilink LFI is restricted to 1-link bundles only.
- Multilink PPP must be enabled on an interface that has interleaving turned on. To enable LFI on an MLPPP interface, use the **ppp multilink interleave** command.
- Multilink PPP over ATM (MLPoATM) is not supported.

- The **show frame-relay fragment** command does not support the following statistics:
 - fragments received (in_frag)
 - fragments dropped (drop_frag)
 - fragments transmitted (out_frag)
- The **show frame-relay fragment interface** command does not support the following statistics:
 - in fragmented pkts
 - in fragmented bytes
 - in dropped reassembling packets
 - in timeouts
 - out interleaved packets
 - out fragmented pkts
 - out fragmented bytes
 - out dropped fragmenting pkts

Cisco 10000 Series ESR LFI Implementation Notes

The following sections provide information about the LFI implementation on the Cisco 10000 series ESR:

- [MLPPP LFI Packet Counts, page 7](#)
- [Configuring LFI on a Frame Relay Interface—Method 1, page 7](#)
- [Configuring LFI on a Frame Relay Interface—Method 2, page 8](#)
- [Configuring LFI on a Multilink PPP Interface, page 9](#)

MLPPP LFI Packet Counts

On an MLPPP interface that has LFI enabled, each packet fragment is counted as a separate packet in the interface's input and output statistics. For example, if a single packet is fragmented into three packets, the **show interface stats** command and the **show interfaces** commands show a count of 3 for Packets In and Packets Out, rather than a count of 1.

Configuring LFI on a Frame Relay Interface—Method 1

One way to configure LFI on a Frame Relay data-link connection identifier (DLCI) is to:

- Use the command **frame-relay interface-queue fair** to enable Frame Relay weighted fair queuing on the physical interface.
- Create a Frame Relay map-class that defines the traffic to prioritize for interleaving. You can also include the following command to define a priority queue for an RTP packet flow:

```
frame-relay ip rtp priority start_port_number port_range [max_bandwidth]
```

- Set the LFI fragmentation size (in bytes) by including the following command in the map-class. Specify a fragmentation size of at least 128 bytes, which is the minimum size supported by the ESR.

```
frame-relay fragment fragment_size
```

- Attach the map-class to the DLCI.
- Use the Frame Relay QoS CLI to configure user-defined queues on the DLCI.

The following is a sample configuration showing this method of configuring LFI. (In this example, fragment size is 384 bytes and the bandwidth for the priority queue is 307 kbps.)

```
interface Serial5/0/0.1/1:0
encapsulation frame-relay
frame-relay class PQ_FR_CLASS_0
frame-relay interface-dlci 17
frame-relay interface-queue fair

map-class frame-relay PQ_FR_CLASS_0
frame-relay fragment 384
frame-relay ip rtp priority 16384 10 307
```

Configuring LFI on a Frame Relay Interface—Method 2

You can also use a nested policy and access lists to configure LFI on a Frame Relay DLCI as follows:

- Use the **priority** command to create a priority queue and use a policy map to associate a class of traffic with the priority queue.

```
policy-map qos_pq_cbwfbq_0
class acl_101
priority percent 10
```

- Create a nested policy and include it in the Frame Relay map class.
- Use the following command to set the fragmentation size (in bytes). Specify a fragmentation size of at least 128 bytes, which is the minimum fragmentation size supported by the ESR.

```
frame-relay fragment fragment_size
```

The following is a sample configuration showing this method of configuring LFI. (In the example, fragmentation size is 768 bytes and bandwidth for the priority queue is 10 percent of the link bandwidth.)

```
interface Serial5/0/0.1/1/1:0
ip address 10.1.1.102 255.255.255.0
no ip directed-broadcast
encapsulation frame-relay
frame-relay class PQ_FR_CLASS_0
frame-relay interface-dlci 17

access-list 101 permit udp any eq 16384 any eq 16384
access-list 102 permit udp any eq 3000 any eq 3000

class-map match-all acl_101
match access-group 101
class-map match-all acl_102
match access-group 102

policy-map qos_pq_cbwfbq_0
class acl_101
priority percent 10
class acl_102
bandwidth percent 30
policy-map outer_policy
class class-default
shape 768
service-policy qos_pq_cbwfbq_0

map-class frame-relay PQ_FR_CLASS_0
service-policy output outer_policy
frame-relay fragment 768
```


