Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 6, 2017

N. Shen
Cisco Systems
S. Amante
Apple, Inc.
M. Abrahamsson
T-Systems Nordic
May 5, 2017

IS-IS Routing with Reverse Metric
draft-ietf-isis-reverse-metric-06

Abstract

This document describes the mechanism to allow IS-IS routing to quickly and accurately shift traffic away from either a point-to-point or multi-access LAN interface by signaling to an adjacent IS-IS neighbor with the metric towards itself during network maintenance or other operational events.

Table of Contents

1. Introduction

The IS-IS [ISO10589] 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 the issues and a new mechanism for improving it.

1.1. Node and Link Isolation

IS-IS routing mechanism has the overload-bit, which can be used by operators to perform disruptive maintenance on the router. But in many operational maintenance cases, it is not necessary to displace all the traffic away from this node. It is useful to augment only a

single link or LAN for the maintenance. More detailed descriptions of the challenges can be found in Appendix A and Appendix B of this document.

1.2. Distributed Forwarding Planes

In a distributed forwarding platform, different forwarding line-cards may have interfaces and IS-IS connections to neighbor routers. If one of the line-card's software resets, it may take some time for the forwarding entries to be fully populated on this line-card, in particular if the router is a PE (Provider Edge) router in ISP's MPLS VPN. The IS-IS adjacency may be established with a neighbor router long before the entire BGP VPN prefixes are downloaded to the forwarding table. It is important to signal to the network not to use this particular IS-IS adjacency inbound to this router if possible. Temporarily pushing out the 'Reverse Metric' over this link to discourage the traffic into this line-card will help to reduce the traffic loss in the network. At the meantime, the remote PE routers will select a different set of PE routers for the BGP best path calculation or use a different link towards the same PE router on which another line-card is recovering.

1.3. Spine-Leaf Applications

In the IS-IS Spine-Leaf extension [I-D.shen-isis-spine-leaf-ext], the leaf nodes will perform equal-cost or unequal-cost load sharing towards all the spine nodes. In certain operational cases, for instance, when one of the backbone links on a spine node is congested, this spine node can push a higher metric towards the connected leaf nodes to reduce the transit traffic through this spine node or link.

1.4. LDP IGP Synchronization

In the [RFC5443], a mechanism is described to achieve LDP IGP synchronization by using the maximum link metric value on the interface. But in the case of a new IS-IS node joining the broadcast network (LAN), it is not optimal to change all the nodes on the LAN to the maximum link metric value, as described in [RFC6138]. This Reverse Metric can be used in this case to discourage both outbound and inbound traffic without affecting the traffic of other existing IS-IS nodes on the LAN.

1.5. IS-IS Reverse Metric

This document proposes that the routing protocol itself be the transport mechanism to allow one IS-IS router to advertise a "reverse metric" in an IS-IS Hello (IIH) PDU to an adjacent node on a point-

to-point or multi-access LAN link. This would allow the provisioning to be performed only on a single node, set a "reverse metric" on a link and have traffic bidirectionally shift away from that link gracefully to alternate, viable paths.

This Reverse Metric mechanism is to be used for both point-to-point and multi-access LAN links. Unlike the point-to-point link, IS-IS protocol currently does not have a way to influence the traffic towards a particular node on LAN links. This proposal enables IS-IS routing the capability of altering traffic in both directions on either a point-to-point link or on a multi-access link of a node.

1.6. 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 a 1 octet field of Flags, a 3 octet field containing an IS-IS Metric, and 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. A sender MUST only transmit a single Reverse Metric TLV in a IS-IS Hello PDU.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Type     |     Length    |    Flags      | Metric Offset
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   Metric Offset (Continue)     | sub-TLV Len   |Optional sub-TLV
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

                     Reverse Metric TLV
```

TYPE: TBD
LENGTH: variable (5 - 255 octets)
VALUE:

Flags (1 octet)
Metric Offset (3 octets)

```
      sub-TLV length (1 octet)
      sub-TLV data (0 - 250 octets)

