

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 08, 2016

C. Franke
D. Lamparter
NetDEF
July 07, 2015

IS-IS Point-to-Multipoint operation
draft-lamparter-isis-p2mp-00

Abstract

This document describes a new mode operation for IS-IS. In addition to the existing LAN and point-to-point modes of operation, a point-to-multipoint mode is defined. This mode is useful for operation both on networks with more than two routers where multicast is expensive in comparison to unicast, as well on networks where creating a LAN pseudonode cannot adequately reflect metrics.

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 <http://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 January 08, 2016.

Copyright Notice

Copyright (c) 2015 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	P2MP Pseudocircuits	3
2.1.	Pseudocircuit behaviour	3
2.1.1.	Representation in LSPs	3
2.1.2.	Forwarding	4
2.2.	Neighbor IS discovery	4
2.2.1.	Manual configuration	4
2.2.2.	Lower layer autodiscovery	4
2.2.3.	Multicast autodiscovery	4
2.3.	Adjacency formation	5
2.4.	Pseudocircuit teardown	5
3.	PDU generation and processing	6
3.1.	LSP, CSNP, PSNP and all other non-Hello PDU types	6
3.2.	LAN and P2P Hellos	6
3.3.	P2MP Hello PDUs	6
3.3.1.	Without Three-way Adjacency TLV	6
3.3.2.	With Three-way Adjacency TLV	7
4.	PDU encoding	7
4.1.	P2MP Hello PDU	7
5.	IANA Considerations	8
6.	Configuration model	9
7.	Security Considerations	9
8.	Privacy Considerations	9
9.	Acknowledgements	9
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	9
	Authors' Addresses	10

[1.](#) Introduction

The core functionality of the protocol outlined in this document consists of an additional Subnetwork dependent function per [Section 8](#) of ISO/IEC 10589:2002 [[IS-IS](#)]. In that regard, the next section can be understood as new section "8.5 Point-to-multipoint networks".

The outlined protocol is remotely similar to the derelict section "8.3 ISO 8208 subnetworks" [[X.25](#)] in that multiple point-to-point adjacencies are established on an interface.

Point-to-multipoint operation of IS-IS is thus not a new idea; in fact [section 6.2](#) already mentions "multipoint links." This document re-enables the concept.

[2.](#) P2MP Pseudocircuits

In place of ISO 8473 call management [[CLNS](#)] establishing sessions, this document establishes pairwise "pseudocircuits" between two IS on a common medium. Multiple such pseudocircuits can normally exist on one P2MP interface.

Each pseudocircuit operates, aside from subnetwork attachment procedures, exactly as a PtP link.

It should be noted that while the Multicast autodiscovery mechanism requires broadcast capability, no other part of the protocol does. The P2MP interface type can be used on non-transitive and/or non-broadcast interfaces.

[2.1.](#) Pseudocircuit behaviour

An implementation maintains a set of pseudocircuits per P2MP interface. Each pseudocircuit is uniquely identified by the local interface and peer's SNPA address.

Each participating system MUST use a consistent SNPA address per local interface. A change in SNPA address results in all neighbors treating the interface as distinct from the previous one. A system MAY support multiple SNPA addresses per interface by treating them as distinct interfaces.

Unnumbered media are not supported by this protocol. Each participant on a link MUST have a unique SNPA address on that link.

As pseudocircuits are dynamic in nature, they must be created and destroyed as needed.

[2.1.1.](#) Representation in LSPs

Pseudocircuits are represented in LSPs as a regular PtP circuit would be. As a result, their treatment in SPF calculations is also identical to PtP circuits.

2.1.2. Forwarding

In scenarios where pseudocircuits do not form a full mesh of all participants on a medium, the best path for a packet may be through the same interface as the one it was received on.

Systems implementing this specification MUST therefore support forwarding packets on the same interface that they were received on. This applies only to interfaces configured for P2MP operation.