Configuring LFI on a Multilink PPP Interface

To configure LFI on an MLPPP interface, do the following:

- Use the following commands to configure multilink operation on the interface:

Command	Purpose
<code>ppp multilink</code>	Enables multilink operation on the interface.
<code>ppp multilink fragmentation</code>	Enables multilink fragmentation on the interface.
<code>ppp multilink fragment-delay milliseconds</code>	(Optional) Specifies the maximum delay (in milliseconds) for fragmentation (for example, you can specify that voice traffic has a maximum fragmentation delay of 20 milliseconds). MLP uses this value to choose a fragment size.
<code>ppp multilink interleave</code>	Enables real-time packet interleaving on the bundle.

- Identify the traffic to prioritize for interleaving.
- Use the **priority** command to create a priority queue and use a policy map to associate that class of traffic with the priority queue. The **priority** command gives priority to a class within the policy map.
- Attach the policy map to an interface to assign that traffic classification and action to the interface.

The following commands show an example of how to configure LFI on an MLPPP interface:

```
class-map match-all VOIP
  match ip rtp 16384 16383
class-map LESS_CRITICAL
  match access-group 101
policy-map VOIP_PRI
  class VOIP
    priority 50
  class LESS_CRITICAL
    set ip precedence 5

interface Multilink1
  ppp chap hostname multilink_name-1
  ppp multilink
  ppp multilink fragment-delay 8
  ppp multilink interleave
  service-policy output VOIP_PRI
  multilink-group 1

interface Serial 2/1
  encapsulation ppp
  ppp multilink
  multilink-group 1
```



Note

The **ip rtp priority** command is not included as a policy action. Instead, the priority configuration defined on the interface determines which traffic is placed in the priority queue.

Single Rate 3-Color Marker for Traffic Policing

The single rate 3-Color Marker feature meters an IP packet stream and marks its packets different colors, based on the Committed Information Rate (CIR) and two associated burst sizes: Committed Burst Size (CBS) and Excess Burst Size (EBS). CIR is measured in bytes of IP packets per second (and it includes the IP header, but not link specific headers). CBS and EBS are measured in bytes.

Previously, the ESR supported a single rate, 2-color marker. The single rate, 3-color marker allows the ESR to classify packets that violate the EBS. This feature is useful, for example, for ingress policing of a service, where service eligibility is determined only by the burst's length, and not its peak rate.

The single rate 3-color marker uses the following colors to classify packets:

- Green (conforming)—Packet size is at most Bc or CBS and within the CIR allowance.
- Yellow (exceeding)—Packet size is greater than the CIR allowance but is at most Be or EBS bytes and within the available surplus.
- Red (violating)—Packet size is greater than both the CIR allowance and the available surplus. This is because either the packet's size exceeds Be or EBS or because a previous packet used some of the surplus and the traffic since then has not slowed sufficiently to acquire the surplus needed for the current packet.

The marker starts with a surplus equal to Be or EBS, and replenishes the surplus by the amount of unused CIR allowance until the surplus reaches Be or EBS.

Configuring the Single Rate 3-Color Marker Feature

To configure this feature, do the following:

- Use the **police** command to set the feature's mode and to assign values for the CIR, CBS, and EBS:

```
Router(config-pmap-c)# police CIR CIR_bps burst-normal burst-max conform-action action
exceed-action action violate-action action
```

For conform, exceed, and violate action, you can specify one of the following actions: **transmit**, **set-dscp-transmit**, **set-prec-transmit**, **set-mpls-exp-transmit**, or **set-qos-transmit**.

Using the Feature

The 3-color marker can be used to mark a packet stream in a service, where different, decreasing levels of assurances (either absolute or relative) are given to packets which are green, yellow, or red. For example, a service might:

- Discard or deny all red packets because they exceed both the committed and excess burst sizes.
- Forward green packets because they are guaranteed for delivery.

