TICTOC Working Group
Internet Draft
Intended status: Informational
Expires: September 2012

S. Jobert
France Telecom Orange
March 5, 2012

Transporting PTP messages over MPLS networks using a link local addressing
draft-jobert-tictoc-ptp-link-local-00.txt

Status of this Memo

Abstract

This document introduces a method for transporting PTP messages over an MPLS network supported by an Ethernet physical layer. The MPLS layer itself is not used to carry the PTP messages with this method; instead, a link local Ethernet channel is used. Several advantages related to this method are highlighted in this document. The method targets in particular telecom applications requiring accurate phase/time synchronization, with "link-by-link" PTP architectures, where all the network nodes support a PTP function, such as Boundary Clock or Transparent Clock.

Table of Contents

1. Introduction

The Precision Time Protocol version 2 (PTPv2), defined by the [IEEE1588-2008] standard, is used to support telecom applications that may include MPLS networks. Telecoms applications may require frequency synchronization only or accurate phase/time synchronization.

This has led to the definition of two PTP telecom profiles at the ITU-T: the Recommendation [G.8265.1] (finalized) defines a PTP telecom profile for frequency synchronization in an “end-to-end” mode (the intermediate network nodes do not support PTP functions) and the future Recommendation G.8275.1 (under development) will define a PTP telecom profile for phase/time synchronization in a “link-by-link” mode (all the intermediate network nodes support PTP functions).

For frequency applications using the ITU-T frequency profile, there is no particular need to identify the PTP messages in case they are carried in an MPLS layer. The use of a high priority class of service is in general sufficient to minimize the Packet Delay Variation (PDV) introduced by the network nodes. The identification of the PTP messages in a network node which does not support PTP functions is not expected in general to provide a better performance than the positioning of the PTP messages in a dedicated high priority queue.

For phase/time applications with stringent requirements (e.g. sub-micro-second accuracy), it is in general recognized that PTP support from the network nodes is required to avoid the generation of Packet Delay Variation. Therefore, being able to identify the PTP messages is considered important. This is the one of the objectives of the definition of a PTP mapping. Some mappings are already defined in the [IEEE1588-2008] standard, and may be applicable to an MPLS network.

This document introduces a method for transporting PTP messages over an MPLS network supported by an Ethernet physical layer. The MPLS layer itself is not used to carry the PTP messages with this method; instead, a link local Ethernet channel is used.

Several advantages related to this method are highlighted in this document. The method targets in particular telecom applications requiring accurate phase/time synchronization, with “link-by-link” PTP architectures, where all the network nodes support a PTP function, such as Boundary Clock (BC) or Transparent Clock (TC).

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

PTP: Precision Time Protocol

PDV: Packet Delay Variation

BC: Boundary Clock

TC: Transparent Clock

3. Analysis of the PTP frequency telecom profile with MPLS networks

For applications requiring frequency synchronization only, when the use of physical layer synchronization methods such as Synchronous Ethernet is not possible, the ITU-T PTP frequency telecom profile defined in the Recommendation G.8265.1 is in general relevant, especially in order to address mobile networks needs.

This PTP telecom profile is based on an “end-to-end” PTP architecture: the intermediate network nodes do not support PTP functions such as Boundary Clock (BC) or Transparent Clock (TC). As such, they generate Packet Delay Variation (PDV). The PTP communication is only performed between a PTP master function and a PTP slave function.

This PTP dialog may involve different layers, due to different encapsulations. In particular, it is common that PTP messages are carried within an MPLS layer when using this PTP profile.

In order to minimize the PDV generated by the intermediate network nodes, PTP messages MUST be marked as high priority traffic, and MUST be positioned in high priority queues. This marking does not involve new PTP functions in the network nodes; it corresponds simply to the usual DiffServ functions supported in these devices.

In particular, the intermediate network nodes do not identify the PTP messages among the rest of the traffic; only the marking of the packets is considered to position them in the relevant queues.

The identification of the PTP messages by an intermediate network node which does not support PTP functions with this PTP frequency telecom profile is not expected in general to provide real performance improvements compared to the prioritization of the PTP traffic and the positioning of the PTP messages in a dedicated high priority queue.

Indeed, more specialized treatment of the PTP messages would make the network node very close to a node supporting PTP functions such as Boundary Clocks or Transparent Clocks. This would be quite contradictory to the architecture assumptions of this PTP frequency telecom profile.

