IS-IS Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 16, 2011

N. Shen
T. Li
Cisco Systems, Inc.
S. Amante
Level 3 Communications
M. Abrahamsson
Tele2

October 13, 2010

# IS-IS Reverse Metric TLV for Network Maintenance Events
## draft-amante-isis-reverse-metric-00

## Abstract

This document describes an improved IS-IS neighbor management scheme which can be used to enhance network performance by allowing operators to quickly and accurately shift traffic away from a point-to-point or multi-access LAN interface by allowing one IS-IS router to signal to a second, adjacent IS-IS neighbor to adjust its IS-IS metric that should be used to temporarily reach the first IS-IS router during network maintenance events.

## 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 16, 2011.

## Copyright Notice

Table of Contents

1. Introduction . . . . . 3
1.1. Link Isolation Challenges . . . . . 3
1.2. IS-IS Reverse Metric . . . . . 4
1.3. Specification of Requirements . . . . . 4

2. IS-IS Reverse Metric TLV . . . . . 4

3. Elements of Procedure . . . . . 6
3.1. Multi-Access LAN Procedures . . . . . 7

4. Reverse Metric TLV Example Use Case . . . . . 8

5. Operational Considerations . . . . . 9

6. Security Considerations . . . . . 9

7. IANA Considerations . . . . . 9

8. Acknowledgements . . . . . 9

9. References . . . . . 10
9.1. Normative References . . . . . 10
9.2. Informative References . . . . . 10

Authors' Addresses . . . . . 10

# 1. Introduction

The IS-IS [ISO 10589] routing protocol has been widely used in Internet Service Provider IP/MPLS networks. Operational experience with the protocol, combined with ever increasing requirements for lossless operations have demonstrated some operational issues. This document describes one issue and a new mechanism for improving it.

## 1.1. Link Isolation Challenges

During network maintenance events, operators substantially increase the IS-IS metric simultaneously on both devices attached to the same link. In doing so, the devices generate new Link State Protocol Data Units (LSP's) that are flooded throughout the network and cause all routers to gradually shift traffic onto alternate paths with very little, to no, disruption to in-flight communications by applications or end-users. When performed successfully, this allows the operator to confidently perform disruptive fault diagnosis and restoration on a link without disturbing ongoing communications in the network.

The challenge with the above solution are as follows. First, it is quite common to have routers with several hundred interfaces onboard and individual interfaces that are transferring several hundred Gigabits/second to Terabits/second of traffic. Thus, it is imperative that operators accurately identify the same point-to-point link on two, separate devices in order to increase (and, afterward, decrease) the IS-IS metric appropriately. Second, the aforementioned solution is very time consuming and even more error-prone to perform when its necessary to temporarily remove a multi-access LAN from the network topology. Specifically, the operator needs to configure ALL devices's that have interfaces attached to the multi-access LAN with an appropriately high IS-IS metric, (and then decrease the IS-IS metric to its original value afterward). Finally, with respect to multi-access LAN's, there is currently no method to bidirectionally isolate only a single node's interface on the LAN when performed more fine-grained diagnosis and repairs to the multi-access LAN.

In theory, use of a Network Management System (NMS) could improve the accuracy of identifying the appropriate subset of routers attached to either a point-to-point link or a multi-access LAN as well as signaling from the NMS to those devices, using a network management protocol, to adjust the IS-IS metrics on the pertinent set of interfaces. The reality is that NMS are, to a very large extent, not used within Service Provider's networks for a variety of reasons. In particular, NMS do not interoperate very well across different vendors or even separate platform families within the same vendor.

The risks of misidentifying one side of a point-to-point link or one

or more interfaces attached to a multi-access LAN and subsequently increasing its IS-IS metric are potentially increased latency, jitter or packet loss. This is unacceptable given the necessary performance requirements for a variety of applications, the customer perception for near lossless operations and the associated, demanding Service Level Agreement's (SLA's) for all network services.

### 1.2. IS-IS Reverse Metric

This document proposes that the routing protocol itself be the transport mechanism to allow one IS-IS router to advertise to an adjacent node on a point-to-point or multi-access LAN link a "reverse metric" in a IS-IS Hello (IIH) PDU. This would allow an operator to only configure a single router, set a "reverse metric" on a link and have traffic bidirectionally shift away from that link gracefully to alternate, viable paths.

