

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 19, 2014

L. Ginsberg
S. Mirtorabi
S. Previdi
A. Roy
Cisco Systems
October 16, 2013

IS-IS Support for Unidirectional Links
draft-ietf-isis-udl-01.txt

Abstract

This document defines support for the operation of IS-IS over Unidirectional Links without the use of tunnels or encapsulation of IS-IS Protocol Data Units. Adjacency establishment when the return path from the router at the receive end of a unidirectional link to the router at the transmit end of the unidirectional link is via another unidirectional link is supported. The extensions defined here are backwards compatible - only the routers directly connected to a unidirectional link need to be upgraded.

Requirements Language

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](#)].

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 April 19, 2014.

Copyright Notice

Copyright (c) 2013 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
2.	Encoding Extensions	4
2.1.	UDL LSPs and the UDL-TLV	4
2.2.	UDL Intermediate System Neighbors sub-TLV	5
2.2.1.	UDL Point-to-Point Intermediate System Neighbor Sub-TLV	5
2.2.2.	UDL LAN Intermediate System Neighbor Sub-TLV	6
2.2.3.	Sub-TLVs Associated w an IS Neighbor	7
2.3.	UDL Manual Area Addresses sub-TLV	10
3.	Adjacency Establishment	11
3.1.	Adjacency Establishment in Point-to-Point Mode	11
3.2.	Adjacency Establishment in Broadcast Mode	12
3.3.	UDL link metric configuration	12
4.	Adjacency Maintenance	13
4.1.	Adjacency Maintenance by IS-T	13
4.2.	Adjacency Maintenance by IS-R	14
4.3.	Use of BFD	14
4.4.	Graceful Restart Support	15
5.	Operation of the Update Process on a UDL	15
6.	Support for UDL on the Return Path	16
7.	IANA Considerations	17
8.	Security Considerations	18
9.	Acknowledgements	18
10.	References	18
10.1.	Normative References	18
10.2.	Informational References	19
	Authors' Addresses	20

1. Introduction

Operation of IS-IS depends upon two-way connectivity. Adjacencies are formed by exchanging hellos on a link, flooding of the link state database is made reliable by exchanges between neighbors on a link, etc. However, there are deployments where operation of the protocol is desired over links which are unidirectional i.e., one end of the link can only send Protocol Data Units (PDUs) and one end of the link can only receive PDUs. Traditional methods of supporting Unidirectional Links (UDLs) have involved establishing a tunnel from the Intermediate System (IS) at the receive end of the UDL to the IS at the transmit end of the UDL, encapsulating/decapsulating the IS-IS PDUs as they enter/exit the tunnel, and associating the PDUs received via the tunnel with the UDL at the transmit end. This typically requires static configuration and may introduce Maximum Transmission Unit (MTU) issues due to the required encapsulation.

This specification defines extensions to the protocol which support correct and reliable operation of IS-IS over UDLs without the need for tunnels or any form of encapsulation.

2. Encoding Extensions

Although the IS at the transmit end of a UDL link (IS-T) can send IS-IS PDUs normally on the link, the IS at the receive end of a UDL link (IS-R) requires assistance from other ISs in the network to pass the information it would normally send directly to IS-T. The Update Process as defined in [[IS-IS](#)] allows information generated by one IS in the network to be reliably flooded to all other ISs in the network using Link State PDUs (LSPs). The extensions defined here utilize LSPs to allow IS-R to send information normally sent in hellos (IIHs) or sequence number PDUs (SNPs) to IS-T in LSPs. As LSPs are flooded to all ISs in an area/sub-domain, care is taken to minimize the LSP churn necessary to support adjacency establishment and maintenance between IS-T and IS-R.

2.1. UDL LSPs and the UDL-TLV

Routers on the receive end of a UDL MUST reserve at least one LSP (for each level supported on the UDL) to advertise the UDL information described below. Such LSPs are referred to as UDL-LSPs although the only distinction between a UDL-LSP and other LSPs is in the TLV information which is present in such an LSP. LSP #0 MUST NOT be used to send UDL information. UDL-LSPs have the following special characteristics:

1. The only TLV which may be advertised in UDL-LSPs is the UDL TLV described below and (optionally) an Authentication TLV and/or Purge Originator Identification TLV [[RFC6232](#)]. This requirement is enforced by the originator of the UDL-LSP but is not checked by receiving systems i.e., other TLVs which are included in a UDL-LSP are processed normally. The reason for the restriction is to minimize the number of LSPs which have UDL information content.
2. Routers on the transmit side of a UDL flood UDL-LSPs regardless of the existence of an adjacency in the UP state on that circuit. Flooding of UDL-LSPs on circuits other than a UDL is as specified in [[IS-IS](#)] i.e., no special handling.