In conclusion, when the ITU-T PTP frequency telecom profile defined in the Recommendation G.8265.1 is used, the identification of the PTP messages among the rest of the MPLS traffic does neither appear necessary, nor providing real performance benefits.

4. Transporting PTP messages over MPLS networks with a “link-by-link” PTP architecture

For applications requiring accurate phase/time synchronization, the use of the future ITU-T PTP phase/time telecom profile under definition in the Recommendation G.8275.1 is foreseen to be relevant to address the needs of mobile networks.

This PTP telecom profile is based on a “link-by-link” PTP architecture: the intermediate network nodes MUST support PTP functions such as Boundary Clock or Transparent Clock. This architecture is considered as necessary to avoid the generation of Packet Delay Variation, due to the stringent accuracy requirements that are targeted. The PTP communication is therefore performed between different PTP entities: PTP master function, PTP slave function, PTP Boundary Clocks, PTP Transparent Clocks.

Hence, being able to identify the PTP messages is considered important, in order to allow the intermediate network nodes to apply the special treatment on the PTP packets corresponding to the PTP function that they implement (BC or TC).

This is one of the objectives of the definition of a PTP mapping. Some mappings are already defined in the [IEEE1588-2008] standard, and may be applicable to an MPLS network. The transport of PTP messages over MPLS networks SHOULD NOT involve the MPLS layer itself in this type of "link-by-link" PTP architecture.

4.1. Need for identifying the PTP messages in MPLS networks

The "link-by-link" PTP architecture described above may be applicable over MPLS networks. As such, it is relevant to discuss the mapping options for transporting the PTP messages over MPLS networks when considering this type of PTP architecture.

Two PTP operations may be necessary in the MPLS nodes in order to handle the PTP packets in the general case:

- PTP packets detection: how to detect that a packet contains PTP payload? (this question is applicable to both Boundary Clock or Transparent Clock types of PTP support)

- PTP payload position in the packet: how to determine where the PTP payload is in the message once the relevant packets have been detected? (this question is applicable only to Transparent Clock PTP support, because Boundary Clocks terminate and process the PTP payload)

Regarding the first point listed above (PTP packets detection), the three following mappings could be considered in the general case:

- in case of an Ethernet mapping, the PTP packets can be detected thanks to a specific Ethertype. Some PTP mappings already defined in [IEEE1588-2008] already cover this point (see Annex F).

- in case of an IP/UDP mapping, the PTP packets can be detected thanks to specific UDP port numbers. Some PTP mapping already defined in [IEEE1588-2008] already cover this point (see Annexes D and E). This mapping corresponds to the mapping specified for the PTP frequency telecom profile defined in [G.8265.1].

- in case of MPLS mapping, if relevant, the draft [4] (“Transporting PTP messages (1588) over MPLS Networks”) currently discussed in the IETF TICTOC Working Group aims at specifying new MPLS mappings enabling to detect the PTP packets among the traffic. Note that these new PTP mappings are not defined in [IEEE1588-2008].

This document advocates that the third type of mapping (MPLS mappings) is not necessary for carrying PTP messages over MPLS networks supported by an Ethernet physical layer when using a “link-by-link” PTP architecture as depicted above in this document. Instead, it is considered that the use of a link local addressing is more relevant when the MPLS network is supported by an Ethernet physical layer. This point will be discussed further in the next sections of this document.

Regarding the second point (PTP payload position in the packet), it should be stressed the network nodes may not know exactly where the PTP payload is in the packet in some cases (e.g. when tunnels are used), because of other potential encapsulations beyond the layer handled by the node. This situation may happen in the case of MPLS network nodes. In particular, as mentioned above, it raises problems for modifying the PTP payload in case of a Transparent Clock PTP support.

This document explains that the use of a link local addressing simplifies this point, since the PTP payload is in this case at a fixed location in the message. It is moreover in line the with the principles of a “link-by-link” PTP architecture, where the PTP messages are sent to the next network node, and are not assumed to be forwarded through a tunnel. This point will be discussed further in the next sections of this document.

4.2. Use of a link local addressing over MPLS networks supported by an Ethernet physical layer

This section introduces a solution to carry PTP messages over an MPLS network supported by an Ethernet physical layer, using a link local Ethernet addressing. This solution fits very well with the “link-by-link” PTP architecture depicted before.