### 1.3. Specification of Requirements

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 [RFC2119].

## 2. IS-IS Reverse Metric TLV

The Reverse Metric TLV is composed of 1 octet for the Type, 1 octet that specifies the number of bytes in the Value field and a variable-length Value field. The Value field starts with a 1 octet field of Flags followed by a 3 octet field containing an IS-IS Metric and, lastly, a 1 octet Traffic Engineering (TE) sub-TLV length field representing the length of a variable number of Extended Intermediate System (IS) Reachability sub-TLV's. If the 'S' bit in the Flags field is set to 1, then the Value field MUST also contain data of 1 or more Extended IS Reachability sub-TLV's.

The Reverse Metric TLV is optional. The Reverse Metric TLV may be present in any IS-IS Hello PDU.

```
TYPE: TBD
LENGTH: variable (5 - 255 octets)
VALUE:
   Flags (1 octet)
   Metric (3 octets)
   TE sub-TLV length (1 octet)
   TE sub-TLV data (0 - 250 octets)
```

Flags

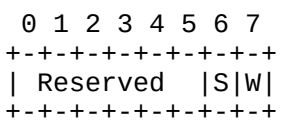


Figure 1: Flags

The Reverse Metric TLV Type is TBD. Please refer to IANA Considerations, in Section 7, for more details.

The Metric field contains a 24-bit unsigned integer equal to the IS-IS metric a neighbor SHOULD set in its own IS Neighbors TLV or Extended IS Reachability TLV for point-to-point links, or Pseudonode LSP by the Designated Intermediate System (DIS) for multi-access LAN's, back toward the router that originated this Reverse Metric TLV. An IS-IS neighbor MUST overwrite the existing IS-IS metric, in its corresponding IS Neighbors, Extended IS Reachability TLV or Pseudonode LSP, with the value it received in the Metric field of the Reverse Metric TLV.

The Metric field may be a "default metric", in the range of 0-63, or a "Traffic Engineering Default Metric" [RFC5305], in the range of 0-2^(24-1) depending on the configuration of router's interface that is originating the Reverse Metric TLV. It is RECOMMENDED that implementations, by default, place the appropriate maximum default metric value, 63 or 2^(24-1), in the Metric field of the Reverse Metric TLV, since the most common use is to remove the link from the topology, except for use as a last-resort path.

There is currently only two Flag bits defined.

W bit (0x01): The "Whole LAN" bit is only used in the context of multi-access LAN's. When a Reverse Metric TLV is transmitted from a (non-DIS) node to the DIS, if the "Whole LAN" bit is set (1), then a DIS MUST replace the IS-IS metric for all nodes in the Pseudonode LSP with the Metric value received in the Reverse Metric TLV. If the "Whole LAN" bit is not set (0), then a DIS MUST replace the IS-IS metric in the Pseudonode LSP for just the node from whom the Reverse Metric TLV was received. Please refer to the Elements of Procedure, in Section 3, for additional details. In addition, the W bit MUST be unset (0) when a Reverse Metric TLV is transmitted in a IIH PDU onto a point-to-point link to an IS-IS neighbor.

S bit (0x02): The "TE sub-TLV" bit MUST be set (1) when an IS-IS router wishes to signal that its neighbor alter parameters contained in the neighbor's IPv4 and/or IPv6 Traffic Engineering "Extended IS

Reachability TLV", as defined in [RFC5305] for IPv4 and [I-D.ietf-isis-ipv6-te] for IPv6. An IS-IS router MUST overwrite only the subset of its own TE sub-TLV's with those sub-TLV's received from a neighbor in the Reverse Metric TLV.

The S bit MUST NOT be set (0) when an IS-IS router does not have TE sub-TLV's that it wishes to send to its IS-IS neighbor.

The values used for the "IPv4 Interface Address" and "IPv6 Interface Address" TE sub-TLV's MUST be set to all zero when sent inside a Reverse Metric TLV. In addition, the "IPv4 Neighbor Address" and "IPv6 Neighbor Address" TE sub-TLV's MUST be set to local node's interface address(es) that is originating a Reverse Metric TLV.

