

Network work group
Internet Draft
Expires: March 2008
Category: Standards Track

Mach Chen
Renhai Zhang
Huawei Technologies Co.,Ltd
Xiaodong Duan
China Mobile
September 6, 2007

ISIS Traffic Engineering (ISIS-TE) Extensions in Support of Inter-AS
Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS)
Traffic Engineering
[draft-chen-ccamp-isis-interas-te-extension-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on March 6, 2008.

Abstract

This document describes extensions to the ISIS Traffic Engineering (ISIS-TE) mechanisms to support Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering (TE) for multiple Autonomous Systems (ASes). It defines ISIS-TE extensions for the flooding of TE information about inter-AS links which can be used to perform inter-AS TE path computation.

Internet-Draft ISIS extensions for Inter-AS TE September 2007

Conventions used in this document

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](#)].

Table of Contents

1.	Introduction.....	2
2.	Problem Statement.....	3
2.1.	A Note on Non-Objectives.....	3
2.2.	Per-Domain Path Determination.....	4
2.3.	Backward Recursive Path Computation.....	5
3.	Extensions to ISIS-TE.....	7
3.1.	Remote AS Number Sub-TLV.....	7
3.2.	Remote ASBR ID Sub-TLV.....	8
3.3.	Inter-AS Reachability TLV.....	9
4.	Procedure for Inter-AS TE Links.....	9
5.	Security Considerations.....	11
6.	IANA Considerations.....	11
6.1.	Inter-AS Reachability TLV.....	11
6.2.	Sub-TLVs for the Inter-AS Reachability TLV.....	11
6.3.	Sub-TLVs for the Extended IS Reachability TLV.....	12
7.	Acknowledgments.....	12
8.	References.....	12
8.1.	Normative References.....	12
8.2.	Informative References.....	12
	Authors' Addresses.....	14
	Intellectual Property Statement.....	14
	Disclaimer of Validity.....	15
	Copyright Statement.....	15

[1.](#) Introduction

[ISIS-TE] defines extensions to the ISIS protocol [[ISIS](#)] to support intra-area Traffic Engineering (TE). The extensions provide a way of encoding the TE information for TE-enabled links within the network (TE links) and flooding this information within an area. The Extended IS reachability TLV and Traffic Engineering Router ID TLV, which are defined in [[ISIS-TE](#)], are used to carry such TE information. The Extended IS reachability TLV has several nested sub-TLVs which describe the TE attributes for a TE link.

[ISIS-TE-V3] and [GMPLS-TE] define similar extensions to ISIS [ISIS] in support of IPv6 and GMPLS traffic engineering respectively.

Requirements for establishing Multiprotocol Label Switching (MPLS) TE Label Switched Paths (LSPs) that cross multiple Autonomous Systems (ASes) are described in [INTER-AS-TE-REQ]. As described in [INTER-AS-TE-REQ], a method SHOULD provide the ability to compute a path spanning multiple ASes. So a path computation entity that may be the head-end Label Switching Router (LSR), an AS Border Router (ASBR), or a Path Computation Element (PCE [PCE]) needs to know the TE information not only of the links within an AS, but also of the links that connect to other ASes.

In this document, some extensions to ISIS-TE are defined in support of carrying inter-AS TE link information for inter-AS Traffic Engineering. Two new sub-TLVs are added to the Extended IS reachability TLV, and a new TLV, which is referred to as inter-AS reachability TLV, is defined. The extensions are equally applicable to IPv4 and IPv6 as identical extensions to [ISIS-TE] and [ISIS-TE-V3]. The detailed definitions and procedures are discussed in the following sections.

[2.](#) Problem Statement

As described in [INTER-AS-TE-REQ], in the case of establishing an inter-AS TE LSP traversing multiple ASes, the Path message [RFC3209] may include the following elements in the Explicit Route Object (ERO) in order to describe the path of the LSP:

- a set of AS numbers as loose hops; and/or
- a set of LSRs including ASBRs as loose hops.

Two methods for determining inter-AS paths are currently discussed. The per-domain method [PD-PATH] determines the path one domain at a time. The backward recursive method [BRPC] uses cooperation between PCEs to determine an optimum inter-domain path. The sections that follow examine how inter-AS TE link information could be useful in both cases.

[2.1.](#) A Note on Non-Objectives


It is important to note that this document does not make any change to the confidentiality and scaling assumptions surrounding the use of ASes in the Internet. In particular, this document is conformant to the requirements set out in [\[INTER-AS-TE-REQ\]](#).

The following lists of features are explicit exclusions.

- o There is no attempt to distribute TE information from within one AS to another AS.
- o There is no mechanism proposed to distribute any form of TE reachability information for destinations outside the AS.
- o There is no proposed change to the PCE architecture or usage.
- o TE aggregation is not supported or recommended.
- o There is no exchange of private information between ASes.
- o No ISIS adjacencies are formed on the inter-AS link.

