



Reducing Latency and Jitter for Real-Time Traffic Using Multilink PPP

Last Updated: December 7, 2011

This module contains information about reducing latency and jitter for real-time traffic on your network. One Cisco mechanism for reducing latency and jitter for real-time traffic is Multilink PPP (MLP), also known as Multilink. This module contains conceptual information about Multilink and describes how Multilink PPP can be used with network peers to reduce latency and jitter for real-time traffic on your network.

- [Finding Feature Information, page 1](#)
- [Information About Multilink, page 1](#)
- [Where to Go Next, page 5](#)
- [Additional References, page 6](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Multilink

- [Restrictions for Multilink, page 2](#)
- [Multilink Functionality, page 2](#)
- [Multiclass Multilink PPP, page 4](#)
- [Distributed Multilink PPP, page 5](#)



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

Restrictions for Multilink

Multilink uses first-in first-out (FIFO) queuing to queue and interleave packets. Alternative mechanisms such as low latency queuing (LLQ), weighted fair queuing (WFQ), or class-based weighted fair queuing (CBWFQ) may be used. If you want to use one of these alternative mechanisms, enable it before configuring Multilink. For more information about queuing mechanisms, see the "Configuring Weighted Fair Queueing" module.

Multilink Functionality

At the top level, Multilink provides packet interleaving, packet fragmentation, and packet resequencing across multiple logical data links. The packet interleaving, packet fragmentation, and packet resequencing is used to accommodate the fast transmission times required for sending real-time packets (for example, voice packets) across the network links. Multilink is especially useful over slow network links (that is, a network link with a link speed less than or equal to 768 kbps).

- [Multilink Interleaving, page 2](#)
- [Multilink Fragmentation, page 2](#)
- [Multilink Resequencing, page 4](#)
- [Multilink Bundles and Their Network Links, page 4](#)

Multilink Interleaving

Multilink interleaving is based upon two other integral Multilink activities:

- The ability to fragment packets (or datagrams)
- The ability to multiplex at least two independent data streams

The term interleaving comes from the latter activity, that is, the interleaving of two (or more) independent data streams which are processed independently by the network peer.

Multilink interleaving is a mechanism that allows short, real-time (that is, time-sensitive) packets to be transmitted to a network peer within a certain amount of time (the "delay budget"). To accomplish this task, Multilink interleaving interrupts the transmission of large non-time-sensitive (sometimes referred to as "bulk") datagrams or packets in favor of transmitting the time-sensitive packet. Once the real-time packet is sent, the system resumes sending the bulk packet.

An example may help to illustrate the concept of delay budget. The network starts transmitting a large datagram to a network peer. This large datagram takes 500 milliseconds (ms) to transmit. Three milliseconds later (while the large datagram is still being transmitted), a voice packet arrives in the transmit queue. By the time the large datagram is completely transmitted (497 ms later) the voice packet (which is highly time-sensitive) is subject to unacceptable delay (that is, its delay budget is exceeded).

Multilink interleaving is particularly useful for applications where too much latency (that is, delay) is detrimental to the function of the application, such as Voice over IP (VoIP). However, it is also beneficial for other forms of "interactive" data, such as Telnet packets where the Telnet packets echo the keystrokes entered by the user at a keyboard.

Multilink Fragmentation

With Multilink fragmentation, the large datagram is fragmented ("chopped") into a number of small packet fragments, Multilink headers are added to the packet fragments, and the packet fragments are transmitted individually to a network peer.

When interleaving is enabled, the packet fragments are small enough so that the time it takes to transmit them does not exceed the time budgeted for transmitting the real-time (time-sensitive) data packet. The real-time data packets are interleaved between the fragments of the large datagram.

Each time Multilink prepares to send another data packet fragment or frame to the receiving network peer, Multilink first checks to see if a real-time (time-sensitive) packet has arrived in the transmit queue. If so, the high-priority packet is sent first before sending the next fragment from the large datagram.