## 3. Elements of Procedure

A router SHOULD first update its own IS-IS metric and/or Traffic Engineering parameters in its IS Neighbors TLV, Extended IS Reachability TLV or Pseudonode LSP, then recompute its SPF tree plus corresponding route metrics and, lastly, flood its updated LSP's, using normal IS-IS mechanisms, as well as start advertising a Reverse Metric TLV in IIH's toward a neighbor. A router MUST advertise a Reverse Metric TLV toward a neighbor only for the period during which it wants a neighbor to temporarily update its IS-IS metric or TE parameters.

When a router receives a Reverse Metric TLV it MUST immediately update its own IS Neighbors TLV, Extended IS Reachability TLV or Pseudonode LSP with the received value(s) in the Metric field or TE sub-TLV's, then recalculate its SPF tree and associated route metrics and, finally, flood its updated LSP's to other IS-IS routers. Note that on a Multi-Access LAN, only the DIS SHOULD act upon information contained in a received Reverse Metric TLV. All non-DIS nodes MUST silently ignore a received Reverse Metric TLV. Please refer to Section 3.1 for additional details with respect to Multi-Access LAN's and the Reverse Metric TLV.

Routers that receive a Reverse Metric TLV MAY send a syslog message or SNMP trap, in order to assist in rapidly identifying the node in the network that is asserting an IS-IS metric or Traffic Engineering parameters different from that which is configured locally on the device. Routers MUST scan the Metric value and TE sub-TLV's in all subsequently received Reverse Metric TLV's. If changes are observed by a receiver of the Reverse Metric TLV in the Metric value, number of TE sub-TLV's or data in the TE sub-TLV's, the receiving router MUST update its advertised IS-IS metric or Traffic Engineering parameters in the appropriate TLV's, recompute its SPF tree and

corresponding metrics to IP prefixes and, finally, flood new LSP's to other IS-IS routers.

When a router stops receiving a Reverse Metric TLV it MUST immediately update its own IS Neighbors TLV, Extended IS Reachability TLV or Pseudonode LSP with the previously configured IS-IS metric value and/or Traffic Engineering parameters, recalculate its SPF and associated route metrics and flood updated LSP's within the IS-IS domain.

It is RECOMMENDED that implementations provide a capability to disable any changes to a node's default metric or Traffic Engineering parameters based upon receipt of properly formatted Reverse Metric TLV's.

If the router does not understand the Reverse Metric TLV or is explicitly configured to ignore received Reverse Metric TLV's, then it will not update nor flood a new IS Neighbors TLV, Extended IS Reachability TLV or Pseudonode LSP and should not recompute its SPF tree or update metrics associated with corresponding routes.

### 3.1. Multi-Access LAN Procedures

In the case of multi-access LAN's, the "W" Flags bit is used to signal from a non-DIS to the DIS whether to change the metric and/or Traffic Engineering parameters for all nodes in the Pseudonode LSP or a single node on the LAN, (the originator of the Reverse Metric TLV).

A non-DIS node, i.e.: Router B, attached to a multi-access LAN will send a Reverse Metric TLV with the W bit set to 0 to the DIS, when Router B wishes the DIS to replace the metric and/or TE parameters contained in the Pseudonode LSP specific to just Router B. Other non-DIS nodes, i.e.: Routers C and D, may simultaneously send a Reverse Metric TLV with the W bit set to 0 to request the DIS replace their respective metric and/or TE parameters contained in the Pseudonode LSP. When the DIS receives a properly formatted Reverse Metric TLV with the W bit set to 0, the DIS MUST only change the metric and/or TE parameters contained in its Pseudonode LSP for the specific neighbor that sent the Reverse Metric TLV.

