

PSEUDOWIRE LAYER NETWORK

Summary

This ITU-T Recommendation defines a layer network that facilitates transport of arbitrary client services over arbitrary server networks, including concatenations of several server networks of different technologies. It describes the encapsulation format.

Source

Editor

Keywords

client/server, interworking, multi-segment pseudowire, network interworking, pseudowire, user plane

<p>TSB Note: All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.</p>

CONTENTS

1	Scope	3
2	References	3
3	Definitions	4
3.1	Terms defined elsewhere	4
3.2	Terms defined in this Recommendation	4
4	Abbreviations and acronyms	4
5	Conventions	5
6	Multi-segment Pseudowires	5
6.1	Reference architecture for multi-segment transport using pseudowires	5
6.2	Advantages of the PW layer	7
6.3	Terminating and stitching IWFs	7
6.4	PW encapsulation format	7
7	General capabilities	8
7.1	User plane capabilities	8
7.2	Control and Management plane capabilities	8
8	The PW encapsulation format	9
8.1	Usage and compatibility	9
8.1.1	PW label	9
8.1.2	Control Word	9
8.1.3	Optional timing information	10
8.2	Server layer packet	11
9	Security considerations	11
Appendix I: Compatibility with existing encapsulations		12
I.1	Encapsulations compatible with this Recommendation	12
I.2	Encapsulations not compatible with this Recommendation	12
Appendix II: Guidance for use of PWs with specific server networks		13
II.1	MPLS	13
II.2	Ethernet	13
II.3	IP	13
II.4	PDH/SDH/OTN using GFP	14
Bibliography		15

Draft Recommendation Y.1418

PSEUDOWIRE LAYER NETWORK

1 Scope

This Recommendation describes a layer network that facilitates transport of arbitrary client services over arbitrary server networks. Example client services include ATM, frame relay, low-rate TDM, and Ethernet. Example server layers include MPLS, Ethernet, IP, and PDH/SDH/OTN using the generic framing procedure (GFP).

The layer network is called the pseudowire (PW) layer and is suitable for both single and multi-segment transport.

This Recommendation does not preclude the use of other adaptations, and in many cases alternative mechanisms are preferable.

This Recommendation is not intended to redefine the behaviour of the protocols specified in the normatively referenced IETF RFCs. In the event of discrepancies between protocols in this Recommendation and the normative RFCs, the IETF RFC is the authoritative source.

2 References

The following ITU-T Recommendations and other references contain provisions, which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; all users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published.

The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- | | |
|----------------|---|
| [ITU-T G.805] | ITU-T Recommendation G.805 (2000), <i>Generic functional architecture of transport networks</i> |
| [ITU-T X.84] | ITU-T Recommendation X.84 (2004), <i>Support of frame relay services over MPLS core networks</i> |
| [ITU-T Y.1401] | ITU-T Recommendation Y.1401 (2008), <i>Principales of interworking</i> |
| [ITU-T Y.1411] | ITU-T Recommendation Y.1411 (2003), <i>ATM-MPLS network interworking – Cell mode user plane interworking</i> |
| [ITU-T Y.1412] | ITU-T Recommendation Y.1412 (2003), <i>ATM-MPLS network interworking – Frame mode user plane interworking</i> |
| [ITU-T Y.1413] | ITU-T Recommendation Y.1413 (2004), <i>TDM-MPLS network interworking - User plane interworking</i> |
| [ITU-T Y.1414] | ITU-T Recommendation Y.1414 (2004), <i>Voice services - MPLS network interworking</i> |
| [ITU-T Y.1415] | ITU-T Recommendation Y.1415 (2005), <i>Ethernet-MPLS network interworking - User plane interworking</i> |

- [ITU-T Y.1452] ITU-T Recommendation Y.1452 (2006), *Voice trunking over IP networks*
- [ITU-T Y.1453] ITU-T Recommendation Y.1453 (2007), *TDM - IP Interworking – user plane interworking.*
- [IETF RFC 3550] IETF RFC 3550 (2004), *RTP: A Transport Protocol for Real-Time Applications*
- [IETF RFC 4385] IETF RFC 4385 (2006), *Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 Interworking: See [ITU-T Y.1411]

3.1.2 Characteristic Information (CI): See [ITU-T G.805]

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 Multi-segment transport: transport of a client service over more than one server network trail.

