

Network Working Group
Request for Comments: 3518
Obsoletes: 2878
Category: Standards Track

M. Higashiyama
Anritsu
F. Baker
T. Liao
Cisco Systems
April 2003

Point-to-Point Protocol (PPP) Bridging Control Protocol (BCP)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

The Point-to-Point Protocol (PPP) provides a standard method for transporting multi-protocol datagrams over point-to-point links. PPP defines an extensible Link Control Protocol (LCP) and proposes a family of Network Control Protocols (NCP) for establishing and configuring different network-layer protocols.

This document defines the NCP for establishing and configuring Remote Bridging for PPP links.

This document obsoletes RFC 2878, which was based on the IEEE 802.1D-1993 MAC Bridge. This document extends that specification by improving support for bridge control packets.

Table of Contents

1.	Historical Perspective	2
1.1	Requirements Keywords	3
2.	Methods of Bridging	3
2.1	Transparent Bridging	3
2.2	Remote Transparent Bridging	4
2.3	Source Routing	5
2.4	Remote Source Route Bridging	6
2.5	SR-TB Translational Bridging	7
3.	Traffic Services	7

3.1	LAN Frame Checksum Preservation	7
3.2	Traffic having no LAN Frame Checksum	7
3.3	Tinygram Compression	8
3.4	Virtual LANs	8
3.5	Bridge Control Packet Indicator	9
4.	A PPP Network Control Protocol for Bridging	10
4.1	Sending Bridge Frames	11
4.1.1	Maximum Receive Unit Considerations	11
4.1.2	Loopback and Link Quality Monitoring	11
4.1.3	Message Sequence	11
4.1.4	Separation of Spanning Tree Domains	12
4.2	Bridged LAN Traffic in IEEE 802 Untagged Frame ...	13
4.3	Bridged LAN Traffic in IEEE 802 Tagged Frame	17
4.4	Bridge management protocol data unit	21
5.	BCP Configuration Options	22
5.1	Bridge-Identification	22
5.2	Line-Identification	24
5.3	MAC-Support	25
5.4	Tinygram-Compression	26
5.5	MAC-Address	27
5.6	Spanning Tree Protocol (old formatted)	28
5.7	IEEE-802-Tagged-Frame	30
5.8	Management-Inline	31
5.9	Bridge-Control-Packet-Indicator	32
6.	Changes From RFC 2878	33
7.	Security Considerations	33
8.	Intellectual Property Notice	33
9.	IANA Considerations	34
10.	Acknowledgments	34
	Appendices	35
A.	Spanning Tree Bridge PDU (old formatted)	35
B.	Tinygram-Compression Pseudo-Code	36
	References	38
	Authors' Addresses	39
	Full Copyright Statement.....	40

1. Historical Perspective

Two basic algorithms are ambient in the industry for Bridging of Local Area Networks. The more common algorithm is called "Transparent Bridging", and has been standardized for Extended LAN configurations by IEEE 802.1. The other is called "Source Route Bridging", and is prevalent on IEEE 802.5 Token Ring LANs.

The IEEE has combined these two methods into a device called a Source Routing Transparent (SRT) bridge, which concurrently provides both Source Route and Transparent bridging. Transparent and SRT bridges are specified in IEEE standard 802.1D-1998 [8].

Although IEEE committee 802.1G is addressing remote bridging [2], neither standard directly defines the mechanisms for implementing remote bridging. Technically, that would be beyond the IEEE 802 committee's charter. However, both 802.1D and 802.1G allow for it. The implementor may model the line either as a component within a single MAC Relay Entity, or as the LAN media between two remote bridges.

The original IEEE 802.1D is augmented by IEEE 802.1Q [9] to provide support for Virtual LAN. Virtual LAN is an integral feature of switched LAN networks.

1.1 Requirements Keywords

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [12].

2. Methods of Bridging

2.1. Transparent Bridging

As a favor to the uninitiated, let us first describe Transparent Bridging. Essentially, the bridges in a network operate as isolated entities, largely unaware of each others' presence. A Transparent Bridge maintains a Forwarding Database consisting of

{address, interface}

or

{address, interface, VLAN ID}

records, by saving the Source Address of each LAN transmission that it receives, along with the interface identifier for the interface it was received on. Bridges which support Virtual LANs additionally keep the Virtual LAN ID in their forwarding database. It goes on to check whether the Destination Address is in the database, and if so, either discards the message when the destination and source are located at the same interface, or forwards the message to the indicated interface. A message whose Destination Address is not found in the table is forwarded to all interfaces except the one it was received on. This behavior applies to Broadcast/Multicast frames as well.

The obvious fly in the ointment is that redundant paths in the network cause indeterminate (nay, all too determinate) forwarding behavior to occur. To prevent this, a protocol called the Spanning Tree Protocol is executed between the bridges to detect and logically remove redundant paths from the network.

One system is elected as the "Root", which periodically emits a message called a Bridge Protocol Data Unit (BPDU), heard by all of its neighboring bridges. Each of these modifies and passes the BPDU on to its neighbors, until it arrives at the leaf LAN segments in the network (where it dies, having no further neighbors to pass it along), or until the message is stopped by a bridge which has a superior path to the "Root". In this latter case, the interface the BPDU was received on is ignored (it is placed in a Hot Standby status, no traffic is emitted onto it except the BPDU, and all traffic received from it is discarded), until a topology change forces a recalculation of the network.

To establish Virtual LANs in an environment of multiple bridges, GVRP (GARP VLAN Registration Protocol) is executed between bridges to exchange Virtual LAN information. GVRP provides a mechanism to dynamically establish and update their knowledge of the set of Virtual LANs that currently have active members.

To reduce unnecessary multicast flooding in the network, bridges exchange group MAC addresses using the GARP Multicast Registration Protocol. GMRP provides a mechanism so that bridges can know which multicast frames should be forwarded on each port.