Cisco 10000 ESR MIB Enhancements

MIB capabilities on the Cisco 10000 series ESR have been enhanced as follows:

- The CISCO-ENTITY-EXT-MIB and CISCO-OAM-MIB were added.
- Support was verified for the following MIBs: ATM-MIB, CISCO-AAL5-MIB, CISCO-ATM-EXT-MIB, SONET-MIB, RFC1315-MIB, CISCO-FRAME-RELAY-MIB, and CISCO-RF-MIB.
- Support was added for MPLS-LSR-MIB mplsInSegmentOctets and mplsInSegmentOctets.
- The IF-MIB was enhanced to support MPLS.
- MIBs were enhanced to support the ESR high-availability feature.

For more information about the MIB capabilities for this release, see the *Cisco 10000 Series ESR Leased Line MIB Specifications Guide* (version 3) at the following URL. (Table 3-1 in the “MIB Specifications” section of the guide lists the MIBs supported in this release.)

<http://www.cisco.com/univercd/cc/td/doc/product/aggr/10000/10kmibs/llguides/llgdv3/index.htm>

SNMP Trap Filtering

Use the SNMP trap filtering feature to filter linkDown traps so that SNMP only sends a linkDown trap if the main interface goes Down. If a Cisco 10000 series ESR interfaces goes Down, all of its subinterfaces go Down, which results in numerous linkDown traps for each subinterface. This feature filters out those subinterface traps. This feature is turned off by default.

To enable the SNMP trap filtering feature, issue the following CLI command. Use the **no** form of the command to disable the feature.

```
[no] snmp ifmib trap throttle
```

Limitations and Restrictions

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.

Controlling the Rate of Logging Messages

It is important that you limit the rate that system messages are logged by the Cisco 10000 series ESR. 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.

We recommend that you configure the **logging rate-limit** command as follows. This rate-limits 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 on the **logging rate-limit command**, see the *Cisco IOS Configuration Fundamentals Command Reference*.

Frame Relay

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

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

Nested Policy Feature

The following limitations and restrictions apply to the Cisco 10000 series ESR nested policy feature:

- Nested policies support no more than two hierarchy levels.
- For additional restrictions and limitations on creating nested policies, refer to the “Configuring Nested Policies on the ESR” section of the Cisco document at the following location:

http://www.cisco.com/univercd/cc/td/doc/product/aggr/10000/10kfm/nest_pol.htm

- The nested-policy shape rate is reserved for nested-policy traffic only. Excess bandwidth is not used for other traffic.
- DotP marking is not supported, nor is 802.1P (including matching and marking the 802.1P header).



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 command **show policy-map interface** to determine the actual shape rate.

Testing Performance of High-Speed Interfaces

Cisco IOS software running on the Cisco 10000 series ESR 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 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. To ensure accurate test results, you should test the throughput of the gigabit Ethernet, 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.

Important Notes

Cisco Discovery Protocol

Unlike other Cisco routers, on the Cisco 10000 series ESR the Cisco Discovery Protocol (CDP) is disabled by default. You can enable CDP on an interface using the **cdp enable** command.

Frame Relay and PPP Sessions

You can run up to 4000 Frame Relay sessions or 4000 PPP sessions, and you can configure up to 800 Border Gateway Protocol (BGP) peers on the Cisco 10000 series ESR. The ESR also supports up to 512 Multilink Point-to-Point (MLP) protocol sessions and up to 1024 MLP bundles.

Performance Routing Engine 1 on the Cisco 10000 Series Edge Services Router

In order for Cisco IOS Release 12.0(22)S and later releases to run on the Cisco 10000 series ESR, the Performance Routing Engine (PRE) installed in the chassis must be the PRE1 version (part number ESR-PRE1). You can verify which PRE is installed in the chassis by using the **show version** command.



Note The Cisco 10000 series ESR does not support mixing two different PRE revisions in the same chassis. Do not install a PRE and PRE1 in the same chassis.

VLAN Session Support

The Cisco 10000 series ESR provides session support for 4000 802.1Q VLANs.

Inserting a New Line Card

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

RBE Subinterfaces

To ensure that the Performance Routing Engine (PRE) microcode has adequate IP address space, use IP addresses in a contiguous address space. Also, use the **ip unnumbered** command on the RBE subinterface.