3.2.2 Pseudowire: a layer network used to emulate the behaviour of another layer network

3.2.3 Stitching IWF: The stitching IWF is a composite transport processing function that recovers the pseudowire characteristic information (PW CI) from a first server layer network updates the PW label and re-encapsulates the PW CI into characteristic information (CI) of a second server layer network. Note that the client of the PW is not recovered.

3.2.4 Terminating IWF: The terminating IWF is a composite transport processing function that encapsulates/decapsulates the client service characteristic information (CI) [ITU-T G.805] to/from the server layer CI.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms.

CI	Characteristic Information
FCS	Frame Check Sequence
IP	Internet Protocol
IWF	Interworking Function
LAN	Local Area Network
MPLS	Multi-protocol label switching
MTU	Maximum Transport Unit
OAM	Operation, Administration and Maintenance
OTN	Optical Transport Network
PW	Pseudowire

PDU	Protocol Data Unit
RTP	Real Time Protocol
QoS	Quality of Service
RFC	Request for Comments
TDM	Time Division Multiplex
VLAN	Virtual LAN

5 Conventions

For brevity, all functional diagrams in this Recommendation assume connection-oriented networks, and thus follow the diagrammatic conventions of [ITU-T G.805]. When either the client network, or server network, or both, are connectionless networks, similar diagrams apply, mutatis mutandis.

In addition, when discussing the encapsulation of a client protocol data unit (PDU) into a server packet we shall consistently mention the pre-pending of server-layer headers. It is recognized that certain server layer encapsulations additionally involve appending of trailers.

6 Multi-segment Pseudowires

6.1 Reference architecture for multi-segment transport using pseudowires

Initial work on client-server interworking focused on client information being carried over a single-server network. However there are many scenarios where it is desirable to traverse multiple concatenated server networks, and in order to facilitate this interworking, the pseudowire (PW) layer network is introduced. Figure 6-1 depicts the general topology for multi-segment pseudowires, depicting the native service, transport segments, the multi-segment path, the terminating and stitching IWFs. The native service is the client of the PW layer network.

Note that pseudowire is an implementation of common adaptation layer (CAL) as described in [ITU-T Y.1401]

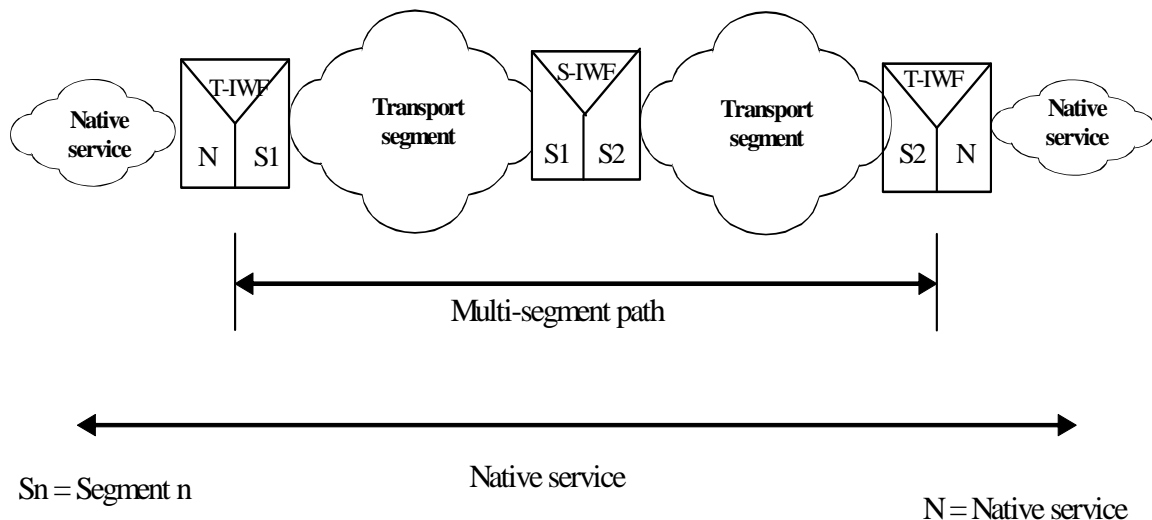


Figure 6-1: Multi-segment interworking topology

In this Recommendation, for the multi-segment case, the PW label may be modified by the interworking function. This is in contrast to the single segment cases.

