



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

First Published: May 2, 2005
Last Updated: July 18, 2008

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.

Contents

- [Information About Multilink, page 1](#)
- [Where to Go Next, page 6](#)
- [Additional References, page 6](#)
- [Glossary, page 8](#)

Information About Multilink

Before configuring Multilink, you should understand the following concepts:

- [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

© 2007 Cisco Systems, Inc. All rights reserved.

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 Queuing”](#) 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

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.

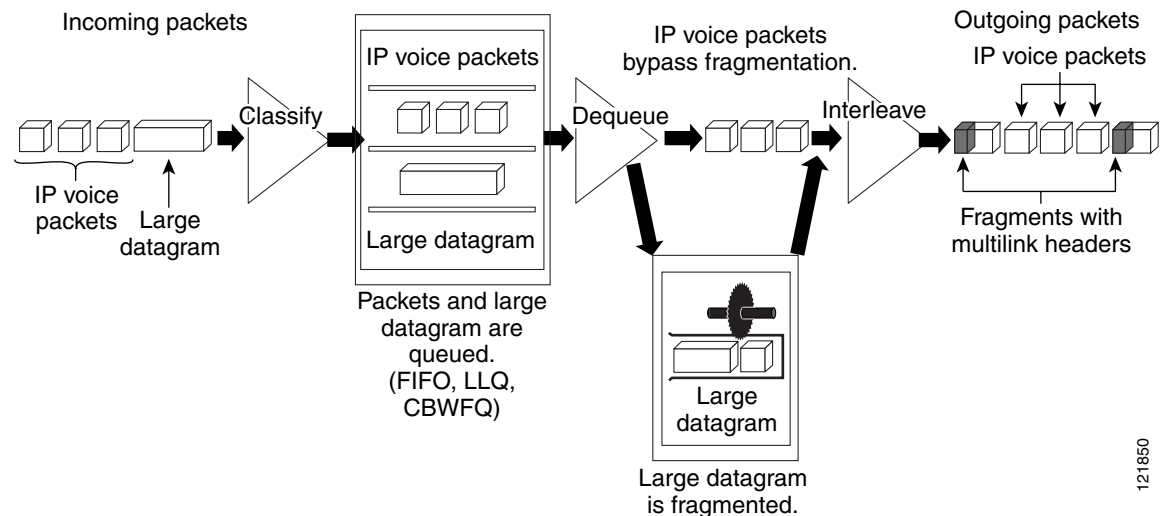
The time delay before the priority packets arrive at the receiving network link is subject to the usual serialization delays at the network link level. That is, any other data already being transmitted has to be finished before the priority packet can be sent. By segmenting long datagrams into small 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.

Figure 1 is a simplified illustration of how Multilink fragments and interleaves packets.

Figure 1 *Multilink Fragmentation and Interleaving*



In Figure 1, 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.

The instructions for using Multilink over PPP on certain links are documented in separate modules (see the [“Where to Go Next”](#) section for more details). Follow the instructions for the type of link you are using.

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.

For more information about Multiclass Multilink PPP, see the [“Multiclass Multilink PPP”](#) module.

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	Cisco IOS Quality of Service Solutions Command Reference
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
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/techsupport

Glossary

ATM—Asynchronous Transfer Mode. The international standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media, such as E3, SONET, and T3.

datagram—Logical grouping of information sent as a network layer unit over a transmission medium without prior establishment of a virtual circuit. IP datagrams are the primary information units in the Internet. The terms cell, frame, message, packet, and segment also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.

jitter—1. The interpacket delay variance; that is, the difference between interpacket arrival and departure. Jitter is an important QoS metric for voice and video applications. 2. Analog communication line distortion caused by the variation of a signal from its reference timing positions. Jitter can cause data loss, particularly at high speeds.

latency—1. Delay between the time a device requests access to a network and the time it is granted permission to transmit. 2. Delay between the time a device receives a frame and the time that frame is forwarded out the destination port.

packet—Logical grouping of information that includes a header containing control information and (usually) user data. Packets most often are used to refer to network layer units of data. The terms datagram, frame, message, and segment also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.

PVC— permanent virtual circuit (or connection). Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection.

PPP—Point-to-Point Protocol. Successor to Serial Line Internet Protocol (SLIP) that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, Internetwork Packet Exchange (IPX), and AppleTalk Remote Access (ARA). PPP also has built-in security mechanisms, such as Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP). PPP relies on two protocols: link control protocol (LCP) and Network Control Protocol (NCP).

VoIP—Voice over IP. The capability to carry normal telephony-style voice over an IP-based internet with plain old telephone service (POTS)-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network. In VoIP, the digital signal processor (DSP) segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with International Telecommunication Union Telecommunication Standardization Sector (ITU-T) specification H.323.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Pulse, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, and Flip Gift Card are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Fast Step, Follow Me Browsing, FormShare, GainMaker, GigaDrive, HomeLink, iLYNX, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo 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. (0908R)

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

© 2005–2008 Cisco Systems, Inc. All rights reserved.