       0 1 2 3 4 5 6 7
      +-+-+-+-+-+-+-+-+
      |   Reserved  |W|
      +-+-+-+-+-+-+-+-+
```

Figure 1: Flags

The Metric Offset field contains a 24-bit unsigned integer of an IS-IS metric that a neighbor SHOULD add to the existing, configured "default metric" contained within its IS Neighbors TLV, Extended IS Reachability TLV's for point-to-point links, or Pseudonode LSP by the Designated Intermediate System (DIS) for multi-access LAN's, back toward the router and the link that originated this Reverse Metric TLV. Refer to "Elements of Procedure", in Section 3 for details on how an IS-IS router should process the Metric Offset field in a Reverse Metric TLV.

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 SHOULD add the received Metric Offset value in the Reverse Metric TLV to each node's existing "default metric" in the Pseudonode LSP. If the "Whole LAN" bit is not set (0), then a DIS SHOULD add the received Metric Offset value in the Reverse Metric TLV to the existing "default metric" in the Pseudonode LSP for the single node from whom the Reverse Metric TLV was received. Please refer to "Multi-Access LAN Procedures", in Section 3.3, for additional details. The W bit MUST be unset when a Reverse Metric TLV is transmitted in a IIH PDU onto a point-to-point link to a neighbor, and the W bit MUST be ignored upon receiving on a point-to-point link.

The "sub-TLV Len" value is non-zero when an IS-IS router wishes to signal that its neighbor alter parameters contained in the neighbor's Traffic Engineering "Extended IS Reachability TLV", as defined in [RFC5305]. This document defines that only the "Traffic Engineering Default Metric" sub-TLV, sub-TLV Type 18, may be sent toward neighbors in the Reverse Metric TLV, because that is used in Constrained Shortest Path First (CSPF) computations. Upon receiving this TE sub-TLV in a Reverse Metric TLV, a node SHOULD add the received TE default metric to its existing, configured TE default metric within its Extended IS Reachability TLV. Use of other sub-TLV's is outside the scope of this document. The "sub-TLV Len" value

MUST be set to zero when an IS-IS router does not have TE sub-TLV's that it wishes to send to its IS-IS neighbor.

3. Elements of Procedure

3.1. Processing Changes to Default Metric

The Metric Offset field, in the Reverse Metric TLV, is a "default metric" that will either be in the range of 0 - 63 when a "narrow" IS-IS metric is used (IS Neighbors TLV, Pseudonode LSP) [RFC1195] or in the range of 0 - (2^24 - 2) when a "wide" Traffic Engineering metric value is used, (Extended IS Reachability TLV) [RFC5305]. It is RECOMMENDED that implementations, by default, place the appropriate maximum default metric value, 63 or (2^24 - 2), in the Metric Offset field and TE Default Metric sub-TLV of the Reverse Metric TLV, since the most common use is to indicate the link of the router is overloaded and to remove the link from the topology, except for use as a last-resort path.

In order to ensure that an individual TE link is used as a link of last resort during SPF computation, its metric MUST NOT be greater than or equal to (2^24 - 1) [RFC5305]. Therefore, a receiver of a Reverse Metric TLV MUST use the numerically smallest value of either the sum of its existing default metric and the Metric Offset value in the Reverse Metric TLV or (2^24 - 2), as the default metric when updating its Extended IS Reachability TLV and TE default-metric sub-TLV's that it will then flood throughout the IS-IS domain, using normal IS-IS procedures. Likewise, originators of a Pseudonode LSP or IS Neighbors TLV MUST use the numerically smallest value of either the sum of its existing default metric and the Metric Offset value it receives in a Reverse Metric TLV or 63 when updating the corresponding Pseudonode LSP or IS Neighbor TLV before they are flooded. This also applies when an IS-IS router is only configured or capable of sending a "narrow" IS-IS default metric, in the range of 0 - 63, but receives a "wide" Metric value in a Reverse Metric TLV, in the range of 64 - (2^24 - 2). In this case, the receiving router MUST use the maximum "narrow" IS-IS default metric, 63, as its IS-IS default metric value in its updated IS Neighbor TLV or Pseudonode LSP that it floods.

If an IS-IS router is configured to originate a TE Default Metric sub-TLV for a link, but receives a Reverse Metric TLV from its neighbor that does not contain a TE Default Metric sub-TLV, then the IS-IS router MUST add the value in the Metric Offset field of the Reverse Metric TLV to its own TE Default Metric sub-TLV for that link. The IS-IS router should then flood the updated Extended IS Reachability TLV, including its updated TE Default Metric sub-TLV, using normal IS-IS procedures.

Routers MUST scan the Metric Offset 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 Offset value or TE Default Metric sub-TLV value, the receiving router MUST update its advertised IS-IS default metric or Traffic Engineering parameters in the appropriate TLV's, recompute its SPF tree and flood new LSP's to other IS-IS routers.

3.2. Processing Changes to Default Metric for Multi-Topology IS-IS