A new TLV is defined in which UDL specific information appears. All information in a UDL-TLV is encoded in sub-TLVs. UDL sub-TLVs are formatted as specified in [[RFC5305](#)]. The format of the UDL-TLV is therefore:

No. of octets	
+-----+	
Type (11)	1
(To be assigned by IANA)	
+-----+	
Length	1
+-----+	
Sub-TLVs	3 - 255
: :	
+-----+	

2.2. UDL Intermediate System Neighbors sub-TLV

UDL links may operate in Point-to-Point mode or in broadcast mode (assuming the subnetwork is a broadcast subnetwork). There are therefore two types of Intermediate System Neighbors sub-TLVs defined. A UDL-TLV MUST NOT contain more than one Intermediate System Neighbors sub-TLV. If multiple Intermediate System Neighbors sub-TLVs appear in a UDL-TLV all information in that UDL-TLV MUST be ignored.

2.2.1. UDL Point-to-Point Intermediate System Neighbor Sub-TLV

The UDL Point-to-Point Intermediate System Neighbor Sub-TLV describes an adjacency on a UDL which is operating in Point-to-Point mode i.e. either a Point-to-Point subnetwork or a LAN subnetwork operating in Point-to-Point mode as described in [[RFC5309](#)]. The information

encoded follows the format for the Point-to-Point Three-Way Adjacency TLV as defined in [[RFC5303](#)] but may also include the local LAN address when the underlying subnetwork is a LAN.

	No. of octets
+-----+	
Type (240)	1
(To be assigned by IANA)	
+-----+	
Length (9 + ID Length)	1
to (15 + ID Length)	
+-----+	
Adjacency 3-way state	1
+-----+	
Extended Local Circuit ID	4
+-----+	
Neighbor System ID	ID Length
+-----+	
Neighbor Extended Local	4
Circuit ID	
+-----+	
Local LAN Address	6
+-----+	

2.2.2. UDL LAN Intermediate System Neighbor Sub-TLV

The UDL LAN Intermediate System Neighbor sub-TLV describes an adjacency on a UDL operating in broadcast mode on a LAN subnetwork.

	No. of octets
+-----+	
Type (6)	1
(To be assigned by IANA)	
+-----+	
Length (7 + ID Length)	1
+-----+	
Neighbor LAN ID	ID Length + 1
+-----+	
Local LAN Address	6
+-----+	

2.2.3. Sub-TLVs Associated w an IS Neighbor

A number of sub-TLVs require the presence of a UDL IS-Neighbor sub-TLV (either Point-to-Point or LAN) in the UDL-TLV in order to provide appropriate context for the information being advertised. These sub-TLVs are described in the sub-sections below.

2.2.3.1. UDL LSP Range sub-TLV

The content of this sub-TLV describes a range of LSPs for which the originating router requires an update. Only the neighbor specified in the associated UDL IS-Neighbor sub-TLV processes the LSP range mentioned in this sub-TLV.

	No. of octets
+-----+	
Type (8)	1
(To be assigned by IANA)	
+-----+	
Length (ID Length + 2)* 2	1
+-----+	
Start LSP ID	ID Length + 2
+-----+	
End LSP ID	ID Length + 2
+-----+	

2.2.3.2. UDL LSP Entry sub-TLV

The content of this sub-TLV describes LSPs for which the originating router requires an update. Only the neighbor specified in the associated UDL IS-Neighbor sub-TLV processes the LSP entries specified in this sub-TLV.