Caveats Resolved in Cisco IOS Release 12.0(23)SX3

The following problems have been fixed in Cisco IOS Release 12.0(23)SX3:

CSCea55465

Each channelized subinterface is assigned a unique interface index (ifIndex) number. Previously, the router sometimes assigned the same ifIndex number to multiple channelized subinterfaces, which caused inconsistency between the routing and Cisco Express Forwarding (CEF) prefix tables for routes through one of the interfaces.

CSCea55641

When Link Fragmentation and Interleaving (LFI) is enabled on a multilink PPP interface, the router no longer experiences problems with priority queue (PQ) traffic latency being higher than expected.

CSCea61833

The Parallel eXpress Forwarding (PXF) processor no longer generates a T1 software exception error when multicast traffic with an IP precedence of 5 is sent over the priority queue. Cisco IOS Release 12.0(23)SX2 provided a partial fix to this problem. This release provides a complete fix for the problem.

CSCea67032

Channelized T3 interfaces no longer fail to initialize. Previously, when a 6-port channelized T3 line card was configured with 16 to 20 time slots, several interfaces might fail to initialize.

CSCea71704

A router reload followed by a configuration download no longer causes the router to stop passing multicast VPN traffic. Previously in this situation, the PXF began to drop packets, which resulted in some interfaces not being enabled for multicast.

CSCea73534

PXF no longer fails when NetFlow is enabled in a configuration that contains a multilink PPP bundle. Previously, the PXF failed (causing a PXF reload) whether NetFlow was enabled on the multilink PPP bundle or on another router interface.

CSCea79268

The Fast Reroute (FRR) feature now functions correctly on the router. Previously, when Multirouter Automatic Protection Switching (MR-APS) was enabled on the router, a FRR did not occur when the router was reloaded, although it should have.

CSCeb01868

In a multicast VPN environment, multicast VPN routing/forwarding (MVRF) now works correctly. Previously, when you configured a new MVRF instance, Protocol Independent Multicast (PIM) neighbors were not detected across the MVRF tunnel and traffic was sometimes dropped.

CSCeb05703

MPLS Traffic Engineering Fast Reroute (FRR) now works correctly on POS interfaces. Previously, when a loss of signal (LOS) alarm was received on a POS interface with an FRR tunnel configured, the state of the tunnel did not immediately change from Ready to Active, although it should have.

Caveats Resolved in Earlier Releases

The following sections present the problems that were fixed in the following earlier releases:

- [Cisco IOS Release 12.0\(23\)SX2, page 14](#)
- [Cisco IOS Release 12.0\(23\)SX1, page 15](#)

Cisco IOS Release 12.0(23)SX2

The following problems were fixed in Cisco IOS Release 12.0(23)SX2:

CSCdx87500

The **mstat** and **mtrace** commands now work correctly over a multicast distribution tree (MDT) tunnel.

CSCdy50036

The **show ip pim mdt history** command now displays the correct reuse value for MDT data groups. Previously, in cases where one MDT data group was being reused and another group was being used only one time, the command incorrectly showed that both groups were being reused.

CSCdz28485

When the default MDT group is removed on the remote provider edge (PE), the Cisco 10000 series ESR no longer stops responding or experiences spurious access.

CSCdz28491

The router no longer loses tunnel information when an MDT default group is removed from a remote PE.

CSCdz30254

The router no longer displays a card provisioning conflict message when the configuration is downloaded after a Cisco IOS software reload. Previously, this message appeared even though the cards were never removed from the chassis and the configuration was not changed.

CSCdz32805

The router now forms data MDTs consistently when configured for use in a multicast VPN environment.

CSCdz55717

The Open Shortest Path First (OSPF) router process no longer experiences a memory leak when OSPF sham links are configured in a Multiprotocol Label Switching (MPLS) VPN environment.

CSCea18595

When configured for multicast VPN, the router now correctly triggers an assertion when it receives data on a tunnel interface.

CSCea27138

When configured for multicast VPN, the router now handles MDT data mappings correctly. Previously, when the router was configured as a receiving PE router, the router deleted MDT mappings too quickly, or did not delete them at all.

CSCea27231

When you use the **priority** command to assign link bandwidth to priority-queue traffic, the traffic is dequeued and scheduled at the specified bandwidth. Previously, the traffic was dequeued and scheduled at the full link bandwidth, which was incorrect.

