

Internet Draft

Puneet Agarwal
Pluris
Bora A. Akyol
Cisco Systems
Informational
October 2001

Document: [draft-agarwal-mpls-ttl-01.txt](#)

Category

Expires: April 2002

TTL Processing in MPLS Networks

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 describes TTL processing in hierarchical MPLS networks.

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](#)].

[1.](#) Introduction and Motivation

This document describes TTL processing in hierarchical MPLS

networks. We believe that this document adds details that have not been addressed in [[MPLS-ARCH](#), [MPLS-ENCAPS](#)], and that the methods presented in this document complement [[MPLS-DS](#)].

[2.](#) Changes from version-00 of draft

Based on the feedback on the -00 version of the draft (including the presentation at IETF-50), the following changes have been incorporated in this draft:

- (a) Added a section to state how the draft enhances previously published RFCs, modifies the functionality specified in the RFCs, adds new functionality not specified in the RFCs and adds details to clarify what is specified in the RFCs
- (b) Added a statement that signaling the TTL mode is out of the scope of this Internet Draft
- (c) Clarified that iTTL can be decremented by a value greater than 1 (depending on configuration)
- (d) Clarified that iTTL should be $\geq N$ before decrementing by N
- (e) Cleared the confusion caused by the use of the term "IP Header/Label" by replacing with the more generic term "header"
- (f) Took out unneeded text in [section 4.5](#)
- (g) Update the terminology for Pipe, Short Pipe and Uniform based on [draft-ietf-mpls-diff-ext-09.txt](#) (instead of the -07 version)
- (h) Add a separate section to visually describe the TTL processing in the above 3 modes (incorporate the pictures from the slides)

[3.](#) TTL Processing in MPLS Networks

3.1. Changes to [RFC 3032](#) [[MPLS-ENCAPS](#)]

- a) [[MPLS-ENCAPS](#)] only covers the Uniform Model and does NOT address the Pipe Model or the Short Pipe Model. This draft will address these 2 models and for completeness will also address the Uniform Model.
- b) [[MPLS-ENCAPS](#)] does not cover hierarchical LSPs. This draft will address this issue.
- c) [[MPLS-ENCAPS](#)] does not define TTL processing in the presence of Penultimate Hop Popping (PHP). This draft will address this issue.

3.2. Terminology and Background

As defined in [[MPLS-ENCAPS](#)], MPLS packets use a MPLS shim header

that indicates the following information about a packet:

- a. MPLS Label (20 bits)
- b. TTL (8 bits)
- c. Bottom of stack (1 bit)
- d. Experimental bits (3 bits)

The experimental bits were later redefined in [\[MPLS-DS\]](#) to indicate the scheduling and shaping behavior that could be associated with a MPLS packet.

[MPLS-DS] also defined two models for MPLS tunnel operation: Pipe and Uniform models. In the Pipe model, a MPLS network acts like a conduit when MPLS packets traverse the network such that only the LSP ingress and egress points are visible to nodes that are outside the tunnel. A "short" variation of the Pipe Model is also defined in [\[MPLS-DS\]](#) to differentiate between different egress forwarding and QoS treatments. On the other hand, the Uniform model makes all the nodes that a LSP traverses visible to nodes outside the tunnel. We will extend the Pipe and Uniform models to include TTL processing in the following sections. Furthermore, TTL processing when performing Penultimate Hop Pop (PHP) is also described in this document. For a detailed description of Pipe and Uniform models, please see [\[MPLS-DS\]](#).

TTL processing in MPLS networks can be broken down into two logical blocks: (i) the incoming TTL determination to take into account any tunnel egress due to MPLS Pop operations; (ii) packet processing of (possibly) exposed packet & outgoing TTL.

We also note here that signaling treatment for TTL behavior using either RSVP-TE or LDP is out of the scope of this document.

3.3. New Terminology

iTTL: The TTL value to use as the incoming TTL. No checks are performed on the iTTL.

oTTL: This is the TTL value used as the outgoing TTL value. It is always (iTTL - 1) unless otherwise stated.

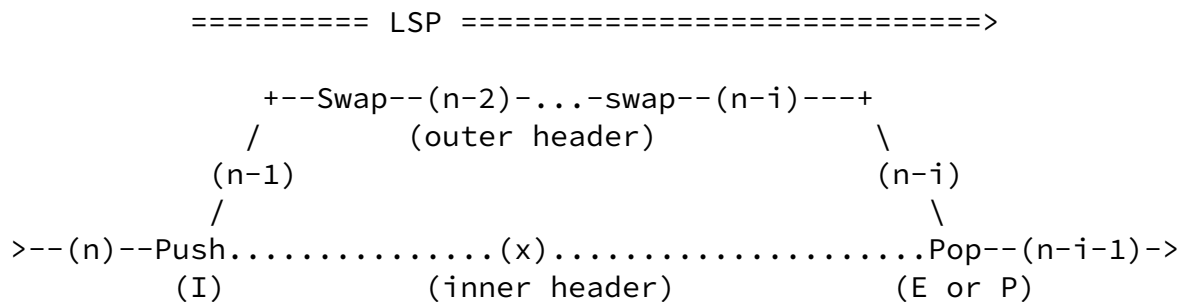
oTTL Check: Check if oTTL is greater than 0. If the oTTL Check is false, then the packet is not forwarded. Note that the oTTL check is performed only if any outgoing TTL (either IP or MPLS) is set to oTTL.