	No. of octets
+-----+	
Type (9)	1
(To be assigned by IANA)	
+-----+	
Length (10 + ID Length)*N	1
+-----+	
: LSP Entries :	
+-----+	

Each LSP Entry has the following format:

+-----+	
Remaining Lifetime	2
+-----+	
LSP ID	ID Length + 2
+-----+	
LSP Sequence Number	4
+-----+	
Checksum	2
+-----+	

[2.2.3.3.](#) Protocols Supported sub-TLV

This sub-TLV specifies the set of Network Layer Protocol Identifiers (NLPIDs) that the originating system is capable of forwarding as defined in [[RFC1195](#)].

	No. of octets
+-----+	
Type (129)	1
(To be assigned by IANA)	
+-----+	
Length	Number of NLPIDs
+-----+	
: NLPIDs :	1 octet/NLPID
+-----+	

[2.2.3.4.](#) IP Address sub-TLV

This sub-TLV specifies the set of IP addresses configured on the interface as defined in [[RFC1195](#)].

No. of octets	
+-----+	
Type (132)	1
(To be assigned by IANA)	
+-----+	
Length	4 * # of addresses
+-----+	
: IP Address(es) :	4 octets/address
+-----+	

[2.2.3.5.](#) Multi-Topology sub-TLV

This sub-TLV specifies the set of topology identifiers supported as defined in [[RFC5120](#)].

No. of octets	
+-----+	
Type (229)	1
(To be assigned by IANA)	
+-----+	
Length	2 * # of MTIDs
+-----+	
: MTIDs :	2 octets/MTID
+-----+	

NOTE: All flag bits defined in [[RFC5120](#)] MUST be transmitted as 0 and ignored on receipt.

[2.2.3.6.](#) IPv6 Interface Address sub-TLV

This sub-TLV specifies the set of IPv6 addresses assigned on the local interface as defined in [[RFC5308](#)]. Addresses MUST be link local addresses.

No. of octets	
+-----+	
Type (232)	1
(To be assigned by IANA)	
+-----+	
Length	16 * # of IPv6 addresses
+-----+	
: IPv6 Addresses :	16 octets/Address
+-----+	

2.2.3.7. IPv6 Global Interface Address sub-TLV

This sub-TLV specifies the set of global IPv6 addresses assigned on the local interface as defined in [RFC6119]. . Addresses MUST be global or unique local addresses.

No. of octets	
+-----+	
Type (233)	1
(To be assigned by IANA)	
+-----+	
Length	16 * # of IPv6 addresses
+-----+	
: IPv6 Addresses :	16 octets/Address
+-----+	

2.3. UDL Manual Area Addresses sub-TLV

This sub-TLV specifies the set of manualAreaAddresses of the originating system. No other sub-TLVs are allowed in a UDL-TLV which has this sub-TLV. Any other sub-TLVs in such a UDL-TLV are ignored on receipt.

No. of octets	
+-----+	
Type (1)	1
(To be assigned by IANA)	
+-----+	
Length	1
+-----+	
: Area Address(es) :	
+-----+	

Each Area Address has the following format:

+-----+	
Address Length	1
+-----+	
Area Address	Address Length
+-----+	

3. Adjacency Establishment

An adjacency over a UDL link may be established over a link operating in Point-to-Point mode (including a LAN subnetwork configured to operate in Point-to-Point mode) or a link operating in broadcast mode. Operation in either mode is identical except for some differences in the manner of adjacency establishment as specified in the following sub-sections.

IS-T utilizes the set of manualAreaAddresses advertised by IS-R in a UDL Manual Area Address sub-TLV in combination with the UDL Intermediate System Neighbor sub-TLV(s) to IS-T advertised by IS-R to determine the level(s) associated with any adjacency to IS-R.

3.1. Adjacency Establishment in Point-to-Point Mode

Adjacency establishment makes use of Three Way Handshake as defined in [RFC5303] when operating in Point-to-Point mode. When operating over a LAN subnetwork, the use of point-to-point operation over LAN as defined in [RFC5309] is also used.

IS-T initiates adjacency establishment by sending Point-to-Point IIHs over the UDL as normal i.e., including Three-Way Handshake TLV. Note that the local circuit ID specified by IS-T need only be unique among the set of Point-to-Point UDL links supported by IS-T on which IS-T is at the transmit end.