2.2. Remote Transparent Bridging

There exist two basic sorts of bridges -- those that interconnect LANs directly, called Local Bridges, and those that interconnect LANs via an intermediate medium such as a leased line, called Remote Bridges. PPP may be used to connect Remote Bridges.

The IEEE 802.1G Remote MAC Bridging committee has proposed a model of a Remote Bridge in which a set of two or more Remote Bridges that are interconnected via remote lines are termed a Remote Bridge Group. Within a Group, a Remote Bridge Cluster is dynamically formed through execution of the spanning tree as the set of bridges that may pass frames among each other.

This model bestows on the remote lines the basic properties of a LAN, but does not require a one-to-one mapping of lines to virtual LAN segments. For instance, the model of three interconnected Remote Bridges, A, B and C, may be that of a virtual LAN segment between A and B and another between B and C. However, if a line exists between

Remote Bridges B and C, a frame could actually be sent directly from B to C, as long as there was the external appearance that it had travelled through A.

IEEE 802.1G thus allows for a great deal of implementation freedom for features such as route optimization and load balancing, as long as the model is maintained.

For simplicity, we discuss Remote Bridging in this document in terms of two Remote Bridges connected by a single line.

2.3. Source Routing

The IEEE 802.1D Committee has standardized Source Routing for any MAC Type that allows its use. Currently, MAC Types that support Source Routing are FDDI and IEEE 802.5 Token Ring.

The IEEE standard defines Source Routing only as a component of an SRT bridge. However, many bridges have been implemented which are capable of performing Source Routing alone. These are most commonly implemented in accordance either with the IBM Token-Ring Network Architecture Reference [1] or with the Source Routing Appendix of IEEE 802.1D-1998 [8].

In the Source Routing approach, the originating system has the responsibility of indicating the path that the message should follow. It does this, if the message is directed off of the local segment, by including a variable length MAC header extension called the Routing Information Field (RIF). The RIF consists of one 16-bit word of flags and parameters, followed by zero or more segment-and-bridge identifiers. Each bridge en route determines from this source route list whether it should accept the message and how to forward it.

In order to discover the path to a destination, the originating system transmits an Explorer frame. An All-Routes Explorer (ARE) frame follows all possible paths to a destination. A Spanning Tree

Explorer (STE) frame follows only those paths defined by Bridge ports that the Spanning Tree Algorithm has put in Forwarding state. Port states do not apply to ARE or Specifically-Routed Frames. The destination system replies to each copy of an ARE frame with a Specifically-Routed Frame, and to an STE frame with an ARE frame. In either case, the originating station may receive multiple replies, from which it chooses the route it will use for future Specifically-Routed Frames.

The algorithm for Source Routing requires the bridge to be able to identify any interface by its segment-and-bridge identifier. When a packet is received that has the RIF present, a boolean in the RIF is inspected to determine whether the segment-and-bridge identifiers are to be inspected in "forward" or "reverse" sense. In its search, the bridge looks for the segment-and-bridge identifier of the interface the packet was received on, and forwards the packet toward the segment identified in the segment-and-bridge identifier that follows it.

GVRP and GMRP are available and effective on Source Routing networks.

2.4. Remote Source Route Bridging

There is no Remote Source Route Bridge proposal in IEEE 802.1 at this time, although many vendors ship remote Source Routing Bridges.

We allow for modelling the line either as a connection residing between two halves of a "split" Bridge (the split-bridge model), or as a LAN segment between two Bridges (the independent-bridge model). In the latter case, the line requires a LAN Segment ID.

By default, PPP Source Route Bridges use the independent-bridge model. This requirement ensures interoperability in the absence of option negotiation. In order to use the split-bridge model, a system MUST successfully negotiate the Bridge-Identification Configuration Option.

Although no option negotiation is required for a system to use the independent-bridge model, it is strongly recommended that systems using this model negotiate the Line-Identification Configuration Option. Doing so will verify correct configuration of the LAN Segment Id assigned to the line.

When two PPP systems use the split-bridge model, the system that transmits an Explorer frame onto the PPP link MUST update the RIF on behalf of the two systems. The purpose of this constraint is to ensure interoperability and to preserve the simplicity of the bridging algorithm. For example, if the receiving system did not know whether the transmitting system had updated the RIF, it would have to scan the RIF and decide whether to update it. The choice of the transmitting system for the role of updating the RIF allows the system receiving the frame from the PPP link to forward the frame without processing the RIF.

Given that source routing is configured on a line or set of lines, the specifics of the link state with respect to STE frames are defined by the Spanning Tree Protocol in use. Choice of the split-

bridge or independent-bridge model does not affect spanning tree operation. In both cases, the spanning tree protocol is executed on the two systems independently.

2.5. SR-TB Translational Bridging

IEEE 802 is not currently addressing bridges that translate between Transparent Bridging and Source Routing. For the purposes of this standard, such a device is either a Transparent or a Source Routing bridge, and will act on the line in one of these two ways, just as it does on the LAN.

3. Traffic Services

Several services are provided for the benefit of different system types and user configurations. These include LAN Frame Checksum Preservation, LAN Frame Checksum Generation, Tinygram Compression, and the identification of closed sets of LANs.

3.1. LAN Frame Checksum Preservation

IEEE 802.1 stipulates that the Extended LAN must enjoy the same probability of undetected error that an individual LAN enjoys. Although there has been considerable debate concerning the algorithm, no other algorithm has been proposed than having the LAN Frame Checksum received by the ultimate receiver be the same value calculated by the original transmitter. Achieving this requires, of course, that the line protocols preserve the LAN Frame Checksum from end to end. The protocol is optimized towards this approach.

3.2. Traffic having no LAN Frame Checksum

The fact that the protocol is optimized towards LAN Frame Checksum preservation raises twin questions: "What is the approach to be used by systems which, for whatever reason, cannot easily support Frame Checksum preservation?" and "What is the approach to be used when the system originates a message, which therefore has no Frame Checksum precalculated?".