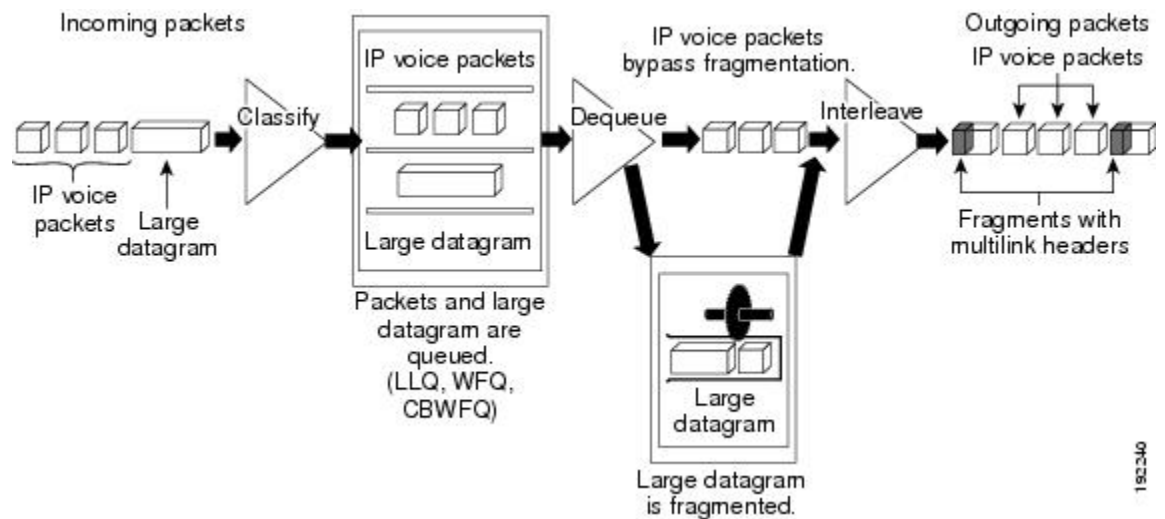
The time delay before the priority packets arrive at the receiving network link is subject to the usual serialization delays at the network link level. That is, any other data already being transmitted has to be finished before the priority packet can be sent. By segmenting long datagrams into small fragments, and checking for newly arrived priority frames between fragments, the priority frame is delayed only by the time it takes to transmit a previously queued fragment rather than a complete large datagram.

Thus, the maximum size of the fragments dictates the responsiveness for insertion of priority packets into the stream. The fragment size can be tuned by adjusting the fragment delay with the **ppp multilink fragment delay** command.

To ensure correct order of transmission and reassembly (which occurs later), multilink headers are added to the large datagram fragments after the packets are dequeued and ready to be sent.

The figure below is a simplified illustration of how Multilink fragments and interleaves packets.

Figure 1 *Multilink Fragmentation and Interleaving*



In the figure above, both IP voice packets and a large datagram arrive at the interface from a single network link. Your network may have multiple links. The IP voice packet and large datagram are queued according to their classification. The large datagram is fragmented (the IP voice packets are not). The IP voice packets are interleaved between the fragments of the large datagram, to which multilink headers are added.

Packets Dequeued and Transmitted

When the large datagram is dequeued, and space becomes available on a member link, Multilink takes a fragment from the original large datagram and transmits the fragments over that link. If an IP voice packet (or other real-time packet) arrives at the transmit queue before Multilink has completely sent the datagram fragment, the next time a link is available to send more packets, Multilink will dequeue and send the high-priority packet. The high-priority packet will be sent instead of another fragment from the large datagram.

Multilink Resequencing

A multilink bundle is a virtual Point-to-Point Protocol (PPP) connection or session over a network link. A multilink bundle at the transmitting end of the network sends the fragments to a multilink bundle on the receiving end of the network link.

The multilink bundle at the receiving end of the network accepts the fragments from the transmitting multilink bundle.

As fragments are received, the multilink bundle reassembles (resequences) the original large datagram from the fragments using the sequence number in the multilink header attached to the fragment by the sender. The reassembled large datagrams are then forwarded in normal fashion.

Multilink Bundles and Their Network Links

As mentioned earlier, a multilink bundle is a virtual PPP connection over a network link. The transmitting multilink bundle transmits the packet over a network link to a receiving multilink bundle, where the multilink bundle reassembles the fragments using the sequence number in the multilink header of the fragment.

The individual member links in a multilink bundle are standard serial PPP connections. Most forms of PPP connections may be used as member links in a bundle, including PPP over ATM, PPP over Frame Relay, and PPP over dial interfaces. However, there may be certain limitations and issues associated with using PPP sessions over certain media types, particularly those for "tunneling" protocols such as PPP over ATM, PPP over Frame Relay, and PPP over Ethernet.

Multiclass Multilink PPP

Multiclass Multilink PPP (MCMP) is based on RFC 2686: *Multi-Class Extension to Multi-Link PPP*. Multiclass Multilink PPP is an extension to the multilink functionality that adds the ability to divide network traffic over the multilink bundle into several independently sequenced streams of fragments. Multilink, as defined by RFC 1990: *The PPP Multilink Protocol (MP)*, provides for one sequenced stream only. RFC 1990 also implicitly allows one additional unsequenced stream, as large datagrams may be transmitted without multilink headers as long as the large datagrams do not need to be fragmented.

In Multiclass Multilink PPP, outgoing packets may be divided into as many as 16 different streams, for which RFC 2686 uses the term classes. Each stream or class has its own governing sequence number, and the receiving network peer (bundle) sorts and processes each stream independently.

Packets can still be sent without multilink headers. However, part of the purpose behind Multiclass Multilink PPP is to reduce or eliminate the need to send unsequenced data.