With this solution, Ethernet interfaces supporting MPLS traffic MUST use the Ethernet multicast address: ‘01-80-C2-00-00-0E’ based on the Annex F of IEEE1588-2008 for all the PTP messages that are sent.

This type of addressing aims at making sure that the PTP messages will be sent to the next network node in the chain (which may be or not an MPLS node).

This solution has several advantages:

- It prevents unwanted forwarding of PTP messages over network nodes which do not provide PTP support: indeed, such a network node is assumed in general to drop the PTP messages, and not to forward them. It is useful in order to avoid the generation of PDV. This property is considered in line with the "link-by-link" PTP architecture principles depicted earlier.

- It facilitates the configuration for the operator, since no particular addressing needs to be configured in the network nodes.

- It allows having a consistent PTP mapping all along the chain: all the PTP messages are transported the same way, using the same mapping, whatever the actual layers used to transport the user plane. In particular, an MPLS node may establish a PTP dialog with an IP node or a node working at the layer 2 with this type of solution.

- It facilitates the PTP payload identification, since the PTP payload is necessarily at a fixed location.

Note: in case of MPLS nodes connected together via a different physical layer than Ethernet, another link local channel linked to the physical layer might be used. This is beyond the scope of this document.

4.3. Use of link local addressing with Transparent Clocks

The case of Transparent Clock type of PTP support deserves a specific analysis when considering the use of a link local addressing. Indeed, some designs of Transparent Clock may not terminate the PTP messages; it creates issues in order to forward the PTP messages when link local addressing is used.

This section highlights however that some simple mechanisms might be implemented in Transparent Clocks to ensure their compatibility with the use of a link local addressing as proposed in the previous

section. It also shows that a link local addressing may avoid the layer violation issues with TCs.

Three main steps are observed in a standard Transparent Clock which does not terminate the PTP messages in order to treat and forward them:

1- Detection of the PTP packet among the rest of the traffic on an active PTP port, and precise timestamping of the arrival instant of the packet in the network node.

2- The PTP packet is treated/forwarded in the network node as a standard packet, e.g. analysis of the network header of the packet corresponding to the layer treated by the network node, in order to determine using the forwarding engine towards which output port the packet must be forwarded (for instance: IP lookup operation in a routing table). In summary: the output port is determined based on information contained in the PTP packet itself, using standard forwarding functions in the network node.

3- Transmission of the PTP packet at the output of the network node on the port determined before, and precise timestamping of the emission instant of the packet in the network. Modification of the "correction field" of the packet to include the residence time calculation.

The layer violation is due here to the fact that the PTP packet has been modified (correction field update) by an intermediate node which was assumed only to forward it. Moreover, there might be some difficulties to determine where the PTP payload is located, as mentioned earlier.

The use of a link local addressing might not be suitable with this model of TC. Indeed, it can be observed that the step 2 requires in the general case that the necessary information (e.g. final destination address) would be contained in the network header of the PTP messages to determine the output port where each PTP message must be forwarded. This is not the case with link local addressing, because each message is sent to the next node over a single link.

However, there are easy ways to overcome this issue. One possible straightforward solution could be to include locally in the network node the necessary information for the forwarding of the PTP messages. This might correspond to a "PTP local forwarding

function", which could be part of the network node configuration (manual configuration would be possible, but automatic procedures would also work).

As for the case of a standard TC, three main steps are observed in order to treat and forward a PTP message in a Transparent Clock implementing a PTP local forwarding function:

- The step 1 is similar in both cases (standard TC and TC with PTP local forwarding function).
- The step 2 would differ in this example (TC with PTP local forwarding function): the standard forwarding function of the network node (forwarding engine) MUST NOT be used in this case to forward the PTP packets; instead, the PTP local forwarding function MUST be used. This allows handling PTP packets without forwarding information in the network header of the packet.
- The step 3 is quite similar in both cases (standard TC and TC with PTP local forwarding function).

It must be stressed that the use of link local addressing leads to terminate the PTP packets that are received by the network node, since the recipient of the PTP messages is the network node itself. The PTP packets sent at the output of the TC with PTP local forwarding function are therefore new PTP packets, similarly to a BC. This is the reason why it can be considered as a way to avoid the layer violation issue.