Upon receipt of a Point-to-Point IIH IS-R creates an adjacency in the INIT state with IS-T and advertises the existence of the adjacency in its UDL-LSP(s) utilizing the UDL Point-to-Point Intermediate System Neighbor sub-TLV. The Local LAN address is included if the link is a LAN subnetwork operating in Point-to-Point mode. UDL-LSPs of the appropriate level(s) are generated according to the type of the adjacency with IS-T.

When IS-T receives the UDL-LSP(s) generated by IS-R containing the UDL Point-to-Point Intermediate System Neighbor sub-TLV it validates the 3 way information and, if valid, transitions its adjacency to UP state. In subsequent Point-to-Point IIHs IS-T includes IS-R's circuit ID information as indicated in the UDL Point-to-Point IS Neighbor sub-TLV in its 3 way handshake TLV. A complete set of CSNPs is sent to IS-R for the level(s) appropriate for the type of adjacency. LSPs which are updated as a result of the existence of the adjacency to IS-R are sent to IS-R, but IS-T does NOT propagate its full LSP Database. This is done to minimize the amount of redundant flooding.

IS-R uses normal adjacency bring up rules based on the 3 way

handshake information it receives in Point-to-Point IIHs from IS-T and advertises its IS neighbor to IS-T in the usual manner i.e. in an LSP other than a UDL-LSP. Following transition of the adjacency to IS-T to the UP state IS-R MAY request IS-T to flood its complete LSP Database by sending an LSP Range sub-TLV to IS-T in a UDL-LSP.

3.2. Adjacency Establishment in Broadcast Mode

IS-T initiates adjacency establishment by sending LAN IIHs of the appropriate level(s) over the UDL as normal. IS-T specifies itself in the LAN ID field of the IIH, including a non-zero circuit ID. Note that the local circuit ID specified by IS-T need only be unique among the set of LAN UDL links supported by IS-T on which IS-T is at the transmit end. This is because pseudo-node LSPs will never be generated for a UDL. Operation in broadcast mode supports a UDL with a single IS-T and multiple IS-Rs.

Upon receipt of a LAN IIH PDU IS-R creates an adjacency in the INIT state with IS-T and advertises the existence of the adjacency in its UDL-LSP(s) utilizing the UDL LAN Intermediate System Neighbor sub-TLV. UDL-LSPs of the appropriate level(s) are generated according to the levels supported by IS-R and IS-T.

When IS-T receives the UDL-LSP(s) generated by IS-R containing the UDL LAN Intermediate System Neighbor sub-TLV(s) it validates the LANID and, if valid, transitions its adjacency to UP state. In subsequent LAN IIH PDUs, IS-T includes IS-R's LAN Address as indicated in the UDL LAN IS Neighbor info. A complete set of CSNPs for the appropriate level is sent over the circuit. LSPs which are updated as a result of the existence of the adjacency to IS-R are sent to IS-R, but IS-T does NOT propagate its full LSP Database. This is done to minimize the amount of redundant flooding.

IS-R uses normal adjacency bring up rules based on the IS Neighbor LAN Address information it receives in LAN IIH PDUs from IS-T and advertises its IS neighbor to IS-T in an LSP other than a UDL-LSP. Note that there is no pseudo-node on a UDL LAN circuit - therefore both IS-T and IS-R MUST advertise an IS Neighbor TLV to each other, not to a pseudo-node. This is identical to what is done on a Point-to-Point subnetwork. Following transition of the adjacency to IS-T to the UP state IS-R MAY request IS-T to flood its complete LSP Database by sending an LSP Range sub-TLV to IS-T in a UDL-LSP.

3.3. UDL link metric configuration

What metrics are configured on a UDL depend upon the intended use of the UDL. If the UDL is to be used for unicast forwarding, then IS-T should be configured with the value appropriate to its intended

preference in the network topology and IS-R should be configured with maximum link metric ($2^{24} - 1$) as defined in [RFC5305] (assuming wide metrics are in use). If the UDL is to be used for building a multicast Reverse Path Forwarding tree, then IS-R should be configured with the value appropriate to its intended preference in the network topology and IS-T should be configured with maximum link metric ($2^{24} - 1$). If the link is to be used for both unicast forwarding and multicast, then it is necessary to have two different metric configurations and perform two different SPF calculations. This may be achieved through the use of multi-topology extensions as defined in [RFC5120]. Note that the configured link metrics have no bearing on adjacency establishment - they only affect the building of a Shortest Path Tree (SPT).