It is possible for one node, Router A, to signal to the DIS with the W bit set to 1, in which case the DIS would replace the metric and/or TE parameters for all neighbor adjacencies in the Pseudonode LSP with the Metric value in the Reverse Metric TLV and transmit a new Pseudonode LSP to all nodes in the IS-IS domain. Later, a second node on the LAN, Router B, could signal to the DIS with the W bit also set to 1. In this case, the DIS MUST use the Reverse Metric TLV Value field(s) advertised by the router with highest MAC address of

the two routers from which it received a Reverse Metric TLV, Router A or B. If Router B's MAC address was highest, then the DIS MUST update the metric and/or Traffic Engineering parameters for all neighbors in its Pseudonode LSP and flood the LSP to all nodes in the IS-IS domain. On the other hand, if Router A's MAC address was highest the DIS will ignore Router B's Reverse Metric TLV and continue to use Router A's Reverse Metric TLV Value field(s) for all neighbors in the Pseudonode LSP. When this occurs, the DIS MAY send a single syslog message or SNMP trap indicating that it has received a Reverse Metric TLV from a neighbor, but is ignoring it due to it being received from a neighbor with a lower MAC address.

Another scenario is that one node, Router A, may signal the DIS with the W bit set to 1. The DIS would update the metric for all neighbors in the Pseudonode LSP and flood the LSP. Later, a second node on the LAN, Router B, could signal the DIS with the W bit set to 0, which indicates to the DIS that Router B is requesting the DIS only update the metric and/or TE parameters for Router B in the Pseudonode LSP. The DIS MUST honor a neighbor's Reverse Metric TLV to update its individual IS-IS metric and/or TE parameters in the Pseudonode LSP even if the DIS receives prior or later requests to assert a Whole LAN metric or TE parameter(s) change from other nodes on the same LAN.

Local configuration on the DIS to adjust the default metric(s) contained in the Pseudonode LSP, as documented in [I-D.shen-isis-oper-enhance] MUST take precedence over received Reverse Metric TLV's.

## 4. Reverse Metric TLV Example Use Case

The following is a brief example illustrating one use case of the Reverse Metric TLV. In order to isolate a point-to-point link from the IS-IS network, an operator would configure one router, Router A, attached to a point-to-point link with a "Reverse Metric". This should not affect the configuration of the existing IS-IS default metric previously configured on the router's interface. Assuming Router A is using IS-IS Extensions for Traffic Engineering [RFC5305], this should trigger Router A to update its Traffic Engineering Default Metric sub-TLV in its own Extended IS Reachability TLV, recompute its SPF tree and corresponding metrics to IP prefixes in the IS-IS domain and begin the process of flooding a new LSP throughout the network. Router A would also begin transmitting a Reverse Metric TLV, with an appropriate Metric value, in an IIH PDU, to its adjacent neighbor, Router B. Upon receipt of the Reverse Metric TLV, Router B would also update its Traffic Engineering Default Metric sub-TLV with the received Metric value in the Reverse

Metric TLV, recalculate its SPF tree and associated route topology as well as start flooding a new LSP containing the updated Extended IS Reachability TLV throughout the network. As nodes in the network receive the associated LSP's from Router A and B and recalculate a new SPF tree, and route topology, traffic should gracefully shift onto alternate paths away from the A-B link; ultimately, after all nodes in the network recompute their SPF tree link A-B should only be used as a link of last-resort. The operator can inspect traffic counters on the A-B interface to determine if the link was successfully isolated from the topology and proceed with necessary fault diagnosis or maintenance of the associated link.

When the maintenance activity is complete, the operator would remove the reverse metric configuration from Router A, which would cease advertisement of the Reverse Metric TLV in IIH PDU's to Router B. Both routers would revert to their originally configured IS-IS metric, recompute new SPF trees and corresponding metrics to IP prefixes and originate new LSP's. As the new LSP's are received and SPF is recalculated by nodes in the IS-IS domain, traffic should gradually shift back onto link A-B.

## 5. Operational Considerations

Since the Reverse Metric TLV may not be recognized by adjacent IS-IS neighbors, operators should inspect input and output traffic throughput counters on the local router to ensure that traffic has bidirectionally shifted away from a link before starting any maintenance activities.

## 6. Security Considerations

This document raises no new security issues for IS-IS.

## 7. IANA Considerations

This document requests that IANA allocate from the IS-IS TLV Codepoints Registry a new TLV, referred to as the "Reverse Metric" TLV, with the following attributes: IIH = y, LSP = n, SNP = n, Purge = n.

## 8. Acknowledgements

The authors would like to thank Mike Shand, Dave Katz, Guan Deng, Ilya Varlashkin, Jay Chen, Les Ginsberg and Peter Ashwood-Smith for

their contributions.

