

Internet Draft
Document: [draft-bryant-pwe3-fr-compare-00.txt](#)
Expires: May 2002

Stewart Bryant
Lloyd Wood
Cisco Systems Ltd
October 2001

A Commentary on Approaches to Frame Relay Encapsulation in PWE3

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Abstract

This document compares and contrasts four approaches to Frame Relay encapsulation over a pseudo-wire that have been proposed in the Pseudo-Wire Edge-to-Edge (PWE3) IETF working group. This document then shows that the general-purpose encapsulation previously proposed for HDLC, Ethernet and VLAN is also the most efficient approach for Frame Relay.

Table of Contents

| | |
|--|--------------------|
| Status of this Memo..... | 1 |
| Abstract..... | 1 |
| 1. Overview..... | 2 |
| 2. The Martini encapsulation..... | 3 |
| 2.1 Comments on the Martini encapsulation..... | 3 |
| 3. The Kawa encapsulation..... | 4 |
| 3.1 Comments on the Kawa encapsulation..... | 5 |
| 4. Bryant Optimized Frame-Relay-specific control word..... | 6 |
| 4.1 Comments on Bryant Optimized Frame-Relay-specific control word..... | 7 |
| 5. General Pseudo-wire Payload Encapsulation..... | 7 |
| 5.1 Comments on General Pseudo-wire Payload Encapsulation..... | 8 |
| 6. Mapping between the Control Word and the Frame Relay Header..... | 8 |
| 6.1 Mapping between Martini control word and Frame Relay header..... | 8 |
| 6.2 Mapping between Kawa control word and Frame Relay header..... | 9 |
| 6.3 Mapping between Bryant optimized Frame-Relay-specific control word and Frame Relay header..... | 10 |
| 6.4 Mapping between the General Pseudo-wire Payload Encapsulation and Frame Relay..... | 11 |
| 7. Frame Relay header translation..... | 11 |
| 8. Conclusions..... | 13 |
| 9. IANA considerations..... | 13 |
| 10. Security considerations..... | 13 |
| 11. References..... | 14 |
| Authors' addresses..... | 14 |
| Full copyright statement..... | 15 |

[1.](#) Overview

The design of the pseudo-wire encapsulation header can have considerable impact on the performance of the system using it. Drafts describing four Frame Relay encapsulation approaches have been presented. This document compares and contrasts these approaches and analyses their computational efficiency. From this analysis it is evident that the general-purpose encapsulation proposed for HDLC, Ethernet and VLAN is also the most efficient approach for Frame Relay.


2. The Martini encapsulation

The following relevant text is extracted from [section 4.1](#) (Frame Relay) of [[MARTINI](#)] for convenient immediate reference.

"A Frame Relay PDU is transported without the Frame Relay header or the FCS. The control word is REQUIRED; however, its use is optional, although desirable. Use of the control word means that the ingress and egress LSRs follow the procedures below. If an ingress LSR chooses not to use the control word, it MUST set the flags in the control word to 0; if an egress LSR chooses to ignore the control word, it MUST set the Frame Relay control bits to 0.

"The BECN, FECN, DE and C/R bits are carried across the network in the control word. The edge routers that implement this document MAY, when either adding or removing the encapsulation described herein, change the BECN and/or FECN bits from zero to one in order to reflect congestion in the network that is known to the edge routers, and the D/E bit from zero to one to reflect marking from edge policing of the Frame Relay Committed Information Rate. The BECN, FECN, and D/E bits MUST NOT be changed from one to zero.