The generic adaptation specified by the PW layer network is used for transport over a single transport segment when there is no well established adaptation, and whenever multi-segment transport is required. Use of the PW layer is not universally recommended, as many client and server networks have well established efficient adaptations.

Multi-segment transport is the transport of a native service over the concatenation of multiple server networks. One case where multi-segment transport must be employed involves native service end-points connected to server networks based on differing technologies. Present interworking Recommendations are specific to the server network technology, and thus the client information can not be carried across such concatenated server networks without client-specific processing between each pair of server networks. Figure 6-2 is a functional diagram, following the diagrammatic conventions of [ITU-T G.805], for a case where it is required to traverse three server networks of different technologies.




Figure 6-2: Functional diagram for multi-segment interworking for the multi-technology case

A second case where multi-segment transport is employed is when native service end-points connect to server networks in different administrative domains. In this case the terminating IWFs will in general not have mutual visibility as internal topologies may not be advertised outside their respective administrative domains. Figure 6-3 is a functional diagram showing the traversal of two server networks belonging to different administrative domains.




Figure 6-3: Functional diagram for multi-segment interworking for the multi-domain case

Multi-segment transport may also be used even when there is single-technology reachability and both terminal IWFs are in a single administrative domain. For example, in Appendix I of [ITU-T Y.1415] a full mesh of interworking LSPs is required to provide a multi-point Ethernet service. If the full mesh is prohibitive from a resource point of view, then a partial mesh may be used if multi-segment transport is supported.

6.2 Advantages of the PW layer

The introduction of an intermediate layer network enables the use of stitching IWFs that operate without knowledge of the native service type, and do not need specific processing functionality to handle different types.

For real-time native services, such as TDM, use of the intermediate layer avoids exposure of the native service between server networks, which would introduce additional packetization delay and clock degradation.

Introduction of an intermediate layer enables use of PW layer OAM. The PW encapsulation described below features an associated channel, which can be used for this purpose.

6.3 Terminating and stitching IWFs

The ingress terminating IWF performs two separate encapsulations. The native service information is first encapsulated along with any needed overhead to form the PW characteristic information (CI). The PW CI is further encapsulated to create the server-specific CI. Similarly, the egress terminating IWF performs two separate decapsulations. The PW CI is first recovered from the server layer CI. The native service information is then recovered from the PW CI.

On the other hand, a stitching IWF recovers the PW CI from a first server layer network, updates the PW label and re-encapsulates the PW CI into CI of the second server layer network. Note that the stitching IWF does not recover native CI, and is agnostic to the native service type.

In addition to their user-plane functionality, stitching IWFs may participate in control and management plane functions. Description of these functions is beyond the scope of this Recommendation.

Note that any additional functionality required for adapting the native service to the generic adaptation layer is only performed by a terminating IWF, and not by stitching IWFs. For example, if the native service is TDM, only an ingress terminating IWF is responsible for segmentation, and only an egress terminating function for clock recovery.

6.4 PW encapsulation format

The PW format is based on that used in [ITU-T Y.1411] through [ITU-T Y.1415] and [ITU-T X.84]. For all server types, the server network is responsible for the transport of the PW PDU from ingress to egress. The format of the PW PDU is given schematically in Figure 6-4. Figure 6-5 shows the packet encapsulations at the different points of the functional diagram for the case of Figure 6-2.