In practice, the operations are similar between standard TC and TC with PTP local forwarding function for generating a new PTP packet based on the PTP packet received (e.g. update of the correction field, etc...).

Moreover, it must also be stressed that the use of link local addressing leads to a fixed location of the PTP payload in the packet. This is expected to greatly simplify the operations.

The PTP local forwarding function includes locally in the network node all the necessary information for forwarding the PTP packets. For instance, it may associate one or several output ports to an input port. An example of what could be a PTP local forwarding function is provided in the figure 1 below.

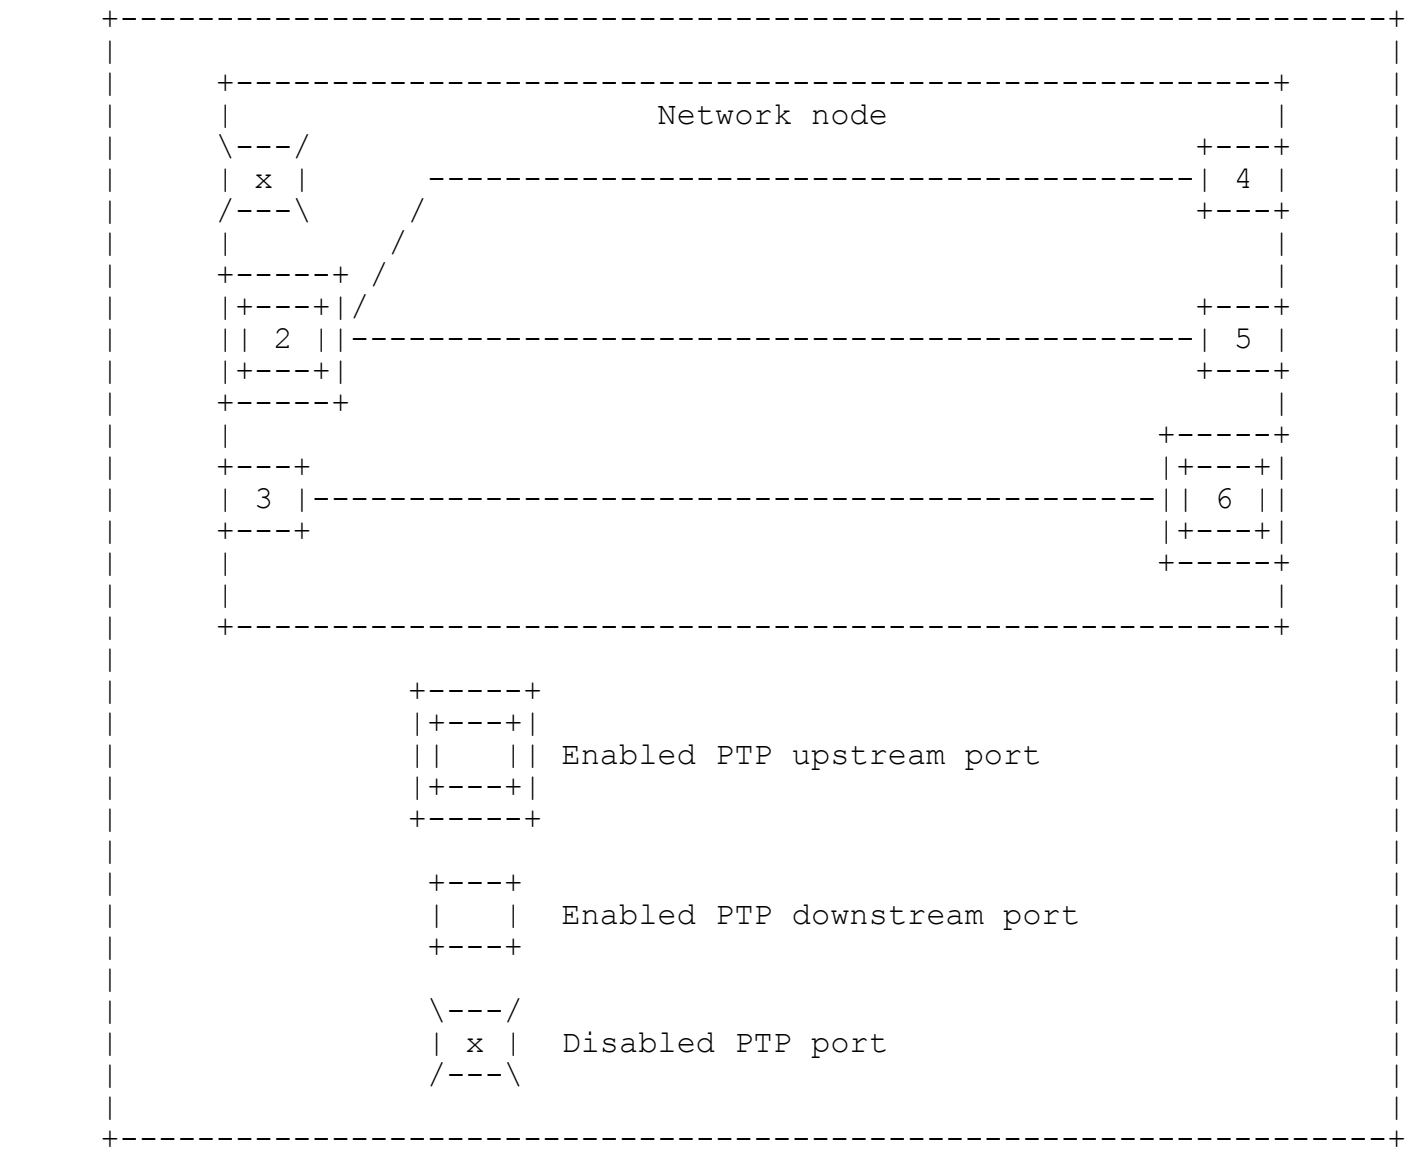

Figure 1 – Example of a possible configuration of the PTP local forwarding function

In the figure 1 above, three configurations are possible for a PTP port in a TC with PTP local forwarding function:

- Disabled PTP port: any potential PTP packet received on this port MUST be discarded.

- Enabled PTP upstream port: corresponds to a port where upstream PTP packets are received (e.g. the PTP packets generated by a PTP master port). When a PTP packet is received on an enabled PTP upstream port, a new PTP packet MUST be transmitted by one or several enabled PTP downstream ports of the network node associated to the enabled PTP upstream port. This/these new PTP packet(s) is/are formed using the information of the original PTP packet that was received, and by modifying the fields normally modified by a TC (the correction field in particular).

- Enabled PTP downstream port: corresponds to a port where downstream PTP packets are received (e.g. the PTP packets generated by a PTP slave port). When a PTP packet is received on an enabled PTP downstream port, a new PTP packet MUST be transmitted by the enabled PTP upstream port of the network node associated to the enabled PTP downstream port. This new PTP packet is formed using the information of the original PTP packet that was received, and by modifying the fields normally modified by a TC (the correction field in particular).