4. Adjacency Maintenance

This section defines how adjacencies are maintained once established. Adjacency maintenance is defined without the need to send periodic UDL-LSP updates as this would be a significant burden on the entire network.

4.1. Adjacency Maintenance by IS-T

IS-T sends IIH PDUs as normal on a UDL. As IS-R does NOT send IIH PDUs to IS-T, IS-T maintains the adjacency to IS-R so long as all of the following conditions are TRUE:

- o IS-T has a valid UDL-LSP from IS-R which includes Point-to-Point UDL IS Neighbor information or LAN UDL IS Neighbor information (as appropriate) regarding the adjacency IS-R has with IS-T on the UDL.
- o IS-T can calculate a return path rooted at IS-R to IS-T which does not traverse the UDL on which the adjacency is associated

When either of the above conditions becomes FALSE, IS-T brings down its adjacency to IS-R. Note that the return path calculation is only required when a topology change occurs in the network. It therefore need only be done in conjunction with a normal event driven SPF calculation.

NOTE: Immediately after the adjacency to IS-R has come up, if the only available return path traverses a UDL link on which the adjacency is still in the process of coming UP, the return path check will fail. This is possible because we bypass normal flooding rules to allow the UDL-LSP to be flooded even when the adjacency is not UP on a UDL link (as described later in this document). If IS-T

immediately brings the adjacency to IS-R down in this case, a circular dependency condition arises. To avoid this, if the return path check fails immediately after the adjacency comes up, a timer T_p is started. The timer is cancelled when a return path check succeeds. If the timer expires, IS-T brings down the adjacency to IS-R. A recommended value for the timer T_p is a small multiple (e.g., "twice") of the estimated time necessary to propagate LSPs across the entire domain.

Although it is unorthodox to bring up an adjacency without confirmed two way connectivity, the extension is well grounded because the receipt of IS-R's UDL-LSP by IS-T is indicative of the existence of a return path even though it cannot yet be confirmed by examination of the LSP database. This unconfirmed two way connectivity is a condition which we do not want to persist indefinitely - hence the use of timer T_p .

[4.2.](#) Adjacency Maintenance by IS-R

IS-R maintains its adjacency with IS-T based on receipt of IIHs from IS-T as normal. So long as IS-T follows the rules for adjacency maintenance described in the previous section this is sufficient.

Further protection against pathological behavior on the part of IS-T (e.g., failure to perform the return path calculation after a topology change) MAY be implemented by IS-R. When IS-R receives a CSNP from IS-T which contains an SNP entry identifying an LSP which is not in IS-R's Link State Database (LSDB) a timer T_f is started for each such LSP. This includes entries which are older than, newer than, or non-existent in IS-R's LSDB. The timer T_f is cancelled if:

- o The associated LSP is received by IS-R on any circuit by normal operation of the Update process or
- o A subsequent set of CSNPs received from IS-T does not include the LSP entry

If any timer T_f expires IS-R brings down the adjacency with IS-T.

In the absence of pathological behavior by IS-T the T_f extension is not required. Its use is therefore optional.

[4.3.](#) Use of BFD

A multi-hop BFD session [[RFC5883](#)] MAY be established between IS-T and IS-R. This can be used to provide fast failure detection. If used, this would also make the calculation by IS-T of a return path from IS-R to IS-T optional.

Support for [[RFC6213](#)] requires that the BFD session come up before the IS-IS adjacency comes up when both neighbors advertise BFD support. In the event that there is a UDL link on the return path from IS-R to IS-T and the adjacency on that link is also in the process of coming up this could introduce a circular dependency between the state of the BFD sessions and the state of the UDL adjacencies. Therefore [[RFC6213](#)] is NOT supported on UDLs.

4.4. Graceful Restart Support

Graceful restart as defined in [[RFC5306](#)] is NOT supported on UDLs.

In the event IS-R is restarting, signaling of restart state would require IS-R to regenerate UDL-LSPs prior to synchronization of the LSPDB. In the event IS-T is restarting, LSPDB synchronization would require the sending of CSNPs from IS-R to IS-T - which is not supported.