"The following is an example of a Frame Relay packet:



"

As a reminder, definitions of bits as used in these drafts and in [[Q.922](#)] are:

B is the BECN (Backward Explicit Congestion Notification) bit
F is the FECN (Forward Explicit Congestion Notification) bit
D is the DE (Discard Eligibility) bit
C is the C/R (Command/Response) bit

2.1 Comments on the Martini encapsulation

[MARTINI] takes the approach of copying Frame Relay payload-dependent bits to fields in the control word, stripping the Frame Relay header from the Frame Relay PDU, and reconstructing the header at the far end of the pseudo-wire. The reason for this approach is not given in [MARTINI].

| | | |
|--------|---|--------------|
| Bryant | internet-draft - expires May 2002 | 3 |
| | draft-bryant-pwe3-fr-compare-00.txt | October 2001 |

The [MARTINI] approach to handling Frame Relay headers is inconsistent with the approaches used in the same draft for other encapsulation types. The encapsulation and de-capsulation processes can be made more efficient if the Frame Relay header is transmitted along with its accompanying PDU, in a manner consistent with the HDLC, Ethernet and VLAN encapsulation types.

The layout and ordering of the B, F, D and C bits in the control word differ from the layout and ordering of these bits in the Frame Relay header. Resulting necessary bit manipulation in both encapsulation and decapsulation introduces unnecessary processing overhead in any implementation. We will show that this overhead can be avoided.

3. The Kawa encapsulation

The following relevant text is extracted from section 6.1 of [[KAWA](#)] for reference:

"FRoPW header structure is shown in Figure 3.




Figure 3 - FRoPW header structure

"The meaning of the fields is as follows:

Res (bits 0 to 2): Reserved bits. They are set to zero on transmission and ignored on reception.

"P - Payload Type (bit 3):

If set to zero then the payload field contains user's data else it contains network maintenance data.

"X, Y (bits 8 and 9):

1 0: First segment of a segmented FROPW packet

0 1: Last segment of a segmented FROPW packet

0 0: Neither the first nor the last segment of a segmented FroPW packet

1 1: Complete FROPW packet

"X and Y bits are used for segmentation and re-assembly of FroPW packet fragments when these capabilities are enabled in the two PEs terminating a PW.

"Length (bits 10 to 15):

The length field is used to pad short FROPW packets over Ethernet PSN. The length field is used as follows: If the length of the layer 2 frame (consisting of layer two control information and layer 2

Bryant internet-draft - expires May 2002

4

[draft-bryant-pwe3-fr-compare-00.txt](#)

October 2001

payload) is less than 64 bytes, the length field MUST be set to the FROPW packet length. Otherwise the length field MUST be set to zero. The value of the length field, if non-zero, is used to remove any padding.

"Sequence number (Bit 16 to 31):

If not used it is set to zero by the sender and ignored by the receiver. Otherwise it specifies the sequence number of a complete FROPW packet or a fragment. A circular list of sequence numbers is used. A sequence number takes a value from 1 to 65535 ($2^{16}-1$). Arithmetic modulo 2^{16} is used to generate a new sequence number.

"The sequence number must be used when segmentation and re-assembly are enabled between two peer PEs terminating a PSN tunnel. The sequence number may be used to detect out-of-order delivery of FROPW packets when the PSN does not guarantee in-order delivery."

3.1 Comments on the Kawa encapsulation

As with the [[MARTINI](#)] approach, the Frame Relay payload-dependent bits are copied to fields in the control word, the Frame Relay header is stripped from the Frame Relay PDU, and the header must be reconstructed at the far end of the pseudo-wire. Again, the reason for doing this is not given.

The grouping and ordering of the payload-dependent bits in the [[KAWA](#)] control word is more consistent with Q.922, allowing them to be processed as a three bit group (FBD) and a single bit (C). This

reduces the processing overhead in comparison with [\[MARTINI\]](#). However, the [\[KAWA\]](#) draft fails to take advantage of the additional performance gains to be made by closely following the layout specified in [\[Q.922\]](#).

We have found the naming and polarity of the fragmentation bits introduced in [\[KAWA\]](#) confusing. Use of the name 'X' for the first segment indicator conflicts with common usage of the symbol 'X' to indicate "don't care" about the value of the bit.

The polarity chosen fails to take advantage of the performance optimizations normally available in processor instruction sets. A common case is likely to be that the frame length is greater than 64 bytes, but less than the MTU. In the scheme described in [\[KAWA\]](#), this is indicated by setting bits 8..15 to the binary value <11000000>.