The Reverse Metric TLV is applicable to Multi-Topology IS-IS (M-ISIS) [RFC5120] capable point-to-point links. If an IS-IS router is configured for M-ISIS it MUST send only a single Reverse Metric TLV in IIH PDU's toward its neighbor(s) on the designated link that is about to undergo maintenance. When an M-ISIS router receives a Reverse Metric TLV it MUST add the received Metric Offset value to its default metric in all Extended IS Reachability TLV's for all topologies. If an M-ISIS router receives a Reverse Metric TLV with a TE Default Metric sub-TLV, then the M-ISIS router MUST add the received TE Default Metric value to each of its TE Default Metric sub-TLV's in all of its MT Intermediate Systems TLV's. If an M-ISIS router is configured to advertise TE Default Metric sub-TLV's for one or more topologies, but does not receive a TE Default Metric sub-TLV in a Reverse Metric TLV, then the M-ISIS router MUST add the value in Metric Offset field of the Reverse Metric TLV to each of the TE Default Metric sub-TLV's for all topologies. The M-ISIS should flood its newly updated MT IS TLV's and recompute its SPF/CSPF accordingly.

Multi-Topology IS-IS [RFC5120] specifies there is no change to construction of the Pseudonode LSP, regardless of the Multi-Topology capabilities of a multi-access LAN. If any MT capable node on the LAN advertises the Reverse Metric TLV to the DIS, the DIS should act according to the "Multi-Access LAN Procedures" in Section 3.3 to update, as appropriate, the default metric contained in the Pseudonode LSP. If the DIS updates the default metric in and floods a new Pseudonode LSP, those default metric values will be applied to all topologies during Multi-Topology SPF calculations.

3.3. Multi-Access LAN Procedures

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. The decision process of the routers on this LAN MUST follow the procedure in section 7.2.8.2 of [ISO10589], and use the "Two-way connectivity check" during the topology and route calculation.

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 optionally 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, e.g.: 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 add the Metric Offset value to the default metric 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 add their own Metric Offset value to their default metric 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 add the default metric contained in its Pseudonode LSP for the specific neighbor that sent the Reverse Metric TLV.

As long as at least one IS-IS node on the LAN sending the signal to DIS with the W bit set, the DIS would add the metric value in the Reverse Metric TLV to all neighbor adjacencies in the Pseudonode LSP, regardless if some of the nodes on the LAN send the Reverse Metric TLV without the W bit set. The DIS MUST use the metric of the highest source MAC address of the node sending the TLV with the W bit set. The DIS MUST use the metric value towards the nodes which explicitly send the Reverse Metric TLV.

Local provisioning on the DIS to adjust the default metric(s) contained in the Pseudonode LSP MUST take precedence over received Reverse Metric TLV's. For instance, local policy of the DIS may be provisioned to ignore the W bit signaling on a LAN.

3.4. Point-To-Point Link Procedures

On a point-to-point link, there is already a "configured" IS-IS interface metric to be applied over the link towards the IS-IS neighbor.

When IS-IS receives the IIH PDU with the "Reverse Metric" on a point-to-point link and if the local policy allows the supporting of "Reverse Metric", it MUST add the metric value in the "Metric" field of the TLV to the locally configured interface metric value to be the metric for this IS-IS adjacency. The metric MUST NOT exceed the maximum allowed value used in either "narrow" (63) or "wide" (2^24 - 2) metric mode.

3.5. LDP/IGP Synchronization on LAN's

As described in [RFC6138] when a new IS-IS node joins a broadcast network, it is unnecessary and sometimes even harmful to put IS-IS maximum link metric on all the nodes. [RFC6138] proposes a solution to have the new node not advertising the adjacency towards the pseudo-node when it is not in a "cut-edge" position.

With the introduction of Reverse Metric in this document, a simpler alternative solution to the above mentioned problem can be used. The Reverse Metric allows the new node on the LAN to have the inbound metric value to be the maximum and this puts the link of this new node in the last resort position without impacting the other IS-IS nodes on the same LAN.

Specifically, when IS-IS adjacencies are being established by the new node on the LAN, besides setting the maximum link metric value (2^24 - 2) on the interface of the LAN for the LDP IGP synchronization as described in [RFC5443], it SHOULD advertise the maximum metric offset value in the Reverse Metric TLV in its IIH PDU to the LAN. It SHOULD continue this advertisement until it completes all the LDP label binding exchanges with all the neighbors over this LAN, either by receiving the LDP End-of-LIB [RFC5919] for all the sessions or by exceeding the provisioned timeout value on the node.

3.6. Link Overload Attribute Bit

Not every TE tunnel is setup using IS-IS link metric or IS-IS link TE metric across the domain. Although the larger than normal link metric or TE metric can be one way to indicate to the PCE controller that the node on the other side of the link is trying to reduce the inbound traffic, but a more explicit way is to have the router set a bit in the "link-attribute" sub-TLV [RFC5029] to express this link is currently overloaded. How the controller or the source of the TE tunnel use the "link overload" information in altering the TE tunnel path is outside the scope of this document.