2.2. Neighbor IS discovery

The discovery machinery produces as output a "candidate neighbor list," which is a list of possible neighbor's SNPAs and is maintained per P2MP interface. Additional information (metric, etc.) MAY be associated with the SNPA, but this is not part of the discovery mechanism.

The candidate neighbor list is conceptual in nature; adding and removing entries results in pseudocircuit creation and deletion. It need not be implemented as an actual list.

The list is the union of the result of each of the following sources of neighbor information. A neighbor may be listed by multiple sources: it MUST NOT be removed while any source still lists it.

The IS-IS implementation SHOULD provide user configuration to disable individual mechanisms and/or filter candidate neighbors both as a security and as misconfiguration preventing measure.

2.2.1. Manual configuration

A list of neighbor IS MAY be configured by the user, providing possible neighbor's SNPAs on an interface.

2.2.2. Lower layer autodiscovery

Lower protocol layers (VPLS, IEEE 802.11) may be able to provide a list of attached neighbors. This list MAY be fed into the candidate neighbor list.

2.2.3. Multicast autodiscovery

For broadcast capable networks, this document defines an autodiscovery mechanism based on multicasting Hello packets. This mechanism MAY be disabled by the user, but MUST be implemented for all lower layer link types with (limited or full) broadcast capability.

Multicast autodiscovery P2MP Hello PDUs are distinguished by not containing a [RFC5303](#) Three-Way Adjacency TLV. Receiving such a Hello places the sending IS on the candidate list for as long as the PDU's holdtime indicates.

A system MAY implement a receive-only passive multicast autodiscovery mode where it adds candidates (forms pseudocircuits) on receiving P2MP Hello PDUs, but does not send such PDUs itself.

If either passive or full multicast autodiscovery is enabled, receiving a P2MP Hello PDU with a Three-Way Adjacency TLV also adds the neighbor to the candidate list.

2.3. Adjacency formation

Since there is no underlying protocol layer partitioning the link into actual PtP circuits in this case, additional handling is required on bringing up the individual pseudocircuit adjacencies.

To prevent different pseudocircuits from "bleeding" into each other, implementations of this protocol MUST enable [\[RFC5303\]](#) on all P2MP pseudocircuits, with changes as follows:

- o Received IIH PDUs on P2MP pseudocircuits without the Point-to-Point Three-Way Adjacency option MUST be discarded. The TLV is not optional on pseudocircuit adjacencies but rather mandatory.

2.4. Pseudocircuit teardown

Pseudocircuits are destroyed when their PtP state machine has transitioned into "Down" state and they are no longer listed as candidates by any discovery mechanism. The conditions for pseudocircuit removal are thus:

- o PtP adjacency timeout functionality (IS-IS [section 8.2.6](#) [\[IS-IS\]](#)) has moved the pseudocircuit to Down state, or it never moved out of Down state
- o The holdtime of any multicast autodiscovery P2MP Hello PDUs has expired.
- o Manual configuration or lower layer autodiscovery no longer lists the neighbor.

As long as a pseudocircuit is present, its PtP state machine will, as expected for PtP circuits, trigger transmission of further Hello PDUs, even when it is in Down state.

3. PDU generation and processing

3.1. LSP, CSNP, PSNP and all other non-Hello PDU types

All PDU types that are not Hello PDUs, e.g. L1/L2 LSP PDUs, L1/L2 CSNP/PSNP PDUs, FS-LSP, FS-CSNP and FS-PSNP PDUs are processed on P2MP interfaces by dispatching them to an individual pseudocircuit by looking up the PDU's source address in the interface's list of P2MP pseudocircuits. This includes all functions defined in [Section 7](#) of ISO/IEC 10589:2002 [[IS-IS](#)] and PDUs added in [[RFC7176](#)] and [[RFC7356](#)].

While possible, this document does not suggest sending such PDUs to multicast destinations. All systems MUST send these PDUs to unicast lower layer destinations so that they are received only by the pseudocircuit's neighbor system. However, systems SHOULD NOT check the destination address on receipt to allow future optimisations.