5. Operation of the Update Process on a UDL

For purposes of LSP propagation IS-T views the UDL as if it were a broadcast subnetwork where IS-T is the Designated Intermediate System (DIS). This is true regardless of the mode of operation of the circuit (point-to-point or broadcast). Therefore, IS-T propagates new LSPs on the UDL as they arrive but after sending an LSP on the UDL the SRM flag for that LSP is cleared i.e. no acknowledgement for the LSP is required or expected. IS-T also sends periodic CSNPs on the UDL.

IS-R cannot propagate LSPs to IS-T on the UDL. IS-R also cannot acknowledge LSPs received from IS-T on the UDL. In this respect IS-R operates on the UDL in a manner identical to a non-DIS on a broadcast circuit. If an LSP entry in a CSNP received from IS-T identifies an LSP which is "newer than" an LSP in IS-R's LSDB, IS-R MAY request the LSP from IS-T by sending a UDL-LSP with an LSP entry as described above. Since IS-R's UDL-LSP(s) will be propagated throughout the network even though the information is only of use to IS-Ts, it is recommended that some small delay occur between the receipt of a CSNP from IS-T and the generation of a UDL-LSP with an updated LSP entry by IS-R so as to allow for the possible receipt of the LSP either from IS-T or on another link.

If the number of LSP entries to be requested exceeds the space available in the UDL TLV associated with the adjacency to IS-T, IS-R MUST NOT generate multiple UDL TLVs associated with the same adjacency. Instead it should maintain the state of SSN flags appropriately for the LSP entries that require updates and send

additional LSP entries (if necessary) in a subsequent UDL-LSP after the previously requested updates arrive.

Use of the LSP Range sub-TLV by IS-R allows more efficient encoding of a request for multiple LSPs. This could be especially useful following an adjacency UP event on a UDL. As described in [Section 3](#), IS-T does NOT propagate its full LSP database following transition of an adjacency to IS-R to the UP state. This is consistent with IS-T operating in the role of DIS on a broadcast circuit. If IS-R has neighbors on other circuits it is possible that it will have received LSPs from other neighbors. In such a case flooding of the full LSP database by IS-T would be redundant. It is therefore left to the discretion of IS-R to request those portions of the LSP database which are not current. This is consistent with IS-R operating as a non-DIS on a broadcast circuit.

On receipt of a UDL-LSP generated by IS-R, IS-T checks the neighbor information in each UDL-TLV. If the information matches an existing adjacency that IS-T has with IS-R then IS-T sets SRM flag on the UDL for any LSPs in its LSDB which are "newer" than the corresponding entries IS-R sent in LSP Entry sub-TLVs in UDL TLVs. SRM flags are also set on the UDL for LSPs which fall in the ranges specified in LSP Range sub-TLVs in UDL TLVs. UDL-TLVs associated with adjacencies to routers other than IS-T are ignored by IS-T.

6. Support for UDL on the Return Path

If all return paths from IS-R to IS-T traverse a UDL, then in order to bring up the adjacency between IS-T and IS-R at least one of the adjacencies on a return path UDL must already be UP. This is required because IS-T relies on receiving the UDL-LSP(s) generated by IS-R in order to bring up its adjacency. In order to overcome a circular dependency in the case where multiple pairs of UDL neighbors are trying to bring up an adjacency at the same time, an extension to LSP propagation rules is required.

When a new UDL-LSP is received by any IS which has one or more active UDLs on which it is operating as an IS-T, the set of neighbors other than the local system which are advertised in UDL-TLVs in the received UDL-LSP is extracted - call this UDL-LSP-ISN-SET. A return path from the originating IS-R to each neighbor in the UDL-LSP-ISN-SET is calculated. If there is no return path to one or more neighbors in this set periodic propagation of that UDL-LSP on all UDLs on which the local system acts as IS-T is initiated regardless of the state of an adjacency on that UDL. Periodic transmission of that UDL-LSP continues until a return path to all neighbors in the UDL-LSP-ISN-SET exists. This calculation is redone whenever the UDL-

LSP is updated and when a topology change in the network occurs as a result of updates to the LSDB. Note that periodic retransmission is only done on UDLs on which the local system acts as IS-T.

