

LSR WG
Internet-Draft
Updates: 5073 (if approved)
Intended status: Standards Track
Expires: 3 September 2026

K. Kompella
Juniper Networks
2 March 2026

Multipath Traffic Engineering Capabilities
draft-kompella-lsr-mpotecap-01

Abstract

Multipath Traffic Engineering (MPTE) combines two approaches to traffic management: equal-cost multipath and constraint-based traffic engineering, offering a powerful new way to engineer networks. To avail of this, a node (possibly an ingress of a MPTE tunnel, or a path computation agent) must have information about the topology, link and node characteristics of a network so that it can compute the components of the MPTE tunnel. One important (node) characteristic is whether a given node supports MPTE, i.e., whether it can participate in the provisioning and maintenance of the tunnel.

This memo shows how this information can be distributed in the IGP via Link State Routing TE Capabilities.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 3 September 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	3
1.1.1. Definition of Commonly Used Terms	3
2. MPTE Capabilities	4
3. IANA Considerations	4
4. Security Considerations	4
5. References	4
5.1. Normative References	4
5.2. Informative References	5
Author's Address	6

1. Introduction

[I-D.kompella-teas-mppte] introduces the notion of multipath traffic engineering (MPTE). It describes how an entity (MPTE DAG computer or MC) can compute a directed acyclic graph (DAG) from one or more ingress nodes to one or more egress nodes that meets given traffic engineering (TE) constraints. This entity (usually one of the ingresses, or a path computation engine) will need information about the network to do the computation, most of which is available in IGP TE extensions. Once the computation is done, the MC communicates the result to the signaling source (SS) which then signals (or provisions) the MPTE tunnel via one of the following protocols: RSVP-TE [RFC3209], PCEP [RFC5440] or BGP [RFC4271].

One key piece of information that is not currently in the IGP extensions is whether or not a given node supports MPTE, i.e., is capable of sending and receiving MPTE updates that create and maintain the tunnel. An MPTE tunnel cannot be setup through such a node, and thus the MC has to take this into account. This memo fills this gap.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here. These words may also appear in this document in lower case as plain English words, absent their normative meanings.

1.1.1. Definition of Commonly Used Terms

This section provides definitions for terms and abbreviations that have a specific meaning to the MPTE protocol and that are used throughout this memo.

constraints: desired properties of paths between ingresses and egresses.

directed acyclic graph: a directed graph that has no cycles.

directed graph (DAG): a set of nodes and directed links. A network is represented by a directed graph.

egress: an end node of an MPTE DAG.

ingress: a starting node of an MPTE DAG.

link: A (directed) edge between two nodes. A pair of nodes may have 0 or more links between them. A link between nodes *u* and *v* will be denoted by (*u*, *v*, *i*), where *i* is *u*'s oif for the link. A link may have associated attributes, in particular, a metric.

MPTE: multipath TE with constraints that uses multiple paths from one or more ingresses to one or more egresses.

MPTED: an MPTE DAG result of computation on MPTE constraints.

MPTED computer (MC): the entity computing the MPTED, typically the ingress (if there is a single ingress) or a Path Computation Element

MPTEP: MPTE protocol: the protocol used to signal MPTEDs.

MPTE tunnel: the signaled entity that carries the traffic from ingresses to egresses along the MPTED.

node: a vertex of a graph. A node may have associated attributes.

PCEP: Path Computation Element communication protocol.

signaling source (SS): the initiator of MPTE signaling to establish, update or destroy an MPTE tunnel.

TE: traffic engineering

2. MPTE Capabilities

[RFC5073] describes IGP protocol extension for the discovery of the TE capabilities of a node. This memo extends that with three new capabilities: whether a node is capable of processing MPTE RSVP-TE messages, and whether a node is capable of processing MPTE PCEP messages. The two capabilities are as follows:

- * MR bit: when set, this flag indicates that the node can process MPTE RSVP-TE messages.
- * MP bit: when set, this flag indicates that the node can process MPTE PCEP messages.
- * MB bit: when set, this flag indicates that the node can process MPTE BGP messages.

These bits are encoded in the TE Node Capability Descriptor defined in [RFC5073].

This Descriptor is carried in ISIS and OSPF as defined in the same RFC.

3. IANA Considerations

IANA is asked to allocate three bits for the above capabilities in the Link State Routing TE Capabilities registry.

4. Security Considerations

This document specifies the content of the TE Node Capability Descriptor TLV in IS-IS and OSPF to be used for MPLS-TE path computation. As this TLV is not used for SPF computation or normal routing, the extensions specified here have no direct effect on IP routing. Tampering with this TLV may have an effect on Traffic Engineering computation. Mechanisms defined to secure IS-IS Link State PDUs [RFC3567], OSPF LSAs [RFC2154], and their TLVs can be used to secure this TLV as well.

5. References

5.1. Normative References

[I-D.kompella-teas-mppte]

Kompella, K., Jalil, L., Khaddam, M., and A. Smith,
"Multipath Traffic Engineering", Work in Progress,
Internet-Draft, draft-kompella-teas-mppte-01, 7 July 2025,
<<https://datatracker.ietf.org/doc/html/draft-kompella-teas-mppte-01>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119,
DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/rfc/rfc2119>>.

[RFC5073] Vasseur, J.P., Ed. and J.L. Le Roux, Ed., "IGP Routing
Protocol Extensions for Discovery of Traffic Engineering
Node Capabilities", RFC 5073, DOI 10.17487/RFC5073,
December 2007, <<https://www.rfc-editor.org/rfc/rfc5073>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

5.2. Informative References

[RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with
Digital Signatures", RFC 2154, DOI 10.17487/RFC2154, June
1997, <<https://www.rfc-editor.org/rfc/rfc2154>>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V.,
and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP
Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001,
<<https://www.rfc-editor.org/rfc/rfc3209>>.

[RFC3567] Li, T. and R. Atkinson, "Intermediate System to
Intermediate System (IS-IS) Cryptographic Authentication",
RFC 3567, DOI 10.17487/RFC3567, July 2003,
<<https://www.rfc-editor.org/rfc/rfc3567>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A
Border Gateway Protocol 4 (BGP-4)", RFC 4271,
DOI 10.17487/RFC4271, January 2006,
<<https://www.rfc-editor.org/rfc/rfc4271>>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation
Element (PCE) Communication Protocol (PCEP)", RFC 5440,
DOI 10.17487/RFC5440, March 2009,
<<https://www.rfc-editor.org/rfc/rfc5440>>.

Author's Address

Kireeti Kompella
Juniper Networks
Sunnyvale, California 94089
United States of America
Email: kireeti.ietf@gmail.com