Surely, one approach would be to require stations to calculate the Frame Checksum in software if hardware support were unavailable; this would meet with profound dismay, and would raise serious questions of interpretation in a Bridge/Router.


However, stations which implement LAN Frame Checksum preservation must already solve this problem, as they do originate traffic. Therefore, the solution adopted is that messages which have no Frame Checksum are tagged and carried across the line.

E1, E2, E3, and E4 are Ethernet LANs (or Token Ring, FDDI, etc.). W1 is a WAN (PPP over T1). B1 and B2 are MAC level bridges.

You want End Stations on E1 and E3 to communicate, and you want End Stations on E2 and E4 to communicate, but you do not want End Stations on E1 and E3 to communicate with End Stations on E2 and E4.

This is true for Unicast, Multicast, and Broadcast traffic. If a broadcast datagram originates on E1, you want it only to be propagated to E3, and not on E2 or E4.

Another way of looking at it is that E1 and E3 form a Virtual LAN, and E2 and E4 form a Virtual LAN, as if the following configuration were actually being used:



3.5. Bridge Control Packet Indicator

The Bridge Control Packet Indicator option is used to classify bridge control packets such as Spanning Tree BPDUs, GARP PDUs, etc. Protocols such as STP and GARP is to the bridging world as OSPF or BGP is to the routing world. Just as IP route update packets are marked with an IP precedence value of 6 or 7 and given preferential forwarding treatment [13], bridge control packets are marked in a similar fashion with the Bridge Control Packet Indicator bit.

If the Bridge Control Packet Indicator option is enabled, a system MUST set a packet's Bridge Control Packet Indicator bit in the flags field to 1 if and only if it is an outgoing bridge control frame. Furthermore, a system MUST avoid dropping or significantly delaying bridge control packets.

If the Bridge Control Packet Indicator option is disabled, a system MUST set the Bridge Control Packet Indicator bit to 0 for all frames. This preserves backward compatibility with RFC 2878 [14]. However, even if this option is disabled, a system SHOULD still avoid dropping or significantly delaying bridge control packets. This can be achieved through parsing the Destination MAC address field.

4. A PPP Network Control Protocol for Bridging

The Bridging Control Protocol (BCP) is responsible for configuring, enabling and disabling the bridge protocol modules on both ends of the point-to-point link. BCP uses the same packet exchange mechanism as the Link Control Protocol. BCP packets may not be exchanged until PPP has reached the Network-Layer Protocol phase. BCP packets received before this phase is reached SHOULD be silently discarded.

The Bridging Control Protocol is exactly the same as the Link Control Protocol [6] with the following exceptions:

Frame Modifications

The packet may utilize any modifications to the basic frame format which have been negotiated during the Link Establishment phase.

Implementations SHOULD NOT negotiate Address-and-Control-Field-Compression or Protocol-Field-Compression on other than low speed links.

Data Link Layer Protocol Field

Exactly one BCP packet is encapsulated in the PPP Information field, where the PPP Protocol field indicates type hex 8031 (BCP).

Code field

Only Codes 1 through 7 (Configure-Request, Configure-Ack, Configure-Nak, Configure-Reject, Terminate-Request, Terminate-Ack and Code-Reject) are used. Other Codes SHOULD be treated as unrecognized and SHOULD result in Code-Rejects.

Timeouts

BCP packets may not be exchanged until PPP has reached the Network-Layer Protocol phase. An implementation SHOULD be prepared to wait for Authentication and Link Quality Determination to finish before timing out waiting for a Configure-Ack or other response. It is suggested that an implementation give up only after user intervention or a configurable amount of time.

Configuration Option Types

BCP has a distinct set of Configuration Options, which are defined in this document.

4.1. Sending Bridge Frames

Before any Bridged LAN Traffic or BPDUs may be communicated, PPP MUST reach the Network-Layer Protocol phase, and the Bridging Control Protocol MUST reach the Opened state.

Exactly one Bridged LAN Traffic or BPDU is encapsulated in the PPP Information field, where the PPP Protocol field indicates type hex 0031 (Bridged PDU).

4.1.1. Maximum Receive Unit Considerations

The maximum length of a Bridged datagram transmitted over a PPP link is the same as the maximum length of the Information field of a PPP encapsulated packet. Since there is no standard method for fragmenting and reassembling Bridged PDUs, PPP links supporting Bridging MUST negotiate an MRU large enough to support the MAC Types that are later negotiated for Bridging support. Because they include the MAC headers, even bridged Ethernet frames are larger than the default PPP MRU of 1500 octets.

4.1.2. Loopback and Link Quality Monitoring

It is strongly recommended that PPP Bridge Protocol implementations utilize Magic Number Loopback Detection and Link-Quality-Monitoring. The 802.1 Spanning Tree protocol, which is integral to both Transparent Bridging and Source Routing (as standardized), is unidirectional during normal operation. Configuration BPDUs emanate from the Root system in the general direction of the leaves, without any reverse traffic except in response to network events.

4.1.3. Message Sequence

The multiple link case requires consideration of message sequentiality. The transmitting system may determine either that the protocol being bridged requires transmissions to arrive in the order of their original transmission, and enqueue all transmissions on a given conversation onto the same link to force order preservation, or that the protocol does NOT require transmissions to arrive in the order of their original transmission, and use that knowledge to optimize the utilization of several links, enqueueing traffic to multiple links to minimize delay.

In the absence of such a determination, the transmitting system MUST act as though all protocols require order preservation. Many protocols designed primarily for use on a single LAN require order preservation.

PPP Multilink [7] and its multi-class extension [11] may be used to allow the use of multiple PPP links between a pair of systems without loss of message sequentiality. It treats the group of links as a single link with speed equal to the sum of the speeds of the links in the group.

4.1.4. Separation of Spanning Tree Domains

It is conceivable that a network manager might wish to inhibit the exchange of BPDUs on a link in order to logically divide two regions into separate Spanning Trees with different Roots (and potentially different Spanning Tree implementations or algorithms). In order to do that, he should configure both ends to not exchange BPDUs on a link. An implementation that does not support any spanning tree protocol MUST silently discard any received IEEE 802.1D BDU packets.