Multiclass Multilink PPP was created explicitly to allow the packets to be divided into several preemptable classes, so that any lower priority class could be interrupted in favor of sending a packet from a higher priority class. Each class of data can be fragmented, and all classes are expected to be fragmented (with the possible exception of the highest priority class). Also, frames from the different streams may be mixed if necessary.

Multiclass Multilink PPP was created as a mechanism to allow implementations to do interleaving, yet without giving up the sequencing of the interleaved packets such as occurs with standard interleaving.

In the Cisco IOS software, when Multilink Multiclass PPP is used instead of standard interleaving, the regular non-priority data is fragmented and transmitted in one class, and interleaved frames are sent in a separate class. Specifically, the regular traffic is sent in class 0 and the interleaved frames are sent in class 1. Thus, interleaving works just as it does with standard interleaving, except that the interleaved frames are

sent in class 1 rather than as unsequenced frames. Multilink does not transmit data using additional classes, although Multilink is capable of receiving data from peers that do.

Multiclass Multilink PPP must be successfully negotiated with the peer system. If interleaving and Multiclass Multilink PPP are both configured, but the use of Multiclass Multilink PPP cannot be negotiated with the peer system, standard interleaving will be used.

Distributed Multilink PPP

Distributed Multilink PPP (dMLP) is an implementation of Multilink on systems that support distributed processing. With distributed processing, packet processing can be handled by "dedicated hardware"--that is, either by the CPU or by another internal device such as a Versatile Interface Processor (VIP) inside the router or a FlexWAN inside the switch. This dedicated hardware can also be referred to as the "dMLP engine."

One system that supports distributed processing is the Cisco 7500 series router with a Versatile Interface Processor (VIP2-40 or higher). Distributed processing is supported on a number of additional routers and switches as well. Refer to the documentation for your specific router or switch to see if it supports distributed processing.



Note

On a Cisco 7500 series router, a VIP2-50 or higher is recommended when the aggregate line rate of the port adapters on the VIP is greater than DS3. A VIP2-50 card is required for OC-3 rates.

With dMLP, packet fragmentation, interleaving, and fragment reassembly are done by the dMLP engine instead of by the Cisco IOS software. However, the Cisco IOS software manages the member links, creates and disassembles the bundles, and handles the control plane processing (including the handling of all PPP control packets).

However, once a bundle is established, the handling of Multilink packets is turned over to the dMLP engine. The dMLP engine handles all the multilink data-path functionality, including fragmentation, interleaving, multilink encapsulation, load balancing among the multiple links, and sorting and reassembly of inbound fragments.

The capabilities of the dMLP engines vary widely, and they may not always behave like the Cisco IOS Multilink feature. The dMLP engine may fragment and load balance using entirely different schemes than those used by the Cisco IOS software, and they may not support the same multilink features. For more information, refer to the documentation for the dMLP engine you are using.

Where to Go Next

The next step is to go to the module containing the instructions for the type of Multilink PPP you want to use, as listed below.

To use Multilink PPP over Frame Relay, see the "Using Multilink PPP over Frame Relay" module.

To use Multilink PPP over ATM links, see the "Using Multilink PPP over ATM Links" module.

To use Multilink PPP over dialer interface links, see the "Using Multilink PPP over Dialer Interface Links" module.

To use Multilink PPP over serial interface links, see the "Using Multilink PPP over Serial Interface Links" module.

**Note**

If you are using an ASR 1000 Series Router, follow the instructions for using Multilink PPP over serial interface links.

Additional References

The following sections provide additional references about Multilink.

Related Documents

Related Topic	Document Title
QoS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS Quality of Service Solutions Command Reference</i>
LLQ, WFQ, CBWFQ, PQ, CQ, FIFO and other queueing mechanisms	"Configuring Weighted Fair Queueing" module
Frame Relay configurations	Configuring Frame Relay" module
ATM configurations	"Configuring ATM" module
Multiclass Multilink PPP	"Multiclass Multilink PPP" module
Multilink PPP configuration information	"Configuring Media-Independent PPP and Multilink PPP" module
Multilink PPP over ATM links (including ATM interfaces and ATM PVCs)	"Using Multilink PPP over ATM Links" module
Multilink PPP over Frame Relay	"Using Multilink PPP over Frame Relay" module
Multilink PPP over dialer interface links	"Using Multilink PPP over Dialer Interface Links" module
Multilink PPP over serial interface links	"Using Multilink PPP over Serial Interface Links" module

Standards

Standard	Title
No new or modified standards are supported, and support for existing standards has not been modified.	--

MIBs

MIB	MIBs Link
No new or modified MIBs are supported, and support for existing MIBs has not been modified.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 1990	The PPP Multilink Protocol (MP)
RFC 2686	Multi-Class Extension to Multi-Link PPP

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned 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. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2011 Cisco Systems, Inc. All rights reserved.