## 9. References

### 9.1. Normative References

[I-D.ietf-isis-ipv6-te]
Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", draft-ietf-isis-ipv6-te-08 (work in progress), September 2010.

[ISO 10589]
ISO, "Intermediate system to Intermediate system routeing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002.

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

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

### 9.2. Informative References

[I-D.shen-isis-oper-enhance]
Shen, N., Li, T., Amante, S., and M. Abrahamsson, "IS-IS Operational Enhancements for Network Maintenance Events", draft-shen-isis-oper-enhance-00 (work in progress), October 2010.

## Authors' Addresses

Naiming Shen
Cisco Systems, Inc.
225 West Tasman Drive
San Jose, CA 95134
USA

Email: naiming@cisco.com

Tony Li
Cisco Systems, Inc.
225 West Tasman Drive
San Jose, CA 95134
USA

Email: tli@cisco.com

Shane Amante
Level 3 Communications
1025 Eldorado Blvd
Broomfield, CO 80021
USA

Email: shane@level3.net

Mikael Abrahamsson
Tele2

Email: swmike@swm.pp.se

IS-IS Working Group
Internet-Draft
Intended status: Informational
Expires: April 3, 2011

N. Shen
T. Li
Cisco Systems, Inc.
S. Amante
Level 3 Communications
M. Abrahamsson
Tele2
September 30, 2010

# IS-IS Operational Enhancements for Network Maintenance Events
draft-shen-isis-oper-enhance-00

## Abstract

This document describes an improved IS-IS neighbor management scheme which can be used to enhance operational experience in terms of convergence speed and finer control of neighbor cost over a LAN.

## 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 3, 2011.

## Copyright Notice

Table of Contents

1. Introduction . . . . . 3
1.1. Interface Shutdown Black Hole . . . . . 3
1.2. LAN of Last Resort . . . . . 3
1.3. Specification of Requirements . . . . . 3

2. Sending Hellos with Fast Exit Notification . . . . . 3

3. Pseudonodes with Non-zero Metrics . . . . . 4
3.1. Operational Considerations . . . . . 5

4. Security Considerations . . . . . 5

5. Acknowledgements . . . . . 5

6. Normative References . . . . . 5

Authors' Addresses . . . . . 6

# 1. Introduction

The IS-IS [ISO 10589] routing protocol has been widely used in Internet Service Provider IP/MPLS networks. Operational experience with the protocol, combined with ever increasing requirements for lossless operations have demonstrated some operational issues. This document describes those issues and some mechanisms for dealing with those issues. These mechanisms do involve implementation support, but do not require protocol changes.

## 1.1. Interface Shutdown Black Hole

One of these operationally problematic issues occurs when IS-IS is disabled on only one side of a link. This can result in a significant delay before neighbor(s) on the other end of the same link notice this change. In turn, this can result in several seconds during which traffic is blackholed, until the IS-IS neighbor(s) time out the adjacency and IS-IS reconverges.

## 1.2. LAN of Last Resort

Another issue stems from a situation when operators want to temporarily make an interface a "last resort" link for transit traffic. This is a straightforward, through cumbersome, operation to perform on a point-to-point link. Each device on the link is reconfigured to use very high metric. This causes traffic to divert to other links in the network. This same operation is more difficult on a multi-access LAN. There, the operator would have to increase the metric on each and every interface attached to the LAN, requiring the reconfiguration of a number of systems.

## 1.3. Specification of Requirements

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 [RFC2119].

# 2. Sending Hellos with Fast Exit Notification

When an operator shuts down IS-IS on an interface, as described in Section 1.1, there is a significant interval before the change is noticed by all adjacencies and traffic is subsequently re-routed around this link. This delay is unnecessary, as neighbors should not have to wait for the adjacency to timeout, particularly when there exist alternate, viable, paths to downstream neighbors. This delay can be eliminated by carefully removing the adjacency between neighbors prior to actually disabling IS-IS on the interface.

An IS-IS adjacency uses the 3-way handshake protocol as defined in [ISO 10589] for multi-access LANs and [RFC3373] for point-to-point links. In both cases, the IS to IS Hello (IIH) message used to establish and maintain the adjacency carries the system identifier of the adjacent systems. The receiving system expects to see its own system identifier listed. If not, then it must drop the adjacency.