4. TTL Processing in different Models

This section describes the TTL processing for LSPs conforming to each of the 3 models (Uniform, Short Pipe and Pipe) in the presence/absence of PHP (where applicable)

4.1. TTL Processing for Uniform Model LSPs (with or without PHP)

(consistent with [\[MPLS-ENCAPS\]](#)):

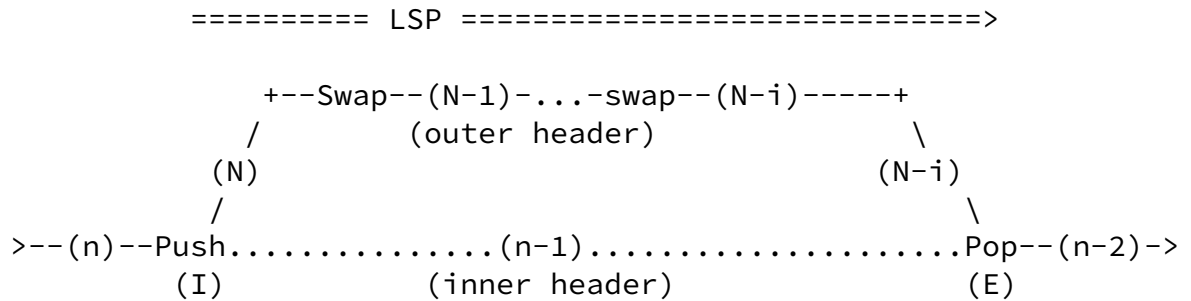


(n) represents the TTL value in the corresponding header
(x) represents non-meaningful TTL information
(I) represents the LSP ingress node
(P) represents the LSP penultimate node
(E) represents the LSP Egress node

This picture shows TTL processing for a uniform model MPLS LSP. Note that the inner and outer TTLs of the packets are synchronized at tunnel ingress and egress.

4.2. TTL Processing for Short Pipe Model LSPs

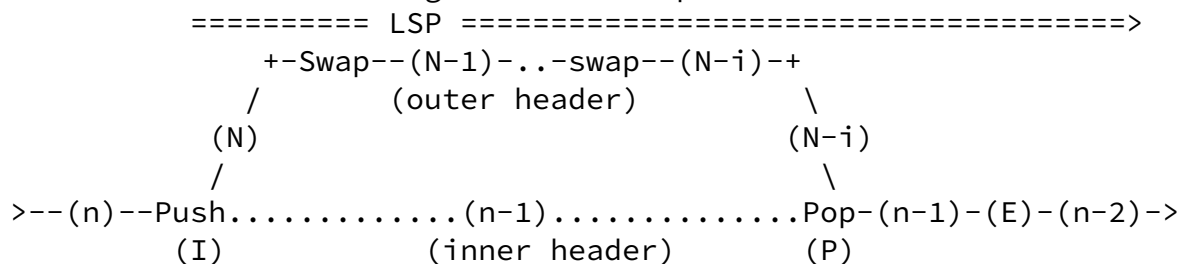
4.2.1. TTL Processing for Short Pipe Model LSPs without PHP



(N) represents the TTL value (may have no relationship to n)
(n) represents the tunneled TTL value in the encapsulated header
(I) represents the LSP ingress node
(E) represents the LSP Egress node

Short Pipe Model was introduced in [[MPLS-DS](#)]. In the short pipe model, the forwarding treatment at the egress LSR is based on the tunneled packet as opposed to the encapsulating packet.

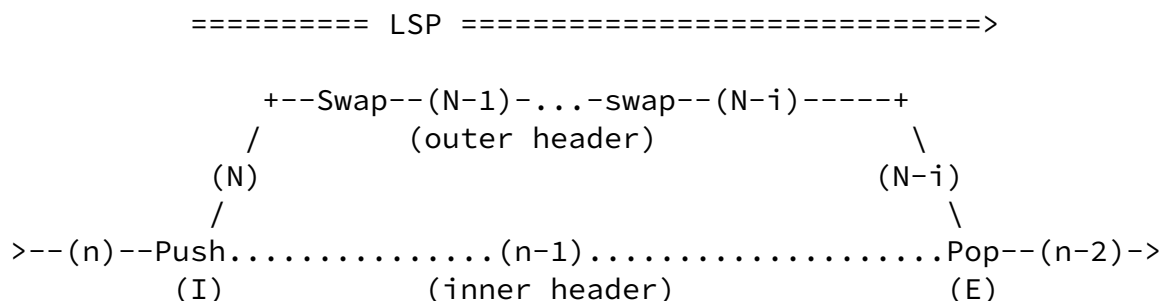
4.2.2. TTL Processing for Short Pipe Model with PHP:



(N) represents the TTL value (may have no relationship to n)
(n) represents the tunneled TTL value in the encapsulated header
(I) represents the LSP ingress node
(P) represents the LSP penultimate node
(E) represents the LSP Egress node

Note that at the end of short pipe model LSP the TTL of the tunneled packet has been decremented by two either with or without PHP.