We propose that the sense of the X and Y bits be reversed, and renamed to notStart (A) and notEnd (B) bits respectively.

This is represented in the following two-bit state table:

| | | |
|--------|---|--------------|
| Bryant | internet-draft - expires May 2002 | 5 |
| | draft-bryant-pwe3-fr-compare-00.txt | October 2001 |

A, B (bits 8 and 9):

- 0 0: Complete packet (start and end)
- 0 1: First segment (start and not end)
- 1 1: Continuing segment (not start and not end)
- 1 0: Last segment (not start and end)

This leads to the use of an ordered stream: 0{1111}0 , where the common case of an unfragmented frame that is greater than 64 bytes, but less than the MTU, is sent as <00000000>. This is straightforward and optimal to detect in normal software implementations.

To provide space for the fragmentation bits, [\[KAWA\]](#) reduces the length of the length field from 8 bits to 6 bits. Some rationale needs to be provided that explains the reason for this approach rather than maintaining 8-bit length compatibility with the other encapsulation drafts, and using an additional two bits from the reserved field instead.

The need to fragment a payload for transmission across a pseudo-wire is common to all payload types. It would therefore seem appropriate to move this to a common PWE3 header to be defined, rather than to use it in a specific Frame Relay method and later reinvent it for

other payload types. This would restore the length field in [\[KAWA\]](#) to the expected 8 bits. The absence of a fragmentation mechanism in MPLS and IPv6 raises the question of whether fragmentation is a unique problem to PWE3 or a general problem that will need to be addressed by other transport types, such as [\[GRE\]](#). If it were a general problem it would be better to further abstract the fragmentation support to a separate fragmentation layer above the PSN and below PWE3.

The [\[KAWA\]](#) control word specifies the length field as a MUST, indicating that it is compulsory. This is wasteful of resources when the PSN is IP, because the IP payload length will always be present.

[4.](#) Bryant Optimized Frame-Relay-specific control word

The following relevant text is extracted from [section 2](#) (An Optimized Frame-Relay-specific control word) of [\[BRYANT\]](#) for reference:

"An optimized Frame-Relay-specific control word would require no bit-manipulation operators to transform to and from the Frame Relay header. Such a control word is presented here:

```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| length      | C|X|X|X|X|X|F|B|D|X|          Sequence Number          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Bryant internet-draft - expires May 2002 6
 [draft-bryant-pwe3-fr-compare-00.txt](#) October 2001

"The meaning of the fields is as follows:

"Length (bits 0 to 5):

The length field is used to indicate the payload length of a packet that may have been padded during transmission, if that information is not provided by the PSN encapsulation. If the length of the layer-two frame (consisting of layer-two control information and layer-two payload) is less than 64 bytes, the length field MUST be set to the Frame Relay PDU length. Otherwise the length field MUST be set to zero."

[4.1](#) Comments on Bryant Optimized Frame-Relay-specific control word

In this proposal the first two octets of the control word and the first two octets of a frame-relay header have the same payload specific bit layout, i.e. the CFB and D bits are located at the same positions in both the control word and header, allowing one to be

constructed from the other using a read, mask, write operation. The length field is placed in the same location at the unused high-order DLCI bits and the remaining bits in the first two octets are unused.

The length field is only used with PSN types that do not provide that information, and is therefore not required when an IP encapsulation is used.

5. General Pseudo-wire Payload Encapsulation

The following relevant text is extracted from [section 3](#) (General Pseudo-wire Payload Encapsulation) of [BRYANT] for reference:

"The most computationally-efficient encapsulation approach is the general pseudo-wire encapsulation approach proposed by [[MARTINI](#)] for HDLC, Ethernet and VLAN. The Frame Relay PDU is transported in its entirety, excluding only HDLC flags and the FCS. Bit stuffing is undone. The control word follows the general control word definition in [[MARTINI](#)]. The control word is OPTIONAL. If the control word is used then the flag bits in the control word are not used, and MUST be set to 0 when transmitting, and MUST be ignored upon receipt.

"If a fragmentation scheme is required within the pseudo-wire protocol layer, then this should be incorporated into the general control word for use by all encapsulation types. A general encapsulation is shown here:

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---------|-----|---|---|---|---|---|-----|--------|---|---|---|---|---|---|-----------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|--|--|--|--|
| | | | | | | | | | 1 | | | | | | | | | 2 | | | | | | | | | 3 | | | | | | | | |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | | | | |
| +-----+ | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | Res | | | | | | A B | Length | | | | | | | Sequence Number | | | | | | | | | | | | | | | | | | | | |
| +-----+ | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

| | | |
|--------|---|--------------|
| Bryant | internet-draft - expires May 2002 | 7 |
| | draft-bryant-pwe3-fr-compare-00.txt | October 2001 |

The meaning of the fields is as follows:

Res (bits 0 to 5):

Reserved bits. These are set to zero on transmission and ignored on reception.

A, B (bits 6 and 7):

0 0: Complete packet (start and end)