If a bridge is connected to an old BCP bridge [10], the other bridge cannot operate according to this specification. Options are therefore to decide that:

- (a) If the bridge wants to terminate the connection, it sends a Terminate-Request and terminate the connection.
- (b) If the bridge wants to run the connection but not receive old BPDUs, its only option is to run without spanning tree on the link at all, which is dangerous. It should Configure-Reject the option and advise the network administration that it has done so.
- (c) If the bridge chooses to be entirely backward compatible, it sends Configure-Ack and operates in the manner described in Appendix A.

In the event that both the new Management-Inline Option and the Spanning-Tree-Protocol-Configuration Option are configure-rejected, indicating that the peer implements no spanning tree protocol at all and doesn't understand the options, it is an incomplete implementation. For safety reasons the system should cease attempting to configure bridging, and log the fact. If the peer was configure-rejecting the options in order to disable spanning tree entirely, it understood the option but could not within its configuration comply. It should have sent the Spanning-Tree-Protocol-Configuration Option with the value NULL.


Implementations SHOULD implement a backward compatibility mode.

4.2. Bridged LAN Traffic (IEEE 802 Untagged Frame)


For Bridging LAN traffic, the format of the frame on the line is shown below. This format is used if the traffic does not include VLAN ID and priority.

The fields are transmitted from left to right.

802.3 Frame format (IEEE 802 Un-tagged Frame)



802.4/802.5/FDDI Frame format (IEEE 802 Un-tagged Frame)



Address and Control

As defined by the framing in use.

PPP Protocol

0x0031 for PPP Bridging

Flags

- bit F: Set if the LAN FCS Field is present
- bit 0: reserved, must be zero
- bit Z: Set if IEEE 802.3 Pad must be zero filled to minimum size
- bit B: Set if the frame is a bridge control packet. See section 3.5 for details.

Pads

Any PPP frame may have padding inserted in the "Optional Data Link Layer Padding" field. This number tells the receiving system how many pad octets to strip off.

MAC Type

Up-to-date values of the MAC Type field are specified in the most recent "Assigned Numbers" RFC [4]. Current values are assigned as follows:

0:	reserved	
1:	IEEE 802.3/Ethernet	with canonical addresses
2:	IEEE 802.4	with canonical addresses
3:	IEEE 802.5	with non-canonical addresses
4:	FDDI	with non-canonical addresses
5-10:	reserved	
11:	IEEE 802.5	with canonical addresses
12:	FDDI	with canonical addresses

"Canonical" is the address format defined as standard address representation by the IEEE. In this format, the bit within each byte that is to be transmitted first on a LAN is represented as the least significant bit. In contrast, in non-canonical form, the bit within each byte that is to be transmitted first is represented as the most-significant bit. Many LAN interface implementations use non-canonical form. In both formats, bytes are represented in the order of transmission.

If an implementation supports a MAC Type that is the higher-numbered format of that MAC Type, then it MUST also support the lower-numbered format of that MAC Type. For example, if an implementation supports FDDI with canonical address format, then it MUST also support FDDI with non-canonical address format. The purpose of this requirement is to provide backward compatibility with earlier versions of this specification.

A system MUST NOT transmit a MAC Type numbered higher than 4 unless it has received from its peer a MAC-Support Configuration Option indicating that the peer is willing to receive frames of that MAC Type.

Frame Control

On 802.4, 802.5, and FDDI LANs, there are a few octets preceding the Destination MAC Address, one of which is protected by the FCS.

The MAC Type of the frame determines the contents of the Frame Control field. A pad octet is present to provide 32-bit packet alignment.

Destination MAC Address

As defined by the IEEE. The MAC Type field defines the bit ordering.

Source MAC Address

As defined by the IEEE. The MAC Type field defines the bit ordering.

LLC data

This is the remainder of the MAC frame which is (or would be were it present) protected by the LAN FCS.

For example, the 802.5 Access Control field, and Status Trailer are not meaningful to transmit to another ring, and are omitted.

LAN FCS

If present, this is the LAN FCS which was calculated by (or which appears to have been calculated by) the originating station. If the LAN FCS flag is not set, then this field is not present, and the PDU is four octets shorter.

Optional Data Link Layer Padding


Any PPP frame may have padding inserted between the Information field and the Frame FCS. The Pads field contains the length of this padding, which may not exceed 15 octets.

The PPP LCP Extensions [5] specify a self-describing pad. Implementations are encouraged to set the Pads field to zero, and use the self-describing pad instead.

Frame FCS

Mentioned primarily for clarity. The FCS used on the PPP link is separate from and unrelated to the LAN FCS.

802.4/802.5/FDDI Frame format (IEEE 802 Tagged Frame)



Address and Control

As defined by the framing in use.

PPP Protocol

0x0031 for PPP Bridging

Flags

- bit F: Set if the LAN FCS Field is present
- bit 0: reserved, must be zero
- bit Z: Set if IEEE 802.3 Pad must be zero filled to minimum size
- bit B: Set if the frame is a bridge control packet. See section 3.5 for details.

Pads

Any PPP frame may have padding inserted in the "Optional Data Link Layer Padding" field. This number tells the receiving system how many pad octets to strip off.

MAC Type

Up-to-date values of the MAC Type field are specified in the most recent "Assigned Numbers" RFC [4]. Current values are assigned as follows:

0:	reserved	
1:	IEEE 802.3/Ethernet	with canonical addresses
2:	IEEE 802.4	with canonical addresses
3:	IEEE 802.5	with non-canonical addresses
4:	FDDI	with non-canonical addresses
5-10:	reserved	
11:	IEEE 802.5	with canonical addresses
12:	FDDI	with canonical addresses

"Canonical" is the address format defined as standard address representation by the IEEE. In this format, the bit within each byte that is to be transmitted first on a LAN is represented as the least significant bit. In contrast, in non-canonical form, the bit within each byte that is to be transmitted first is represented as the most-significant bit. Many LAN interface implementations use non-canonical form. In both formats, bytes are represented in the order of transmission.