4.3. TTL Processing for Pipe Model LSPs (without PHP only):



(N) represents the TTL value (may have no relationship to n)

(n) represents the tunneled TTL value in the encapsulated header

(I) represents the LSP ingress node

(E) represents the LSP Egress node

From the TTL perspective, the treatment for a Pipe Model LSP is identical to the Short Pipe Model without PHP.

4.4. Incoming TTL (iTTL) determination

If the incoming packet is an IP packet, then the iTTL is the TTL value of the incoming IP packet.

If the incoming packet is a MPLS packet and we are performing a Push/Swap/PHP, then the iTTL is the TTL of the topmost incoming label.

If the incoming packet is a MPLS packet and we are performing a Pop (tunnel termination), the iTTL is based on the tunnel type (Pipe or Uniform) of the LSP that was popped. If the popped label belonged to a Pipe model LSP, then the iTTL is the value of the TTL field of the

Agarwal & Akyol [draft-agarwal-mpls-ttl-01.txt](#)

TTL Processing in MPLS Networks

5

October 2001

header exposed after the label was popped (note that for the purpose of this draft, the exposed header may be either an IP header or an MPLS label). If the popped label belonged to a Uniform model LSP, then the iTTL is equal to the TTL of the popped label. If multiple Pop operations are performed sequentially, then the procedure given above is repeated with one exception: the iTTL computed during the previous Pop is used as the TTL of subsequent label being popped; i.e. the TTL contained in the subsequent label is essentially ignored and replaced with the iTTL computed during the previous pop.

4.5. Outgoing TTL Determination and Packet Processing

After the iTTL computation is performed, the oTTL check is performed.

If the oTTL check succeeds, then the outgoing TTL of the (labeled/unlabeled) packet is calculated and packet headers are updated as defined below.

If the packet was routed as an IP packet, the TTL value of the IP packet is set to oTTL (iTTL - 1). The TTL value(s) for any pushed label(s) are determined as described in [section 4.6](#).

For packets that are routed as MPLS, we have four cases:

- 1) Swap-only: The routed label is swapped with another label and the TTL field of the outgoing label is set to oTTL.
- 2) Swap followed by a Push: The swapped operation is performed as described in (1). The TTL value(s) of any pushed label(s) are determined as described in [section 4.6](#).
- 3) Penultimate Hop Pop (PHP): The routed label is popped. The oTTL check should be performed irrespective of whether the oTTL is used to update the TTL field of the outgoing header. If the PHPed label belonged to a short Pipe model LSP, then the TTL field of the PHP exposed header is neither checked nor updated. If the PHPed label was a Uniform model LSP, then the TTL field of the PHP exposed header is set to the oTTL. The TTL value(s) of additional labels are determined as described in [section 4.6](#)
- 4) Pop: The pop operation happens before routing and hence it is not considered here.

4.6. Tunnel Ingress Processing (Push)

For each pushed Uniform model label, the TTL is copied from the label/IP-packet immediately underneath it.

For each pushed Pipe model or Short Pipe model label, the TTL field is set to a value configured by the network operator. In most implementations, this value is set to 255 by default.

4.7. Implementation Remarks

- 1) Although iTTL can be decremented by a value larger than 1 while it is being updated or oTTL is being determined, this feature should be only used for compensating for network nodes that are not capable of

- decrementing TTL values such as OXCs.
- 2) Whenever iTTL is decremented, the implementor must make sure that the value does not go negative.
 - 3) In the short pipe model with PHP enabled, the TTL of the tunneled packet is unchanged after the PHP operation.

[5.](#) Conclusion

This Internet Draft describes how TTL field can be processed in a MPLS network. We clarified the various methods that are applied in the presence of hierarchical tunnels and completed the integration of Pipe and Uniform models with TTL processing.

6. Security Considerations

This document does not add any new security issues other than the ones defined in [[MPLS-ENCAPS](#), [MPLS-DS](#)].

[7.](#) References

[MPLS-ARCH] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label Switching Architecture," [RFC 3031](#).

[MPLS-ENCAPS] E. Rosen, D. Tappan, G. Fedorkow, Y. Rekhter, D. Farinacci, T. Li, A. Conta, "MPLS Label Stack Encoding," [RFC3032](#).

[MPLS-DS] F. Le Faucheur, L. Wu, B. Davie, S. Davari, P. Vaananen, R. Krishnan, P. Cheval, J. Heinanen, "MPLS Support of Differentiated Services," [draft-ietf-mpls-diff-ext-09.txt](#). (Work in progress)

[8.](#) Author's Addresses

Puneet Agarwal
Pluris
10455 Bandley Drive
Cupertino, CA 95014
Email: puneet@pluris.com

Bora Akyol
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
Email: bora@cisco.com