If the network is partitioned the lack of a return path from a given IS-R to a given IS-T may persist. It is therefore recommended that the periodic retransmission employ an exponential backoff timer such that when the partition persists the periodic retransmission period is long enough so as to not represent a significant burden. It is recommended that the periodic retransmission be initially set to the locally configured CSNP interval. Note that periodic retransmission is only performed on UDL links and if an IS-R has previously received the same UDL-LSP it will silently ignore the retransmission since the UDL-LSP will already be in its LSDB. Unnecessary reflooding of the retransmitted UDL-LSP beyond the UDL does not occur.

IS-R MUST accept and propagate UDL-LSPs received on a UDL even when there is no adjacency in the UP state on the UDL circuit. Flooding of UDL-LSPs by IS-R uses normal flooding rules. LSPs received by IS-R on the UDL which do NOT include UDL TLVs are discarded unless the adjacency is UP (normal processing).

This extension allows establishment of an adjacency on a UDL even when the return path transits another UDL which is also in the process of bringing up an adjacency. The periodic nature of the flooding is meant to compensate for the unreliability of the flooding. After the adjacency is UP, IS-R can request LSPs from IS-T by putting LSP entries into UDL-LSPs - but that ability is not available until the adjacency is UP.

7. IANA Considerations

This document requires the definition of a new IS-IS TLV to be reflected in the "IS-IS TLV Codepoints" registry:

Type	Description	IIH	LSP	SNP	Purge
----	-----	---	---	---	-----
11	Unidirectional Link Information	N	Y	N	Y

This document requires that a new IANA registry be created to control the assignment of sub-TLV code points to be advertised within a Unidirectional Link Information TLV. The registration procedure is "Expert Review" as defined in [RFC5226]. The following sub-TLVs are defined by this document. Values are suggested values subject to assignment by IANA.

Value	Description
-----	-----
1	Manual Area Addresses
6	LAN IS Neighbor
8	LSP Range
9	LSP Entry
129	Protocols Supported
132	IP Interface Address
229	Multi-topology
232	IPv6 Interface Address
233	IPv6 Global Interface Address
240	Point-to-Point IS Neighbor

8. Security Considerations

Security concerns for IS-IS are addressed in [\[IS-IS\]](#), [\[RFC5304\]](#), and [\[RFC5310\]](#).

9. Acknowledgements

The idea of supporting IS-IS on UDLs without using tunnels or encapsulation was originally introduced in the US patent "Support of unidirectional link in IS-IS without IP encapsulation and in presence of unidirectional return path" (patent number: 7,957,380), by Sina Mirtorabi, Abhay Kumar Roy, Lester Ginsberg.

10. References

10.1. Normative References

- [IS-IS] "Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473), ISO/IEC 10589:2002, Second Edition.", Nov 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), December 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), February 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5303] Katz, D., Saluja, R., and D. Eastlake, "Three-Way Handshake for IS-IS Point-to-Point Adjacencies", [RFC 5303](#), October 2008.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), October 2008.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", [RFC 5308](#), October 2008.
- [RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", [RFC 6119](#), February 2011.

[10.2.](#) Informational References

- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", [RFC 5304](#), October 2008.
- [RFC5306] Shand, M. and L. Ginsberg, "Restart Signaling for IS-IS", [RFC 5306](#), October 2008.
- [RFC5309] Shen, N. and A. Zinin, "Point-to-Point Operation over LAN in Link State Routing Protocols", [RFC 5309](#), October 2008.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), February 2009.
- [RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", [RFC 5883](#), June 2010.
- [RFC6213] Hopps, C. and L. Ginsberg, "IS-IS BFD-Enabled TLV", [RFC 6213](#), April 2011.
- [RFC6232] Wei, F., Qin, Y., Li, Z., Li, T., and J. Dong, "Purge Originator Identification TLV for IS-IS", [RFC 6232](#), May 2011.

Authors' Addresses

Les Ginsberg
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
USA

Email: ginsberg@cisco.com

Sina Mirtorabi
Cisco Systems
3800 Zankar Road
San Jose, CA 95134
USA

Email: smirtora@cisco.com

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

Abhay Roy
Cisco Systems
560 McCarthy Blvd.
Milpitas, CA 95135
USA

Email: akr@cisco.com