If an implementation supports a MAC Type that is the higher-numbered format of that MAC Type, then it MUST also support the lower-numbered format of that MAC Type. For example, if an implementation supports FDDI with canonical address format, then it MUST also support FDDI with non-canonical address format. The purpose of this requirement is to provide backward compatibility with earlier versions of this specification.

A system MUST NOT transmit a MAC Type numbered higher than 4 unless it has received from its peer a MAC-Support Configuration Option indicating that the peer is willing to receive frames of that MAC Type.

Frame Control

On 802.4, 802.5, and FDDI LANs, there are a few octets preceding the Destination MAC Address, one of which is protected by the FCS.

The MAC Type of the frame determines the contents of the Frame Control field. A pad octet is present to provide 32-bit packet alignment.

Destination MAC Address

As defined by the IEEE. The MAC Type field defines the bit ordering.

Source MAC Address

As defined by the IEEE. The MAC Type field defines the bit ordering.

Pri

3 bit priority value as defined by IEEE 802.1D.

C

Canonical flag as defined by IEEE 802.1Q. It must be set if RIF data is present in the LLC data.

VLAN ID

12 bit VLAN identifier number as defined by IEEE 802.1Q.

LLC data

This is the remainder of the MAC frame which is (or would be were it present) protected by the LAN FCS.

For example, the 802.5 Access Control field, and Status Trailer are not meaningful to transmit to another ring, and are omitted.

LAN FCS

If present, this is the LAN FCS which was calculated by (or which appears to have been calculated by) the originating station. If the LAN FCS flag is not set, then this field is not present, and the PDU is four octets shorter.

Optional Data Link Layer Padding

Any PPP frame may have padding inserted between the Information field and the Frame FCS. The Pads field contains the length of this padding, which may not exceed 15 octets.

The PPP LCP Extensions [5] specify a self-describing pad. Implementations are encouraged to set the Pads field to zero, and use the self-describing pad instead.

Frame FCS

Mentioned primarily for clarity. The FCS used on the PPP link is separate from and unrelated to the LAN FCS.

4.4. Bridge protocols and GARP protocols

To avoid network loops and improve redundancy, Bridges exchange a Spanning Tree Protocol data unit known as BPDU. Bridges also exchange a Generic Attributes Registration Protocol data unit to carry the GARP VLAN Registration Protocol (GVRP) data and GARP Multicast Registration Protocol (GMRP). GVRP allow the Bridges to create VLAN groups dynamically. GMRP allows bridges to filter Multicast data if the receiver is absent from the network. These Bridge protocols include Spanning Tree Protocol and GARP protocols data units are carried with a special destination address assigned by the IEEE.

These bridge protocols data units and GARP protocol data units must be carried in the frame format shown in section 4.2 or 4.3. The Bridge that receives these data units identifies these protocols based on the destination address in the frame format, just like the operation of receiving frames from a LAN segment.

Bridge protocols and GARP protocols data units MUST be recognized by checking the destination addresses, which are assigned by IEEE.

01-80-c2-00-00-00	Bridge Group Address (used by STP)
01-80-c2-00-00-01	IEEE Std. 802.3x Full Duplex PAUSE operation
01-80-c2-00-00-10	Bridge Management Group Address
01-80-c2-00-00-20	GARP Multicast Registration Protocol (GMRP)
01-80-c2-00-00-21	GARP VLAN Registration Protocol (GVRP)

But there is one exception to this rule: if the bridge is connected to an old BCP bridge [10] and can support backward compatibility, it MUST send the BPDU in the old format described in Appendix A.

5. BCP Configuration Options

BCP Configuration Options allow modifications to the standard characteristics of the network-layer protocol to be negotiated. If a Configuration Option is not included in a Configure-Request packet, the default value for that Configuration Option is assumed.

BCP uses the same Configuration Option format defined for LCP [6], with a separate set of Options.

Up-to-date values of the BCP Option Type field are specified in the most recent "Assigned Numbers" RFC [4]. Current values are assigned as follows:

1	Bridge-Identification
2	Line-Identification
3	MAC-Support
4	Tinygram-Compression
5	LAN-Identification (obsoleted)
6	MAC-Address
7	Spanning-Tree-Protocol (old formatted)
8	IEEE 802 Tagged Frame
9	Management Inline
10	Bridge Control Packet Indicator

5.1. Bridge-Identification

Description

The Bridge-Identification Configuration Option is designed for use when the line is an interface between half bridges connecting virtual or physical LAN segments. Since these remote bridges are modeled as a single bridge with a strange internal interface, each remote bridge needs to know the LAN segment and bridge numbers of the adjacent remote bridge. This option MUST NOT be included in the same Configure-Request as the Line-Identification option.

The Source Routing Route Descriptor and its use are specified by the IEEE 802.1D Appendix on Source Routing. It identifies the segment to which the interface is attached by its configured segment number, and itself by bridge number on the segment.

The two half bridges MUST agree on the bridge number. If a bridge number is not agreed upon, the Bridging Control Protocol MUST NOT enter the Opened state.


Since mismatched bridge numbers are indicative of a configuration error, a correct configuration requires that either the bridge declare the misconfiguration or choose one of the options. To allow two systems to proceed to the Opened state despite a mismatch, a system MAY change its bridge number to the higher of the two numbers. A higher-numbered system MUST NOT change its bridge number to a lower number. It should, however, inform the network administration of the misconfiguration in any case.

By default, a system that does not negotiate this option is assumed to be configured not to use the model of the two systems as two halves of a single source-route bridge. It is instead assumed to be configured to use the model of the two systems as two independent bridges.