CSCea33501

Setting the MAC rewrite index to zero no longer causes the PXF to reload.

CSCea54257

Link Fragmentation and Interleaving (LFI) now correctly handles 1-part dequeuing with two buffers.

CSCea61833

The PXF no longer generates a T1 software exception error when multicast traffic with an IP precedence of 5 is sent over the priority queue.

CSCea70433

The PXF no longer fails when the router is passing multicast VPN traffic.

CSCea71781

When multicast VPN and Multicast Source Discovery Protocol (MSDP), the router no longer generates a LINK-2-LINEST message when it receives a broadcast IP packet on a multicast VPN tunnel.

Cisco IOS Release 12.0(23)SX1

The following problems were fixed in Cisco IOS Release 12.0(23)SX1:

CSCdy01660

MIB walks on the sonetVTIntervalTable and sonetFarEndVTIntervalTable (SONET-MIB) now work correctly. Previously, information in the tables was deleted after 24 hours, which caused MIB walks to loop.

CSCdz00466

SNMP now provides correct values for the MIB objects atmTrafficDescrType and atmTrafficQoSClass for all types of PVCs. Note that this problem was a duplicate of CSCdp53356, which has been fixed.

CSCdz29005

On Cisco channelized OC-12 and channelized STM line cards, the SNMP MIB object ifOperStatus now correctly shows T1 layers as Down when all other layers are Down.

CSCdz30172

A traceback message no longer appears during stateful switchover (SSO) testing of PPP.

CSCdz30287

The Cisco 10000 series ESR now shapes traffic correctly when LFI is enabled on an MLP bundle running over a T1 channel.

CSCdz36334

Previously, when a priority queue (PQ) was underloaded and fragmentation queues (FQs) were overloaded, the PQ Max latency and Avg latency values were higher than they should have been. The latency values are now correct. Note that this problem was a duplicate of CSCdz70148, which has been fixed.

CSCdz40022

The **show interface** command now shows correct input and output byte counts for MLP interfaces with LFI enabled. The **clear counters** command also clears the output packet and byte counters as it should.

CSCdz49398

A traceback message no longer appears while using the 3-color policer to police 1000 VLANs.

CSCdz50858

A stateful switchover (SSO) no longer causes a Cisco 10000 series ESR with redundant PRE-1 cards and channelized OC-12 or channelized STM line cards to experience a traceback or %CEF fibidb error.

CSCdz51638

The **show frame-relay fragment** command increments the “out-frag” and “out fragmented” counts as it should. Previously, the counts did not increment.

CSCdz56052

Cisco channelized T3 (CT3) line cards no longer stop responding when 1000 Frame Relay or MLP interfaces (with LFI enabled) are configured over two CT3 line cards. Note that CSCdz49702 is a duplicate of this.

CSCdz58839

A PXF failure no longer occurs when a fast reroute (FRR) is performed after a primary link fails. This problem was observed when the Cisco 10000 series ESR was configured as an MPLS traffic engineering (TE) tunnel head and FRR break point, and the primary link had 500 tunnels.

CSCdz70148

The Cisco 10000 series ESR now processes priority traffic correctly. Previously, when priority traffic arrived in a burst or exceeded the configured rate, the traffic experienced a longer-than-expected delay.

CSCdz77838

A PXF microcode reload no longer breaks MPLS VPN multicast traffic forwarding. Previously, a reload could cause Protocol Independent Multicast (PIM) VPN routing/forwarding (VRF) neighbor relationships to be lost. When this happened, the router stopped forwarding VRF multicast traffic.

CSCea13379

Previously, when the Cisco 10000 series ESR received PIM (*,G) joins with (S,G)R prune, the router incorrectly forwarded only the (*,G) joins to the rendezvous point (RP). Now, the router correctly forwards the (*,G) joins with (S,G)R prune.

CSCea15963

When two redundant provider edge (PE) Cisco 10000 series ESRs are connected to a single VPN, both routers might send different multicast distributed tunnel-Join (MDT-Join) updates to the VPN source. This can cause the receivers to toggle between different MDT data groups, which results in intermediate data loss for the receivers.