An implementation that wishes to avoid the issue in Section 1.1 can do so by sending out a final IIH that includes no neighboring system IDs. When this is received, it should cause all neighbors to drop their adjacencies with the router that sent the IIH. This will also cause the systems to update their Link State Protocol Data Units (LSPs), flood them and reconverge to new paths. The technique is known as Fast Exit Notification.

This approach is not guaranteed. If the final IIH is lost on the link, then the neighboring systems will have to wait to time out the adjacency. Since this is unlikely, it is still a useful optimization. Implementations that require an even higher degree of assurance can retransmit the final IIH, possibly multiple times.

## 3. Pseudonodes with Non-zero Metrics

If an operator wishes to reconfigure a multi-access LAN so that it is only used as a resource of the last resort, then with current mechanisms, the operator must reconfigure each node on the LAN to give the LAN a high metric, as described in Section 1.2. It would be much easier for the operator if they could make a single configuration change that would cause IS-IS to treat the multi-access LAN as a link of last resort.

[ISO 10589] defines the pseudonode LSP as having a metric of zero. This implies that during the Shortest Path First (SPF) calculation, the metric for traversing the LAN is solely based on the metric set by the IS used to access the LAN. Thereby, the pseudonode does not contribute to the cost of traversing the LAN.

However, from the point of view of the SPF calculation, the metric in the pseudonode LSP does not have to be zero. Instead, the metric in a pseudonode LSP could be treated just like a normal LSP and have non-zero metrics to some or all of the systems on the LAN. This can then be used to simplify the operation for turning a LAN into a link of last resort. This could be done by having the Designated Intermediate System (DIS) change all of the metrics within the pseudonode LSP to a high value. This would effectively make the LAN look very ’expensive’ and cause SPF calculations to converge to alternate links, if at all possible.

Because this change to the usage of the pseudonode LSP is in direct contradiction to the existing IS-IS specification, extreme caution is necessary. Implementations that would not interpret a non-zero pseudonode metric correctly might cause forwarding loops. As of this writing, we are actively surveying existing known implementations to determine if setting a non-zero metric in a pseudonode LSP will be interpreted properly.

This technique can also be used to divert traffic away from a subset of the nodes on the LAN. If the DIS increases the metric from the pseudonode to a subset of the systems on the LAN, then traffic will avoid exiting the LAN via that subset of systems.

### 3.1. Operational Considerations

A further simplification is to allow any system to temporarily become the DIS, when it is directed to, and set a non-zero metric in the pseudonode. This is beneficial because the operator would otherwise first have to determine the current DIS, access that system and reconfigure it. If an implementation wishes to support this, then it can provide an operation that both changes its priority on the LAN so that a node first becomes DIS and then generates a new pseudonode LSP with the non-zero metric.

If there is a concern that the DIS may change, it is prudent to define another node on the same LAN with the second highest priority for becoming DIS. This node can be configured to also set the metric in its pseudonode LSP appropriately if it becomes the new DIS.

## 4. Security Considerations

This document raises no new security issues for IS-IS.

## 5. Acknowledgements

The authors would like to thank Mike Shand, Dave Katz, Guan Deng, Ilya Varlashkin, Jay Chen, Peter Ashwood-Smith and Les Ginsberg for their contributions.

## 6. Normative References

[ISO 10589]
ISO, "Intermediate system to Intermediate system routeing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network

Service (ISO 8473)", ISO/IEC 10589:2002.

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

[RFC3373] Katz, D. and R. Saluja, "Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies", RFC 3373, September 2002.

## Authors' Addresses

Naiming Shen
Cisco Systems, Inc.
225 West Tasman Drive
San Jose, CA 95134
USA

Email: naiming@cisco.com

Tony Li
Cisco Systems, Inc.
225 West Tasman Drive
San Jose, CA 95134
USA

Email: tli@cisco.com

Shane Amante
Level 3 Communications
1025 Eldorado Blvd
Broomfield, CO 80021
USA

Email: shane@level3.net

Mikael Abrahamsson
Tele2

Email: swmike@swm.pp.se