Example

If System A announces LAN Segment AAA, Bridge #1, and System B announces LAN Segment BBB, Bridge #1, then the resulting Source Routing configuration (read in the appropriate direction) is then AAA,1,BBB.

A summary of the Bridge-Identification Option format is shown below. The fields are transmitted from left to right.



Type

1

Length

4

LAN Segment Number

A 12-bit number identifying the LAN segment, as defined in the IEEE 802.1D Source Routing Specification.

Bridge Number

A 4-bit number identifying the bridge on the LAN segment, as defined in the IEEE 802.1D Source Routing Specification.

5.2. Line-Identification

Description

The Line-Identification Configuration Option is designed for use when the line is assigned a LAN segment number as though it were a two system LAN segment in accordance with the Source Routing algorithm.


The Source Routing Route Descriptor and its use are specified by the IEEE 802.1D Appendix on Source Routing. It identifies the segment to which the interface is attached by its configured segment number, and itself by bridge number on the segment.

The two bridges MUST agree on the LAN segment number. If a LAN segment number is not agreed upon, the Bridging Control Protocol MUST NOT enter the Opened state.

Since mismatched LAN segment numbers are indicative of a configuration error, a correct configuration requires that either the bridge declare the misconfiguration or choose one of the options. To allow two systems to proceed to the Opened state despite a mismatch, a system MAY change its LAN segment number to the higher of the two numbers. A higher-numbered system MUST NOT change its LAN segment number to a lower number. It should, however, inform the network administration of the misconfiguration in any case.

By default, a system that does not negotiate this option is assumed to have its LAN segment number correctly configured by the user.

A summary of the Line-Identification Option format is shown below. The fields are transmitted from left to right.



Type

2

Length

4

LAN Segment Number

A 12-bit number identifying the LAN segment, as defined in the IEEE 802.1D Source Routing Specification.

Bridge Number

A 4-bit number identifying the bridge on the LAN segment, as defined in the IEEE 802.1D Source Routing Specification.

5.3. MAC-Support

Description

The MAC-Support Configuration Option is provided to permit implementations to indicate the sort of traffic they are prepared to receive. Negotiation of this option is strongly recommended.

By default, when an implementation does not announce the MAC Types that it supports, all MAC Types are sent by the peer which are capable of being transported given other configuration parameters. The receiver will discard those MAC Types that it does not support.


A device supporting a 1600 octet MRU might not be willing to support 802.5, 802.4 or FDDI, which each support frames larger than 1600 octets.

By announcing the MAC Types it will support, an implementation is advising its peer that all unspecified MAC Types will be discarded. The peer MAY then reduce bandwidth usage by not sending the unsupported MAC Types.

Announcement of support for multiple MAC Types is accomplished by placing multiple options in the Configure-Request.

The nature of this option is advisory only. This option MUST NOT be included in a Configure-Nak.

A summary of the MAC-Support Option format is shown below. The fields are transmitted from left to right.



Type

3

Length

3

MAC Type

One of the values of the PDU MAC Type field (previously described in the "Bridged LAN Traffic" section) that this system is prepared to receive and service.

5.4. Tinygram-Compression

Description


This Configuration Option permits the implementation to indicate support for Tinygram compression.

Not all systems are prepared to make modifications to messages in transit. On high speed lines, it is probably not worth the effort.

This option MUST NOT be included in a Configure-Nak if it has been received in a Configure-Request. This option MAY be included in a Configure-Nak in order to prompt the peer to send the option in its next Configure-Request.

By default, no compression is allowed. A system which does not negotiate, or negotiates this option to be disabled, should never receive a compressed packet.

A summary of the Tinygram-Compression Option format is shown below. The fields are transmitted from left to right.



Type

4

Length

3

Enable/Disable

If the value is 1, Tinygram-Compression is enabled. If the value is 2, Tinygram-Compression is disabled, and no decompression will occur.

The implementations need not agree on the setting of this parameter. One may be willing to decompress and the other not.

5.5. MAC-Address

Description


The MAC-Address Configuration Option enables the implementation to announce its MAC address or have one assigned. The MAC address is represented in IEEE 802.1 Canonical format, which is to say that the multicast bit is the least significant bit of the first octet of the address.

If the system wishes to announce its MAC address, it sends the option with its MAC address specified. When specifying a non-zero MAC address in a Configure-Request, any inclusion of this option in a Configure-Nak MUST be ignored.

If the implementation wishes to have a MAC address assigned, it sends the option with a MAC address of 00-00-00-00-00-00. Systems that have no mechanism for address assignment will Configure-Reject the option.

A Configure-Nak MUST specify a valid IEEE 802.1 format physical address; the multicast bit MUST be zero. It is strongly recommended (although not mandatory) that the "locally assigned address" bit (the second least significant bit in the first octet) be set, indicating a locally assigned address.

A summary of the MAC-Address Option format is shown below. The fields are transmitted from left to right.



Type

6

Length

8

MAC Byte

Six octets of MAC address in 802.1 Canonical order. For clarity, the position of the Local Assignment (L) and Multicast (M) bits are shown in the diagram.

5.6. Spanning-Tree-Protocol (old format)

Description

The Spanning-Tree-Protocol Configuration enables a Bridge to remain compatible with older implementations of BCP [10]. This configuration option is, however, incompatible with the Management-Inline option, which enables a bridge to implement the many protocols that IEEE now expects a bridge to be able to use.

If the peer rejects the Management-Inline configuration option, by sending configure-reject, it must be an implementation of [10], which is described in Appendix A. The system may optionally terminate the negotiation or offer to negotiate in that manner.


In this case, if both bridges support a spanning tree protocol, they MUST agree on the protocol to be supported. The old BPDU described in Appendix A MUST be used rather than the format shown in section 4.2 or 4.3. When the two disagree, the lower-numbered of the two spanning tree protocols should be used. To resolve the conflict, the system with the lower-numbered protocol SHOULD Configure-Nak the option, suggesting its own protocol for use. If

a spanning tree protocol is not agreed upon, except for the case in which one system does not support any spanning tree protocol, the Bridging Control Protocol MUST NOT enter the Opened state.