CSCea18756

Previously, the NetFlow aggregation test failed to match the correct prefix, mask, and AS values for source, destination, prefix, and AS types of aggregation schemes. Now, the values match as expected.

Open Caveats

Table 1 describes the open caveats in Cisco IOS Release 12.0(23)SX3.

Table 1 Open Caveats in Cisco IOS Release 12.0(23)SX3

Caveat	Description
General	
CSCdz02942	<p>After a stateful switchover (SSO) in a redundant PRE system, the MIB objects <code>ifLastChange</code> (IF-MIB) and <code>atmVclLastChange</code> (ATM-MIB) might not match <code>sysUpTime.0</code> (MIB-2), but they should. This problem can occur:</p> <ul style="list-style-type: none"> • After a second switchover • When the standby PRE is reset, followed by a switchover • When the standby PRE is inserted in the chassis later, after insertion of the primary PRE <p>Workaround: None.</p>
CSCdz26218	<p>PXF diversion statistics do not clear when you issue the microcode reload all command, although they should.</p> <p>Workaround: Reload the Cisco IOS software.</p>
CSCdz29077	<p>After you issue the commands shutdown and no shutdown on a Cisco channelized OC-12 line card, the show controller t3 command incorrectly displays the T3 controller as Down. The controller is really Up.</p> <p>Workaround: Issue the shutdown command, wait several seconds, and then issue the no shutdown command on the SONET interface. This sets the T3 interface to Up in the show controller t3 command.</p>
CSCdz32795	<p>A routing policy does not count packets of 65 bytes or larger. The packets are transported, but not counted.</p> <p>Workaround: None.</p>
CSCdz35018	<p>A traceback occurred during testing of IP multicast.</p> <p>Workaround: None.</p>
CSCdz54597	<p>The Cisco 4-port channelized OC-3 line card experiences low throughput and packet loss when configured with 252 SONET interfaces (VT E1 unframed).</p> <p>Workaround: None.</p>
CSCdz56101	<p>On a Cisco channelized 4-port STM-1 line card with 756 interfaces, each at 512 Kbps with SDH framing and AU-4-TUG-3 controllers, the line protocol failed to initialize for several interfaces.</p> <p>Workaround: None.</p>

Table 1 Open Caveats in Cisco IOS Release 12.0(23)SX3 (continued)

Caveat	Description
General (continued)	
CSCea66250	When an ICMP echo request (ping) packet is received on an interface that has access control list (ACL) logging enabled, the PXF can fail and a Local Bus Exception error message appears. Workaround: Disable ACL logging on the interface.
CSCea74742	When the traffic policing max burst size is set to a small value, the policing exceed function does not work correctly. Workaround: None.
SNMP and MIBs (see also General)	
CSCdz50531	The SONET-MIB does not update the intervals in the sonetVTIntervalTable or sonetVTFarEndIntervalTable for channelized STM and channelized OC-12 AU-3 TU controllers. Workaround: None.
CSCdz65705	When SNMP trap filtering is enabled (snmp ifmib trap throttle), SNMP does not generate a notification when a subinterface is shut down, although it should. Workaround: None.
CSCea08828	When SNMP trap filtering is enabled on a Cisco 1-port channelized OC-12 line card, the T3 layer generates a linkDown trap when the SONET layer is shut down, although it should not. Workaround: None.
Stateful Switchover	
CSCdy66774	After a stateful switchover (SSO), the first packet sent through UDP is lost. Workaround: Configure the ESR to check that the first IP packet has an ARP entry defined for its destination (which must be connected through a fast Ethernet interface). To do this, issue the following CLI command (where <i>A.B.C.D</i> is the destination IP address and <i>H.H.H</i> is the 48-bit hardware address): <pre>conf t arp A.B.C.D H.H.H</pre>
eiBGP Multipath Load Sharing	
CSCdy88837	If eiBGP multipath is configured in an MPLS network and one of the links is configured as eBGP, traffic may not be distributed through all links. This problem occurs when MAX Paths EIBGP is enabled under the address family. In this case, the IP BGP table displays multiple routes to a destination network (eBGP and iBGP paths), but the IP Route VPN table only shows iBGP paths. Workaround: Disable eiBGP and use the weight command to forward traffic through eBGP or iBGP paths.
CSCdz10264	When the ESR is configured as a provider edge (PE) router with 100 VPNs, the number of packets reported at an ingress port might not reflect the actual number of packets forwarded by the ESR. This problem occurs on gigabit Ethernet (gigE) line cards when a gigE port is configured as a subinterface on a VLAN and traffic is sent to each VLAN ID from another router. Workaround: None.