3.7. Operational Guidelines

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 towards it.

The use of Reverse Metric does not alter IS-IS metric parameters stored in a router's persistent provisioning database.

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.

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

4. Security Considerations

The enhancement in this document makes it possible for one IS-IS router to manipulate the IS-IS default metric or optionally Traffic Engineering parameters of adjacent IS-IS neighbors. Although IS-IS routers within a single Autonomous System nearly always reside under the control of a single administrative authority, it is highly RECOMMENDED that operators configure authentication of IS-IS PDU's to mitigate use of the Reverse Metric TLV as a potential attack vector, particularly on multi-access LAN's.

5. 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.

This document also request that IANA allocate from the link-attribute bit value for sub-TLV 19 of TLV 22. This new bit is referred to as the "Link Overload" bit.

6. Acknowledgments

The authors would like to thank Mike Shand, Dave Katz, Guan Deng, Ilya Varlashkin, Jay Chen, Les Ginsberg, Peter Ashwood-Smith, Uma Chunduri, Alexander Okonnikov, Jonathan Harrison, Dave Ward, Himanshu Shah, Wes George, Danny McPherson, Ed Crabbe, Russ White and Robert Razsuk for their contributions.

This document was produced using Marshall Rose's xml2rfc tool.

7. References

7.1. Normative References

[I-D.shen-isis-spine-leaf-ext]
Shen, N., Ginsberg, L., and S. Thyamagundalu, "IS-IS Routing for Spine-Leaf Topology", draft-shen-isis-spine-leaf-ext-03 (work in progress), March 2017.

[ISO10589]
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.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <http://www.rfc-editor.org/info/rfc1195>.

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

[RFC5029] Vasseur, JP. and S. Previdi, "Definition of an IS-IS Link Attribute Sub-TLV", RFC 5029, DOI 10.17487/RFC5029, September 2007, <http://www.rfc-editor.org/info/rfc5029>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <http://www.rfc-editor.org/info/rfc5120>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <http://www.rfc-editor.org/info/rfc5305>.

7.2. Informative References

[RFC5443] Jork, M., Atlas, A., and L. Fang, "LDP IGP Synchronization", RFC 5443, DOI 10.17487/RFC5443, March 2009, <http://www.rfc-editor.org/info/rfc5443>.

[RFC5919] Asati, R., Mohapatra, P., Chen, E., and B. Thomas, "Signaling LDP Label Advertisement Completion", RFC 5919, DOI 10.17487/RFC5919, August 2010, <http://www.rfc-editor.org/info/rfc5919>.

[RFC6138] Kini, S., Ed. and W. Lu, Ed., "LDP IGP Synchronization for Broadcast Networks", RFC 6138, DOI 10.17487/RFC6138, February 2011, <http://www.rfc-editor.org/info/rfc6138>.

Appendix A. Node Isolation Challenges

On rare occasions it is necessary for an operator to perform disruptive network maintenance on an entire IS-IS router node, i.e.: major software upgrades, power/cooling augments, etc. In these cases, an operator will set the IS-IS Overload Bit (OL-bit) within the Link State Protocol Data Units (LSP's) of the IS-IS router about to undergo maintenance. The IS-IS router immediately floods the updated LSP's to all IS-IS routers throughout the IS-IS domain. Upon receipt of the updated LSP's, all IS-IS routers recalculate their Shortest Path First (SPF) tree excluding IS-IS routers whose LSP's have the OL-bit set. This effectively removes the IS-IS router about to undergo maintenance from the topology, thus preventing it from forwarding any transit traffic during the maintenance period.

After the maintenance activity is completed, the operator resets the IS-IS Overload Bit within the LSP's of the original IS-IS router causing it to flood updated IS-IS LSP's throughout the IS-IS domain. All IS-IS routers recalculate their SPF tree and now include the original IS-IS router in their topology calculations, allowing it to be used for transit traffic again.

Isolating an entire IS-IS router from the topology can be especially disruptive due to the displacement of a large volume of traffic through an entire IS-IS router to other, sub-optimal paths, (i.e.: those with significantly larger delay). Thus, in the majority of network maintenance scenarios, where only a single link or LAN needs to be augmented to increase its physical capacity or is experiencing an intermittent failure, it is much more common and desirable to gracefully remove just the targeted link or LAN from service, temporarily, so that the least amount of user-data traffic is affected while intrusive augment, diagnostic and/or replacement procedures are being executed.

Appendix B. Link Isolation Challenges

Before network maintenance events are performed on individual physical links or LAN's, operators substantially increase the IS-IS metric simultaneously on both devices attached to the same link or LAN. 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 augmentation, fault diagnosis or repairs 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.

Appendix C. Contributors' Addresses

Tony Li

Email: tony.li@tony.li

Authors' Addresses

Naiming Shen
Cisco Systems
560 McCarthy Blvd.
Milpitas, CA 95035
USA

Email: naiming@cisco.com

Shane Amante
Apple, Inc.
1 Infinite Loop
Cupertino, CA 95014
USA

Email: samante@apple.com

Mikael Abrahamsson
T-Systems Nordic
Kistagangen 26
Stockholm
SE

Email: Mikael.Abrahamsson@t-systems.se