Most systems will only participate in a single spanning tree protocol. If a system wishes to participate simultaneously in more than one spanning tree protocol, it MAY include all of the appropriate protocol types in a single Spanning-Tree-Protocol Configuration Option. The protocol types MUST be specified in increasing numerical order. For the purpose of comparison during negotiation, the protocol numbers MUST be considered to be a single number. For instance, if System A includes protocols 01 and 03 and System B indicates protocol 03, System B should Configure-Nak and indicate a protocol type of 03 since 0103 is greater than 03.

By default, an implementation MUST either support the IEEE 802.1D spanning tree or support no spanning tree protocol. An implementation that does not support any spanning tree protocol MUST silently discard any received IEEE 802.1D BPDU packets, and MUST either silently discard or respond to other received BPDU packets with an LCP Protocol-Reject packet in this case.

A summary of the Spanning-Tree-Protocol Option format is shown below. The fields are transmitted from left to right.



Type

7

Length

2 octets plus 1 additional octet for each protocol that will be actively supported. Most systems will only support a single spanning tree protocol, resulting in a length of 3.

Protocol n

Each Protocol field is one octet and indicates a desired spanning tree protocol. Up-to-date values of the Spanning-Tree-Protocol field are specified as PPP DLL numbers in the most recent "Assigned Numbers" RFC [4]. Current values are assigned as follows:

Value	Protocol
0	Null (no Spanning Tree protocol supported)
1	IEEE 802.1D spanning tree
2	IEEE 802.1G extended spanning tree protocol
3	IBM Source Route Spanning tree protocol
4	DEC LANbridge 100 Spanning tree protocol

5.7. IEEE-802-Tagged-Frame


Description

This configuration option permits the implementation to indicate support for IEEE 802 Tagged Frame. Negotiation of this option is strongly recommended.

A device supporting IEEE 802 Tagged Frame must be willing to support IEEE 802 Tagged Frame shown in section 4.3.

By default, IEEE 802 Tagged Frame is not supported. A system which does not negotiate, or negotiates this option to be disabled, should never receive a IEEE 802 Tagged Frame.

A summary of the IEEE 802 Tagged Frame Option format is shown below. The fields are transmitted from left to right.



Type

8

Length

3

Enable/Disable

If the value is 1, IEEE-802-Tagged-Frame is enabled. If the value is 2, IEEE-802-Tagged-Frame is disabled, and MUST not send any IEEE-802-Tagged-Frame packet.

5.8. Management-Inline

Description

The Management-Inline Configuration Option indicates that the system is willing to receive any IEEE-defined inter-bridge protocols, such as bridge protocol data units and GARP protocol data units, in the frame format shown in section 4.2 or 4.3.


Old BCP [10] implementations will use the negotiation procedure described in section 5.6. Implementations of this procedure will use this option to indicate compliance with the new BCP and may optionally negotiate the section 5.6 procedure, either on the same configure-request or in response to a configure-reject, as well. It is recommended that the configure-request only show this option when it is relevant, and that it reply with the Spanning-Tree-Protocol (old formatted) option if a configure-reject is received, as in the normal case one can expect it to be the quickest negotiation.

If a system receives a configure-request offering both alternatives, it should accept this procedure and reject the Spanning-Tree-Protocol (old format) option.

One can expect old BCP [10] implementations to not understand the option and issue a configure-reject.

By default, Management-Inline is not allowed. A system which does not negotiate, or negotiates this option to be disabled, should never receive a Bridge Protocol data unit or GARP protocol data unit inline.

A summary of the Management-Inline Option format is shown below. The fields are transmitted from left to right.



Type

9

Length

2

5.9 Bridge-Control-Packet-Indicator


Description

This configuration option permits the implementation to indicate support for Bridge Control Packet Indicator. Negotiation of this option is strongly recommended.

By default, Bridge Control Packet Indicator is not supported. Negotiating this option enables the Bridge Control Packet Indicator. Not negotiating this option disables the Bridge Control Packet Indicator.

A system which does not negotiate MUST never send or receive a frame with the Bridge Control Packet Indicator bit set to 1.

A summary of the Bridge Control Packet Indicator option format is shown below. The fields are transmitted from left to right.



Type

10

Length

2

6. Changes From RFC 2878

This section enumerates changes made to the old RFC [14] to produce this document.

- (1) Add Bridge Control Packet Indicator to configuration option.
- (2) Modify meaning of one of the reserved bits in the flags field.

7. Security Considerations

This network control protocol compares the configurations of two devices and seeks to negotiate an acceptable subset of their intersection, to enable correct interoperation even in the presence of minor configuration or implementation differences. In the event that a major misconfiguration is detected, the negotiation will not complete successfully, resulting in the link coming down or not coming up. It is possible that if a bridged link comes up with a rogue peer, network information may be learned from forwarded multicast traffic, or denial of service attacks may be created by closing loops that should be detected and isolated or by offering rogue load.

Such attacks are not isolated to this NCP; any PPP NCP is subject to attack when connecting to a foreign or compromised device. However, no situations arise which are not common to all NCPs; any NCP that comes up with a rogue peer is subject to snooping and other attacks. Therefore, it is recommended that links on which this may happen should be configured to use PPP authentication during the LCP start-up phase.

8. Intellectual Property Notice

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

9. IANA Considerations

This document proposes one new BCP option number to be maintained by the IANA. This option, described in Section 5.9, is Bridge-Control-Packet-Indicator. The IANA has assigned the value 10 for this option.

10. Acknowledgments

This document is a product of the Point-to-Point Protocol Extensions Working Group.

This document is based on the PPP Bridging Control Protocol, RFC 2878 [14], edited by Higashiyama and Baker and produced by the Point-to-Point Protocol Extensions Working Group. It extends that document by providing support for Bridge Control Packet Indicator as outlined in section 3.5 and 5.9.