3.2. LAN and P2P Hellos

Received LAN and P2P Hello PDUs MUST be discarded when received on an interface configured for P2MP operation. The system MAY notify the operator of such a misconfiguration.

TODO: hitless migration possible?

3.3. P2MP Hello PDUs

P2MP Hello PDUs are divided in two categories depending on the presence of a [RFC5303](#) Three-way Adjacency TLV.

3.3.1. Without Three-way Adjacency TLV

For each P2MP interface that has the multicast autodiscovery mechanism enabled, a system periodically sends P2MP Hello PDUs without Three-way Adjacency TLV to a lower layer specific multicast address. The periodic timer SHOULD be configurable and is subject to jitter per [Section 10.1](#) of ISO/IEC 10589:2002 [[IS-IS](#)].

On receipt, such PDUs are not processed on any pseudocircuit's PtP state machine. They are processed on pseudocircuit management level as described in :

- o If no pseudocircuit matches the PDU's lower layer source address, one is created for that address
- o The pseudocircuit is held in existence at least until the PDU's holdtime expires.

No information other than the neighbor's SNPA is carried from the PDU onto the pseudocircuit, thus most TLVs that are normally valid in Hellos become pointless. However, all TLVs that affect PDU processing (in particular authentication TLVs) are processed normally.

On IEEE 802 lower layers that support 48bit MAC addresses, the AllIntermediateSystems multicast address (09-00-2B-00-00-05) is used as destination when sending these PDUs.

3.3.2. With Three-way Adjacency TLV

P2MP Hello PDUs with a [RFC5303](#) Three-way Adjacency TLV are both processed and generated by the PtP Hello state machine ([Section 8.2](#) of ISO/IEC 10589:2002 [[IS-IS](#)]) of their associated pseudocircuit.

This implies that received PDUs are dispatched by looking up their source SNPA address among the pseudocircuits associated with the receiving interface (possibly creating a pseudocircuit).

For Hello PDUs sent by the pseudocircuit's PtP Hello state machine, the destination SNPA address is set to the pseudocircuit's neighbor SNPA address.