```
0 1: First segment (start and not end)
```

```
1 1: Continuing segment (not start and not end)
```

```
1 0: Last segment (not start and end)"
```

Use of the length and sequence number fields is similar to that of

the Bryant optimised Frame-relay-specific header discussed in [section 4](#) above, except that we adopt the 8-bit length and position preferred by [\[MARTINI\]](#).

5.1 Comments on General Pseudo-wire Payload Encapsulation

This approach treats Frame Relay as a general packet transport and transports and uses an encapsulation common to HDLC, VLAN and Ethernet. The [\[KAWA\]](#) fragmentation scheme is optionally included, with the changes to [\[KAWA\]](#) recommended in this document.

Note that this proposal makes no comment about the mapping between L2TP sessions (MPLS labels) and Frame Relay DLCIs. This approach, with the associated computational efficiencies demonstrated below, is valid both when there is a one-to-one relationship between DLCI and L2TP session (MPLS labels), and when multiple DLCIs share a common L2TP session (MPLS label).

6. Mapping between the Control Word and the Frame Relay Header

This section explores the computational complexity of the mapping between the proposed control words and the Frame Relay header.

6.1 Mapping between Martini control word and Frame Relay header

This section looks at the issue of mapping between the Martin control word and the Frame Relay header in software.

The control word defined in [\[MARTINI\]](#) is:

| | | | |
|---------------------------|---------------------|--------------------------|-----|
| 0 | 1 | 2 | 3 |
| 0 1 2 3 4 5 6 7 8 9 | 0 1 2 3 4 5 6 7 8 9 | 0 1 2 3 4 5 6 7 8 9 | 0 1 |
| +-----+-----+-----+-----+ | | | |
| Rsvd D B F C | | Length Sequence Number | |
| +-----+-----+-----+-----+ | | | |

The two-octet Frame Relay header in normal IETF (DoD) notation is:

Bryant internet-draft - expires May 2002 8
[draft-bryant-pwe3-fr-compare-00.txt](#) October 2001

| | |
|-------------------------------|-------------|
| 0 | 1 |
| 0 1 2 3 4 5 6 7 8 9 | 0 1 2 3 4 5 |
| +-----+-----+ | |
| hi dlci C 0 lo dlci F B D 1 | |
| +-----+-----+ | |

Bits 7 and 15 are the Q.922 EA (Extended Address) bits used for address field extension. In the unextended two-octet header these

are defined to be 0 and 1 respectively.

Processors rarely have efficient bit manipulation operations, so you cannot normally just assign the bits to their new positions easily. It therefore becomes necessary to isolate each of the DBFC bits in the control word using an AND operation, to use a shift operator to move each bit to its correct position in the Frame Relay header, and then to load the isolated bit into the header buffer using an OR operator. This requires four operations on each of four bits:

- Move first octet of control word to temporary location.
- AND to isolate bit.
- Shift bit to corresponding location in Frame Relay header octet.
- Write or OR bit into Frame Relay header.

A similar number of operations is needed in constructing the Control Word from the received Frame Relay header. For comparison we can regard this construction as having a cost of:

$$(4 \text{ bits} * 4 \text{ operations}) + (4 \text{ bits} * 4 \text{ operations}) = 32 \text{ operations.}$$

Note that we have not included the processing of the DLCI bits (normally an OR operation) in this analysis, since that is an overhead common to all transmission formats. Also note that, for the purposes of Frame Relay header manipulation, the EA bits can normally be included in the pro-forma DLCI definition.

6.2 Mapping between Kawa control word and Frame Relay header

This section looks at the issue of mapping between the [\[KAWA\]](#) control word and the Frame Relay header in software. The Kawa control word is:

[illegible]

Again for comparison, the two-octet Frame Relay header in normal IETF (DoD) notation is:

```

0                                     1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+---+---+---+---+---+---+---+
| hi dlc i      | C | 0 | lo dlc i | F | B | D | 1 |
+-+-+---+---+---+---+---+---+---+

```

In this case we can isolate the FBD bits as group, and knowing that they are in the same position and order in the octet in both Control

If two-octet operations are available to the processor, it is simply necessary to isolate the CFB and D bits in the control word through an AND operation and write them into the Frame Relay header.

- Move first word of control word to temporary location.
- AND to isolate CFBD bits.
- Write bits into Frame Relay header.

The operation to create the Control Word has similar steps, making a comparative total cost of 6 operations, rising to a comparative cost of 12 operations if only single-octet, rather than word, operators are available.

This is the fastest control word design for use with the [\[MARTINI\]](#) and [\[KAWA\]](#) approach of transmitting Frame Relay over pseudo-wire via an intermediate format.

[6.4](#) Mapping between the General Pseudo-wire Payload Encapsulation and Frame Relay

The general pseudo-wire encapsulation technique also discussed in [\[BRYANT\]](#) transports the Frame Relay header intact, and does not move any of the payload specific bits to the control word. This effectively has a Frame Relay header bit processing cost of zero.

[7.](#) Frame Relay header translation

A common justification for the use of an intermediate format for the transmission of a Frame Relay header is the need to translate the header from one format to another across the pseudo-wire.

There are two Frame Relay header formats in common usage: the two-octet version used in the examples in this draft, and the four-octet version.

If the intermediate format is used, the encapsulator has to implement two cases (two-octet to intermediate and four-octet to intermediate), and the decapsulator has to implement two cases (intermediate to two-octet, and intermediate to four-octet). If the frame is transmitted across the pseudo-wire un-translated, then there are three translation cases (Nothing, 2->4 and 4->2). This results in a net reduction in translation and implementation complexity, and an increase in performance.

It is useful to review the memory organization of the two-octet and four-octet Frame Relay headers, to fully appreciate the complexity of the translation operation. In normal IETF notation, the two-

octet frame relay header is (yet again):

Bryant internet-draft - expires May 2002 11
 [draft-bryant-pwe3-fr-compare-00.txt](#) October 2001

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+
| hi dlci   |C|0|lo dlci|F|B|D|1|
+---+---+---+---+---+---+---+---+---+---+
```

and for comparison the four-octet frame relay header is:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| hi dlci   |C 0| dlci  |F|B|D|0|   dlci   |0| dlci_lo  |0|1|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Bit 30 is the D/C bit. This is set to zero to indicate that the last octet contains DLCI, rather than control, information.

Note that the two least-significant (leftmost) octets in the header have an almost identical format, the only difference being that the EA bit in bit 15 is a zero to indicate a four-octet header.

Translating from a two-byte Frame Relay header to a four-byte frame relay header requires the following operations:

- Move first word of the header to temporary location.
- AND to clear bit 15 (EA = 0).
- Write word into Frame Relay header at new buffer offset.

Correct setting of the remaining EA bits (bits 23 and 31) and the DLCI Core indicator (D/C) bit (bit 30) can be considered an integral part of writing the least significant thirteen bits of the DLCI. Translation from a two-octet Frame Relay header to a four-octet Frame relay header has similar complexity:

- Move first word of the header to temporary location.
- OR to set bit 15 (EA = 0).
- Write word into Frame Relay header at new buffer offset.