Table 1 Open Caveats in Cisco IOS Release 12.0(23)SX3 (continued)

Caveat	Description
eiBGP Multipath Load Sharing (continued)	
CSCdz26257	Traffic might not be load balanced through iBGP paths when MPLS eiBGP or iBGP max-paths is configured. For example, two iBGP paths might exist to the next hop but the PXF only forwards traffic through one of them. This means one of the links is not being used. Workaround: If traffic is being forwarded through the slower links, you can shut down the slower link to allow the faster link to forward traffic.
Link Fragmentation and Interleaving	
CSCdy75500	The show frame-relay fragment command truncates the Frame Relay interface name for channelized 4-port STM-1 line cards. Workaround: None.
CSCdz27628	When class-based weighted fair queuing (CBWFQ) is used with LFI enabled and multiple queues are configured for a link, packets are dropped from underloaded queues rather than overloaded (congested) queues. This problem occurs when traffic with different packet sizes is sent to different queues, which causes fragmentation on some queues but not others. Workaround: Be sure to use a priority queue when you enable LFI.
CSCdz42829	When LFI is enabled, the show frame pvc command sometimes shows invalid counts for input and output packets and bytes. Workaround: None.

Obtaining Documentation

Cisco provides several ways to obtain documentation, 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 web sites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which may have shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Registered Cisco.com users can order a single Documentation CD-ROM (product number DOC-CONDOCCD=) through the Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/ordering_place_order_ordering_tool_launch.html

All users can order monthly or quarterly subscriptions through the online Subscription Store:

<http://www.cisco.com/go/subscription>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

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 Networking Products MarketPlace:

<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, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. On the Cisco Documentation home page, click **Feedback** at the top of the page.

You can e-mail your comments to bug-doc@cisco.com.

You can submit your comments 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

Cisco provides Cisco.com, which includes the Cisco Technical Assistance Center (TAC) Website, as a starting point for all technical assistance. Customers and partners can obtain online documentation, troubleshooting tips, and sample configurations from the Cisco TAC website. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC website, including TAC tools and utilities.

Cisco.com

Cisco.com offers a suite of interactive, networked services that let you access Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com provides a broad range of features and services to help you with these tasks:

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

To obtain customized information and service, you can self-register on Cisco.com at this URL:

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

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two levels of support are available: the Cisco TAC website and the Cisco TAC Escalation Center. The avenue of support that you choose depends on the priority of the problem and the conditions stated in service contracts, when applicable.

We categorize Cisco TAC inquiries according to urgency:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration. There is little or no impact to your business operations.
- Priority level 3 (P3)—Operational performance of the network is impaired, but most business operations remain functional. You and Cisco are willing to commit resources during normal business hours to restore service to satisfactory levels.
- Priority level 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively impacted by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.
- Priority level 1 (P1)—An existing 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.

Cisco TAC Website

The Cisco TAC website provides online documents and tools to help troubleshoot and resolve technical issues with Cisco products and technologies. To access the Cisco TAC website, go to this URL:

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

All customers, partners, and resellers who have a valid Cisco service contract have complete access to the technical support resources on the Cisco TAC website. Some services on the Cisco TAC website require a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to this URL to register:

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

If you are a Cisco.com registered user, and you cannot resolve your technical issues by using the Cisco TAC website, you can open a case online at this URL:

<http://www.cisco.com/en/US/support/index.html>

If you have Internet access, we recommend that you open P3 and P4 cases through the Cisco TAC website so that you can describe the situation in your own words and attach any necessary files.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses priority level 1 or priority level 2 issues. These classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer automatically opens a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to this URL:

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

Before you call, check with your network operations center to determine the level of Cisco support services to which your company is entitled: for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). When you call the center, have your service agreement number and your product serial number available.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site 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. (0304R)

Copyright © 2003, Cisco Systems, Inc.
All rights reserved.