PW label
Control Word (CW)
Optional Timing Information
Native Service CI

Figure 6-4: Format of PW packet format

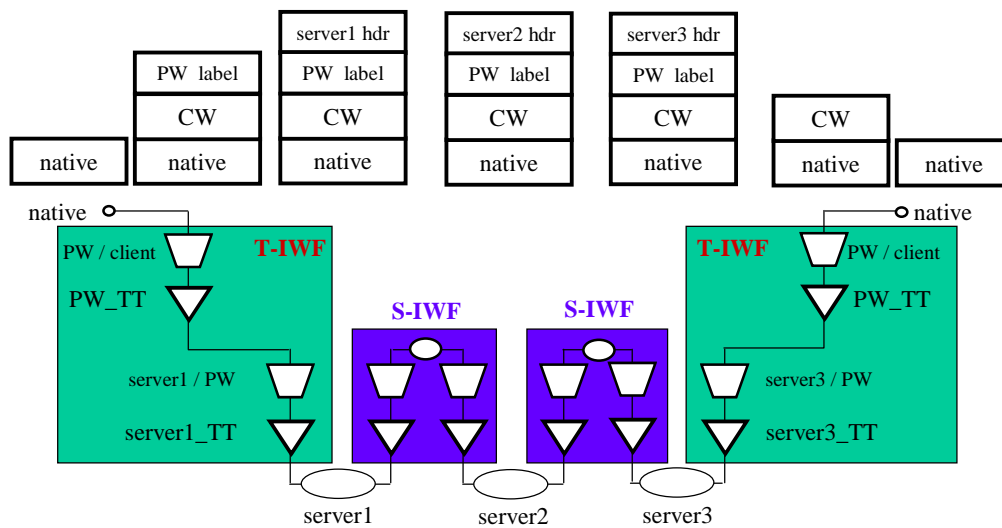


Figure 6-5: Encapsulations at different points of the functional diagram

7 General capabilities

7.1 User plane capabilities

The following user plane capabilities are required:

- The ability to transparently transport any native service across any server layer network.
- The ability to generate and monitor OAM from terminating IWF to terminating IWF, irrespective of server network technology.
- The ability to maintain clock performance if needed by the native service.

The following additional capabilities are only required for multi-segment transport:

- The ability to transparently transport native service CI across a concatenation of server networks, irrespective of server network technology.
- The ability to stitch transport segments independent of native service technology.
- The ability to flexibly interconnect between transport segments.

7.2 Control and Management plane capabilities

The following control plane capabilities are required:

- The ability to setup transport segments by manual provisioning.
- The optional ability to setup transport segments by signalling.

The following additional capabilities are not required for the single segment but are optional only for multi-segment transport:

- The ability of stitching IWFs to maintain control signalling with both server networks.
- The ability to setup the entire multi-segment path by signalling.

8 The PW encapsulation format

8.1 Usage and compatibility

The PW layer enables transport of an arbitrary native service over an arbitrary packet-based server layer network. When there is no well-established adaptation mechanism, and/or when multi-segment transport is required, the PW encapsulation format shall be used to carry native service CI across a packet-based server network.

The PW encapsulation format is compatible with existing encapsulation formats enumerated in Appendix I.1.

A non-exhaustive list of encapsulation formats not compatible with PWs is given in Appendix I.2.

The PW encapsulation is performed by first creating a PW PDU from the native service CI, as depicted in Figure 8-1. Thereafter a server layer packet is created from the PW PDU, as described in sub-clause 8.2, below.




Figure 8-1: PW PDU

8.1.1 PW label

The PW label consists of 4 octets and contains information necessary to distinguish the particular native service instantiation from all others, and must be unique for a given transport segment. Its format is given in Figure 8-2.

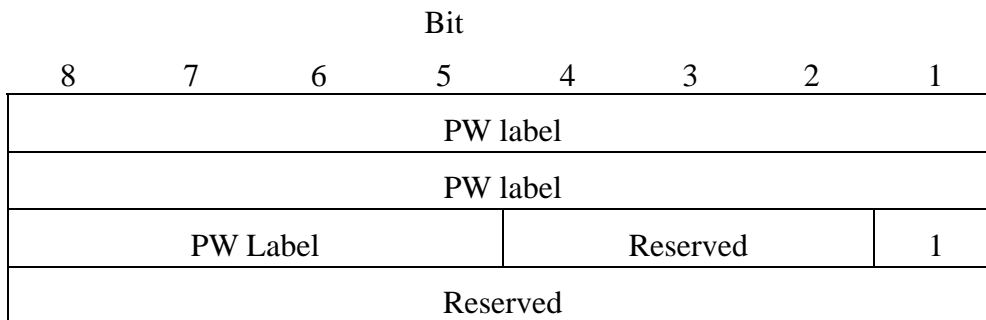


Figure 8-2: PW label format

The PW label consists of 20 bits and must be unique for the transport segment in question.