Note that the case of a two-port device is an example where implicit PTP local forwarding function exists: every port PTP packet received on one port must be forwarded by the other port.

The advantages of this type of mechanism are that it allows mixing BCs and TCs in a chain in a consistent way, using link local addressing. It also allows avoiding layer violation issues, since the PTP messages are terminated and processed by each network node, including the TC with PTP local forwarding function.

5. Security Considerations

<Add any security considerations>

6. IANA Considerations

<Add any IANA considerations>

7. References

7.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Crocker, D. and Overell, P.(Editors), "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, Internet Mail Consortium and Demon Internet Ltd., November 1997.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2234] Crocker, D. and Overell, P.(Editors), "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, Internet Mail Consortium and Demon Internet Ltd., November 1997.

[IEEE1588-2008] IEEE 1588-2008, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", July 2008.

[G.8265.1] ITU-T Recommendation G.8265.1 "Precision time protocol telecom profile for frequency synchronization", October 2010.

7.2. Informative References

[3] Faber, T., Touch, J. and W. Yue, "The TIME-WAIT state in TCP and Its Effect on Busy Servers", Proc. Infocom 1999 pp. 1573-1583.

[4] Davari, Oren, Bhatia, Roberts, Montini "Transporting PTP messages (1588) over MPLS Networks", October 2011

[Fab1999] Faber, T., Touch, J. and W. Yue, "The TIME-WAIT state in TCP and Its Effect on Busy Servers", Proc. Infocom 1999 pp. 1573-1583.

8. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Sébastien Jobert
France Telecom Orange
2 avenue Pierre Marzin
22300 LANNION, FRANCE

Email: sebastien.jobert@orange.com