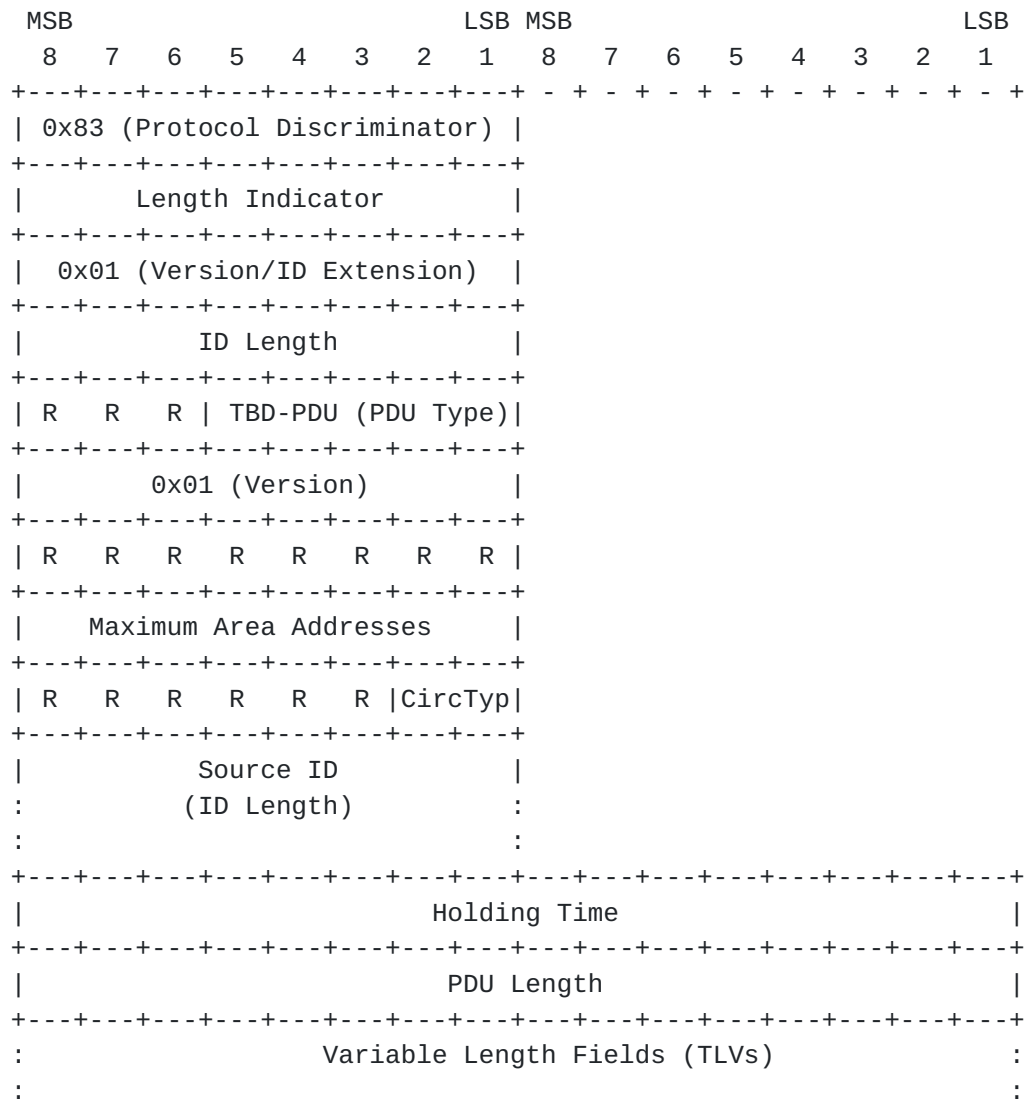
TLV validity and processing is exactly as for PtP Hellos.

4. PDU encoding

(Note this section follows ISO 10589 [section 9](#), particularly in numbering bits starting at 1 for the LSB.)

4.1. P2MP Hello PDU

Figure 1: P2MP Hello PDU format



All fields function exactly as for Point-to-Point IS-IS hello PDUs, as specified in [Section 9.7](#) of ISO/IEC 10589:2002 [[IS-IS](#)].

Note the Local Circuit ID field is absent. Its function is replaced by [[RFC5303](#)].

5. IANA Considerations

IANA is requested to allocate a value from the "IS-IS PDU Registry" as follows:

Type: TBD-PDU

Description: P2MP-HELLO-PDU

Reference: This document

Applicability of sub-TLVs for this PDU (per "TLV Codepoints Registry") is per the "IIH" column.

6. Configuration model

TODO: YANG model.

7. Security Considerations

TODO.

8. Privacy Considerations

TODO.

9. Acknowledgements

TODO.

10. References

10.1. Normative References

- [IS-IS] ISO/IEC, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, 2002.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5303] Katz, D., Saluja, R., and D. Eastlake, "Three-Way Handshake for IS-IS Point-to-Point Adjacencies", [RFC 5303](#), October 2008.

10.2. Informative References

- [CLNS] ISO/IEC, "Protocol for providing the connectionless-mode network service: Protocol specification", ISO/IEC 8473-1:1998, 1998.
- [RFC7176] Eastlake, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", [RFC 7176](#), May 2014.

[RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", [RFC 7356](#), September 2014.

[X.25] ISO/IEC, "X.25 Packet Layer Protocol for Data Terminal Equipment", ISO/IEC 8208:2000, 2000.

Authors' Addresses

Christian Franke
NetDEF
Leipzig
DE

Email: chris@opensourcerouting.org

David Lamparter
NetDEF
Leipzig 04103
Germany

Email: david@opensourcerouting.org