The two reserved fields may be used by specific server networks for needed functionality such as marking class of service and loop avoidance.

8.1.2 Control Word

The control word consists of 4 octets as depicted in Figure 8-3. Note that this control word is fully compatible with the control word specified in [IETF RFC 4385] and with the common interworking indicators of Recommendations [ITU-T Y.1411]-[ITU-T Y.1415], [ITU-T Y.1452] and [ITU-T Y.1453].

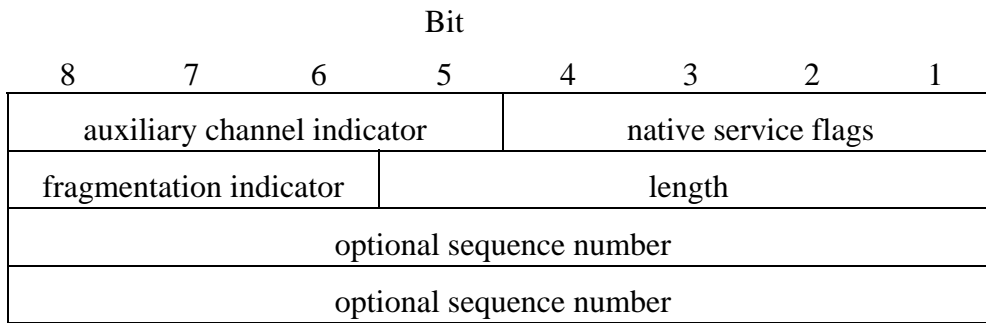


Figure 8-3: Control word (CW) format

The auxiliary channel indicator consists of 4 bits and is used to differentiate between native service payloads and other payload types. Two values are presently defined:

0000: native service payload

0001: auxiliary channel.

The native service flags consist of 4 bits whose meanings are specific to the native service technology.

The use of the rest of the control word is specific to either the native service, the server or both.

The two-bit fragmentation indicator is used for native services that require marking of fragments (see [ITU-T Y.1413]). When not used it is set to zero.

The six-bit length is used for server networks with minimum transmission unit sizes (e.g. Ethernet), and it enables the PW adaptation function to correctly recover the native service CI when padding has been employed. Its use is described in [ITU-T Y.1413].

The sequence number is used for server networks that may misorder packets and native services that are sensitive to packet misordering yet have no inherent reordering mechanism.

8.1.3 Optional timing information

The optional timing information, when required, is provided by using the fixed RTP header described in [IETF RFC3550]. Guidance as to its use can be found in [ITU-T Y.1413].

8.2 Server layer packet

The server layer packet is created from the PW-PDU by pre-pending headers, as schematically depicted in Figure 8-4.

Server Layer Header(s)
PW Label
Control Word (CW)
Optional RTP header
Native Service PDU

Figure 8-4: Server network PDU

The server layer headers contain all information needed to enable forwarding of the packet from server layer ingress to egress, as well as to maintain required QoS. Specific examples of server layer headers are given in Appendix II.

9 Security considerations

Security aspects have not been addressed in this Recommendation.

Appendix I: Compatibility with existing encapsulations

(This appendix does not form an integral part of this Recommendation)

I.1 Encapsulations compatible with this Recommendation

The following adaptations have been identified as being compatible with the PW adaptation defined in this Recommendation:

- ATM over MPLS, per [ITU-T Y.1411]
- AAL5 over MPLS, per [ITU-T Y.1412]
- TDM over MPLS, per [ITU-T Y.1413]
- Voice services over MPLS, per [ITU-T Y.1414]
- Ethernet over MPLS, per [ITU-T Y.1415]
- Frame relay over MPLS, per [ITU-T X.84]
- Ethernet over MPLS, per [b_IETF RFC 4448]
- Structure agnostic transport of TDM over MPLS, per [b_IETF RFC4553]
- Structure aware transport of TDM over MPLS, per [b_IETF RFC 5086] and [b_IETF RFC 5087]
- PPP/HDLC over MPLS, per [b_IETF RFC 4618]
- Frame relay over MPLS, per [b_IETF RFC 4619]
- ATM over MPLS, per [b_IETF RFC4717]
- ATM cell transport over MPLS, per [b_IETF RFC 4816]
- TDM over MPLS using AAL1, per [b_MFA4.1]
- TDM voice services over MPLS, per [b_MFA5.1]
- TDM over Ethernet, per [b_MEF8]