The minimal additional computational cost of translation is therefore one additional branch plus one read plus one write, i.e. branch plus 3 operations. This cost is much lower than the cost Of translating to and from an intermediate format.

8. Conclusions

Based on this analysis the authors make the following recommendations to the IETF PWE3 Working Group:

Based on the reduced computational complexity, the simplicity of translation and the compatibility with the VLAN encapsulation type, the pseudo-wire Frame Relay service should transmit the frame as received rather than in an intermediate format, and perform any necessary translations as a local operation during the decapsulation processing. The general-purpose encapsulation presented in [[BRYANT](#)] is ideal for this, and should be used.

If the use of an intermediate Frame Relay format is objectively judged desirable, the optimised frame-relay-specific intermediate control word proposed by [[BRYANT](#)] is computationally more efficient than that proposed by [[KAWA](#)] and [[MARTINI](#)], and should therefore be adopted.

If the use of an intermediate Frame Relay format is objectively judged desirable, and there are also good reasons why the reserved bits are required and why the length field should remain unchanged, then the payload-dependent bit layout proposed in [[KAWA](#)] is computationally preferable to that proposed by [[MARTINI](#)], so [[KAWA](#)] is therefore preferred.

The naming and polarity of the fragmentation bits proposed by [[KAWA](#)] should be changed to the definitions proposed in this draft to reduce computational complexity. (We understand from recent list discussion that this is in progress.)

If a fragmentation mechanism is needed in PWE3, then it should be defined as a general method used by all encapsulation types. The general-purpose encapsulation presented in [[BRYANT](#)] is ideal for this, and should be used.

9. IANA considerations

There are no IANA considerations for this document.

10. Security considerations

There are no security implications for this document.

| | | |
|--------|---|--------------|
| Bryant | internet-draft - expires May 2002 | 13 |
| | draft-bryant-pwe3-fr-compare-00.txt | October 2001 |

11. References

Internet-drafts are works in progress available from
<http://www.ietf.org/internet-drafts/>

[BRYANT] S. Bryant and Wood, L. 'Two Efficient PWE3 Frame Relay Encapsulations', work in progress as an internet-draft:
<[draft-bryant-pwe3-fr-encap-00.txt](http://www.ietf.org/internet-drafts/draft-bryant-pwe3-fr-encap-00.txt)>

[GRE] D. Farinacci et al, 'Generic Routing Encapsulation (GRE)',
IETF [RFC2784](http://www.ietf.org/rfc/rfc2784.txt), standards-track.

[KAWA] C. Kawa, Malis, A., et al., 'Frame relay over Pseudo-Wire Emulation Edge-to-Edge', work in progress as an internet-draft:
<[draft-kamapabhava-fr-pwe3-00.txt](http://www.ietf.org/internet-drafts/draft-kamapabhava-fr-pwe3-00.txt)>

[MARTINI] L. Martini, Tappan, D., et al. 'Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks', work in progress as an internet-draft:
<[draft-martini-l2circuit-encap-mpls-03.txt](http://www.ietf.org/internet-drafts/draft-martini-l2circuit-encap-mpls-03.txt)>

[Q.922] ITU-T Recommendation Q.922, ISDN Data Link Layer Specification for Frame Mode Bearer Services, ITU, Geneva, 1992.

Authors' addresses

Stewart Bryant
Cisco Systems Ltd
4, The Square
Stockley Park
Uxbridge UB11 1BL
United Kingdom

Email: stbryant@cisco.com
Phone: +44-20-8756-8828

Lloyd Wood
Cisco Systems Ltd
4, The Square
Stockley Park
Uxbridge UB11 1BL
United Kingdom

Email: lwood@cisco.com
Phone: +44-20-8734-4236

Bryant internet-draft - expires May 2002
[draft-bryant-pwe3-fr-compare-00.txt](#)

14
October 2001

Full copyright statement

Copyright (C) The Internet Society (2001).
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.