Appendices

A. Spanning Tree Bridge PDU (old format)


By default, Spanning Tree BPDUs MUST be encoded with a MAC or 802.2 LLC header as described in section 4.2 or 4.3 of this document. However, should the remote entity Configure-Reject the Management-Inline option, thereby indicating that it is a purely RFC 1638 compliant device, the local entity may subsequently encode BPDUs as described in section 4.3 of RFC 1638 provided that use of a suitable non-NULL STP protocol across the link is successfully negotiated using the (old) Spanning-Tree-Protocol option.

This is the Spanning Tree BPDU used in RFC 1638, without any MAC or 802.2 LLC header (these being functionally equivalent to the Address, Control, and PPP Protocol Fields). The LAN Pad and Frame Checksum fields are likewise superfluous and absent.

The Address and Control Fields are subject to LCP Address-and-Control-Field-Compression negotiation.

A PPP system which is configured to participate in a particular spanning tree protocol and receives a BPDU of a different spanning tree protocol SHOULD reject it with the LCP Protocol-Reject. A system which is configured not to participate in any spanning tree protocol MUST silently discard all BPDUs.

Spanning Tree Bridge PDU



Address and Control

As defined by the framing in use.

Spanning Tree Protocol

Up-to-date values of the Spanning-Tree-Protocol field are specified in the most recent "Assigned Numbers" RFC [4]. Current values are assigned as follows:

Value (in hex)	Protocol
0201	IEEE 802.1 (either 802.1D or 802.1G)
0203	IBM Source Route Bridge
0205	DEC LANbridge 100

The two versions of the IEEE 802.1 spanning tree protocol frames can be distinguished by fields within the BPDU data.

BPDU data

As defined by the specified Spanning Tree Protocol.

B. Tinygram-Compression Pseudo-Code

PPP Transmitter:

```

if (ZeroPadCompressionEnabled &&
    BridgedProtocolHeaderFormat == IEEE8023 &&
    PacketLength == Minimum8023PacketLength) {
/*
 * Remove any continuous run of zero octets preceding,
 * but not including, the LAN FCS, but not extending
 * into the MAC header.
 */
    Set (ZeroCompressionFlag);                /* Signal receiver */
    if (is_Set (LAN_FCS_Present)) {
        FCS = TrailingOctets (PDU, 4);        /* Store FCS */
        RemoveTrailingOctets (PDU, 4);        /* Remove FCS */
        while (PacketLength > 14 &&          /* Stop at MAC header or */
              TrailingOctet (PDU) == 0) /* last non-zero octet */
            RemoveTrailingOctets (PDU, 1); /* Remove zero octet */
        Appendbuf (PDU, 4, FCS);             /* Restore FCS */
    }
    else {
        while (PacketLength > 14 &&          /* Stop at MAC header */
              TrailingOctet (PDU) == 0) /* or last zero octet */
            RemoveTrailingOctets (PDU, 1); /* Remove zero octet */
    }
}

```

PPP Receiver:

```
if (ZeroCompressionFlag) {                               /* Flag set in header? */
/* Restoring packet to minimum 802.3 length */
  Clear (ZeroCompressionFlag);
  if (is_Set (LAN_FCS_Present)) {
    FCS = TrailingOctets (PDU, 4); /* Store FCS */
    RemoveTrailingOctets (PDU, 4); /* Remove FCS */
    Appendbuf (PDU, 60 - PacketLength, zeroes); /* Add zeroes */
    Appendbuf (PDU, 4, FCS); /* Restore FCS */
  }
  else {
    Appendbuf (PDU, 60 - PacketLength, zeroes); /* Add zeroes */
  }
}
```

References

- [1] IBM, "Token-Ring Network Architecture Reference", 3rd edition, September 1989.
- [2] IEEE 802.1, "Draft Standard 802.1G: Remote MAC Bridging", P802.1G/D7, December 30, 1992.
- [3] IEEE 802.1D-1993, "Media Access Control (MAC) Bridges", ISO/IEC 15802-3:1993 ANSI/IEEE Std 802.1D, 1993 edition., July 1993.
- [4] Reynolds, J., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, January 2002.
- [5] Simpson, W., "PPP LCP Extensions", RFC 1570, January 1994.
- [6] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [7] Sklower, K., Lloyd, B., McGregor, G., Carr, D. and T. Coradetti, "The PPP Multilink Protocol (MP)", RFC 1990, August 1996.
- [8] IEEE 802.1D-1998, "Information technology - Telecommunications and Information exchange between systems - Local and metropolitan area networks - Common Specifications - Part 3: Media Access Control (MAC) Bridges: Revision. This is a revision of ISO/IEC 10038: 1993, 802.1j-1992 and 802.6k-1992. It incorporates P802.11c, P802.1p and P802.12e." ISO/IEC 15802-3: 1998.
- [9] IEEE 802.1Q, ANSI/IEEE Standard 802.1Q, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks", 1998.
- [10] Baker, F. and R. Bowen, "PPP Bridging Control Protocol (BCP)", RFC 1638, June 1994.
- [11] Bormann, C., "The Multi-Class Extension to Multi-Link PPP", RFC 2686, September 1999.
- [12] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [13] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the "Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

- [14] Higashiyama, M. and F. Baker, "PPP Bridging Control Protocol (BCP)", RFC 2878, July 2000.

Authors' Addresses

Mitsuru Higashiyama
Anritsu Corporation
1800 Onna
Atsugi-shi
Kanagawa-prf.
243-8555 Japan

Phone: +81 (46) 296-6625
EMail: Mitsuru.Higashiyama@yy.anritsu.co.jp

Fred Baker
1121 Via Del Rey
Santa Barbara, California
93117 USA

Phone: (408) 526-4257
EMail: fred@cisco.com

Tawei Liao
cisco Systems, Inc.
170 W. Tasman Drive
San Jose, CA 95134

Phone: (408) 853-8905
EMail: tawei@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