I.2 Encapsulations not compatible with this Recommendation

This Recommendation does not preclude the use of other adaptations for specific cases, and in many cases well-established mechanisms are preferable.

The following is a non-exhaustive list of such cases:

- IP over Ethernet, per [b_IETF RFC 894]
- Frame-relay over Ethernet, per [b_IETF RFC 1701]
- Ethernet over Ethernet, per [802.1ah]
- IP over MPLS, per [b_IETF RFC 3032]
- MPLS over MPLS via label stacking [b_IETF RFC 3032]
- TDM over MPLS, per [ITU-T Y.1413] or over IP per [ITU-T Y.1453]
- Voice over IP, per [IETF RFC 3550]

Appendix II: Guidance for use of PWs with specific server networks

(This Appendix does not form an integral part of this Recommendation)

II.1 MPLS

Interworking for this server network is discussed in [ITU-T Y.1411] - [ITU-T Y.1415], [ITU-T X.84], [b_IETF RFC 4448], [b_IETF RFC 4553], [b_IETF RFC 4618], [b_IETF RFC 4619], [b_IETF RFC4717], [b_IETF RFC4816], [b_IETF RFC5086], [b_IETF FC5087], [b_MFA4.1] and [b_MFA5.1].

In order to transport native service CI across an MPLS network, we need only pre-pend the MPLS label stack [b_IETF RFC 3032] to the PW PDU, see Figure II-1. The PW label in the PW PDU becomes the bottom label of the MPLS stack. Note that as such, it has its “Bottom of Stack” indicator set to one.




Figure II-1: Encapsulation format for PW over MPLS

Note that when the native service is MPLS, label stacking as specified in [b_IETF RFC 3032] is always used. That adaptation can be distinguished from the generic adaptation by observing the “Bottom of Stack” indicator. When the native service is IPv4 or IPv6, the adaptation specified in [b_IETF RFC 3032] is used. That adaptation can be distinguished by the values of 0100 and 0110 in the control word switch field, respectively.

II.2 Ethernet

Interworking for this server network is discussed in MEF8.

In order to transport a native service CI across an Ethernet network, an Ethernet header and FCS trailer are inserted before and after the PW PDU, respectively, see Figure II-2. The source address (SA) and destination address (DA) are set to the Ethernet MAC addresses of the encapsulating and decapsulating IWFs, respectively. Virtual LAN (VLAN) tagging may be used and padding shall be applied if needed. The Ethertype should be set to 8847 (hex), the value assigned for unicast MPLS. The Ethertype may be set to a value specific to the native service, when such an Ethertype is defined, e.g. 88D8 (hex) for TDM over Ethernet.




Figure II-2: Encapsulation format for PW over Ethernet

Note that when the native service is Ethernet, the “MAC-in-MAC” adaptation specified in [b_IEEE 802.1ah] is preferred. That adaptation can be distinguished by Ethertype of 88A8 (hex) or 8100 (hex).

II.3 IP

In order to transport native service CI across an IP network, we need only pre-pend the IPv4 [b_IETF RFC 791] or IPv6 header [b_IETF RFC 2460] to the PW PDU, as is done for transport of MPLS over IP [b_IETF RFC 4023], see Figure III-3. The source and destination addresses are set to

addresses of the encapsulating and decapsulating IWFs, respectively. As the PW label is of the same format as the bottom label on an MPLS stack, the IPv4 Protocol Number field or the IPv6 Next Header field is set to 137, indicating an MPLS unicast packet.




Figure II-3: Encapsulation format for PW over IP

Note that this adaptation is different from the adaptations specified in Y.1452, Y.1453 (which lack the PW label) and the adaptation defined in [IETF RFC 3550] (which lacks the PW label and control word). As these adaptation use UDP, they can be distinguished by the IP protocol field being set to 11 (hex). Furthermore, this adaptation is different from those defined in [b_IETF RFC 4349], [b_IETF RFC 4454], [b_IETF RFC 4591] and [b_IETF RFC 4719], which may run without UDP. Those adaptations can be distinguished from the generic adaptation by its having the IP protocol field set to 73 (hex).

II.4 PDH/SDH/OTN using GFP

GFP provides a packet interface to TDM networks. In order to transport native service CI over a TDM network, the PW-PDU is encapsulated using GFP-F [b_ITU-T G.7041/Y.1303], see Figure II-4. The octets of the PW-PDU are placed octet-aligned into the GFP payload information field. The user payload identifier (UPI) is set to that that defined for unicast MPLS. As PW-PDUs do not have error control, the GFP FCS should be used.




Figure II-4: Encapsulation format for PW over GFP

Note that there are additional adaptations, such as those defined in [b_ITU-T X.85], [b_ITU-T X.86], [b_IETF RFC 1619], [b_IETF RFC 2615], and Clause 50 of [802.3-2005]. The generic adaptation can be distinguished by its having the path signal label in the path overhead indicate GFP, e.g., the SDH C2 byte must equal to 1B (hex)

Bibliography

- [b_ITU-T G.7041] ITU-T Recommendation G.7041/Y.1303 (2005), *Generic framing procedure (GFP)*
- [b_ITU-T X.85] ITU-T Recommendation X.85/Y.1321 (2001), *P over SDH using LAPS*
- [b_ITU-T X.86] ITU-T Recommendation X.86/Y.1323 (2001), *Ethernet over LAPS*
- [b_IETF RFC 791] IETF RFC 791 (1981), *Internet Protocol*
- [b_IETF RFC 1619] IETF RFC 1619 (1994), *PPP over SONET/SDH*
- [b_IETF RFC 2460] IETF RFC 2460 (1998), *Internet Protocol, Version 6 (IPv6) specification*
- [b_IETF RFC 2615] IETF RFC 2615 (1999), *PPP over SONET/SDH*
- [b_IETF RFC 3032] IETF RFC 3032 (2001), *MPLS label stack encoding*
- [b_IETF RFC 4023] IETF RFC 4023 (2005), *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*
- [b_IETF RFC 4349] IETF RFC 4349 (2006), *High-Level Data Link Control (HDLC) Frame over Layer 2 Tunneling Protocol, Version 3 (L2TPv3)*
- [b_IETF RFC 4448] IETF RFC 4448 (2006), *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
- [b_IETF RFC 4454] IETF RFC 4454 (2006), *Asynchronous Transfer Mode (ATM) over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*
- [b_IETF RFC 4553] IETF RFC 4553 (2006), *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*
- [b_IETF RFC 4591] IETF RFC 4591 (2006), *Frame Relay over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*
- [b_IETF RFC 4618] IETF RFC 4618 (2006), *Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks*
- [b_IETF RFC 4619] IETF RFC 4619 (2006), *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*
- [b_IETF RFC 4717] IETF RFC 4717 (2006), *Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks*
- [b_IETF RFC 4719] IETF RFC 4719 (2006), *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*
- [b_IETF RFC 4816] IETF RFC 4816 (2007), *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*
- [b_IETF RFC 5086] IETF RFC 5086 (2007), *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*
- [b_IETF RFC 5087] IETF RFC 5087 (2007), *Time Division Multiplexing over IP (TDMoIP)*
- [b_MEF 8] MEF 8 (2004), *Implementation agreement for the emulation of PDH circuits over Metro Ethernet network.*
- [b_MFA4.1] Implementation Agreement, MPLS/Frame Relay Alliance 4.0, MPLS/Frame Relay Alliance Technical Committee, (2003), *TDM Transport over MPLS using AAL1 Implementation Agreement*
- [b_MFA5.1] Implementation Agreement, MPLS /Frame Relay Alliance 5.0.0 MPLS /Frame Relay Alliance Technical Committee (2003), *I.366.2 Voice Trunking Format over MPLS Implementation Agreement*
-