[2.2](#). Per-Domain Path Determination

In the per-domain method of determining an inter-AS path for an MPLS-TE LSP, when an LSR that is an entry-point to an AS receives a PATH message from an upstream AS with an ERO containing a next hop that is an AS number, it needs to find which LSRs (ASBRs) within the local AS are connected to the downstream AS so that it can compute a TE LSP segment across the AS to one of those LSRs and forward the PATH message to the LSR and hence into the next AS. See the figure below for an example:



<-- AS1 -->:<---- AS2 ---->:<--- AS3 --->

Figure 1: Inter-AS Reference Model

The figure shows three ASes (AS1, AS2, and AS3) and twelve LSRs (R1 through R12). R3 and R4 are ASBRs in AS1. R5, R6, R7, and R8 are ASBRs in AS2. R9 and R10 are ASBRs in AS3.

If an inter-AS TE LSP is planned to be established from R1 to R12, the AS sequence is limited as: AS1, AS2, AS3.

Suppose that the Path message enters AS2 from R3. The next hop in the ERO shows AS3, and R5 must determine a path segment across AS2 to reach AS3. It has a choice of three exit points from AS2 (R6, R7, and

R8) and it needs to know which of these provide TE connectivity to AS3, and whether the TE connectivity (for example, available bandwidth) is adequate for the requested LSP.

Alternatively, if the next hop in the ERO is the entry ASBR for AS3 (say R9), R5 needs to know which of its exit ASBRs has a TE link that connects to R9. Since there may be multiple exist ASBRs that are connected to R9 (both R7 and R8 in this example), R5 also needs to know the TE properties of the inter-AS TE links so that it can select the correct exit ASBR.

Once the path message reaches the exit ASBR, any choice of inter-AS TE link can be made by the ASBR if not already made by entry ASBR that computed the segment.

More details can be found in the Section 4.0 of [[PD-PATH](#)], which clearly points out why advertising of inter-AS links is desired.

To enable R5 to make the correct choice of exit ASBR the following information is needed:

- o List of all inter-AS TE links for the local AS.
- o TE properties of each inter-AS TE link.
- o AS number of the neighboring AS connected to by each inter-AS TE link.

- o Identity (TE Router ID) of the neighboring ASBR connected to by each inter-AS TE link.

In GMPLS networks further information may also be required to select the correct TE links as defined in [\[GMPLS-TE\]](#).


The example above shows how this information is needed at the entry point ASBRs for each AS (or the PCEs that provide computation services for the ASBRs), but this information is also needed throughout the local AS if path computation function is fully distributed among LSRs in the local AS, for example to support LSPs that have start points (ingress nodes) within the AS.

[2.3.](#) Backward Recursive Path Computation

Another scenario using PCE techniques has the same problem. [\[BRPC\]](#) defines a PCE-based TE LSP computation method (called Backward Recursive Path Computation) to compute optimal inter-domain constrained MPLS-TE or GMPLS LSPs. In this path computation method, a

specific set of traversed domains (ASes) are assumed to be selected before computation starts. Each downstream PCE in domain(i) returns to its upstream neighbor PCE in domain(i-1) a multipoint-to-point tree of potential paths. Each tree consists of the set of paths from all Boundary Nodes located in domain(i) to the destination where each path satisfies the set of required constraints for the TE LSP (bandwidth, affinities, etc.).

So a PCE needs to select Boundary Nodes (that is, ASBRs) that provide connectivity from the upstream AS. In order that the tree of paths provided by one PCE to its neighbor can be correlated, the identities of the ASBRs for each path need to be referenced, so the PCE must know the identities of the ASBRs in the remote AS reached by any inter-AS TE link, and, in order that it provides only suitable paths in the tree, the PCE must know the TE properties of the inter-AS TE links. See the following figure as an example:






Figure 2: BRPC for Inter-AS Reference Model

The figure shows three ASes (AS1, AS2, and AS3), three PCEs (PCE1, PCE2, and PCE3) and twelve LSRs (R1 through R12). R3 and R4 are ASBRs in AS1. R5, R6, R7, and R8 are ASBRs in AS2. R9 and R10 are ASBRs in AS3. PCE1, PCE2, and PCE3 cooperate to perform inter-AS path computation and are responsible for path segment computation within their own domains.

If an inter-AS TE LSP is planned to be established from R1 to R12, the traversed domains are assumed to be selected: AS1->AS2->AS3, and the PCE chain is: PCE1->PCE2->PCE3. First, the path computation request originated from the PCC (R1) is relayed by PCE1 and PCE2 along the PCE chain to PCE3, then PCE3 begins to compute the path segments from the entry boundary nodes that provide connection from AS2 to the destination (R12). But, to provide suitable path segments, PCE3 must determine which entry boundary nodes provide connectivity to its upstream neighbor AS (identified by its AS number), and must know the TE properties of the inter-AS TE links. In the same way,

PCE2 also needs to determine the entry boundary nodes according to its upstream neighbor AS and the inter-AS TE link capabilities.

Thus, to support Backward Recursive Path Computation the same information as listed in [Section 2.2](#) is required.

[3.](#) Extensions to ISIS-TE

Note that this document does not define mechanisms for distribution of TE information from one AS to another, does not distribute any form of TE reachability information for destinations outside the AS, does not change the PCE architecture or usage, does not suggest or recommend any form of TE aggregation, and does not feed private information between ASes. See [section 2.1](#).

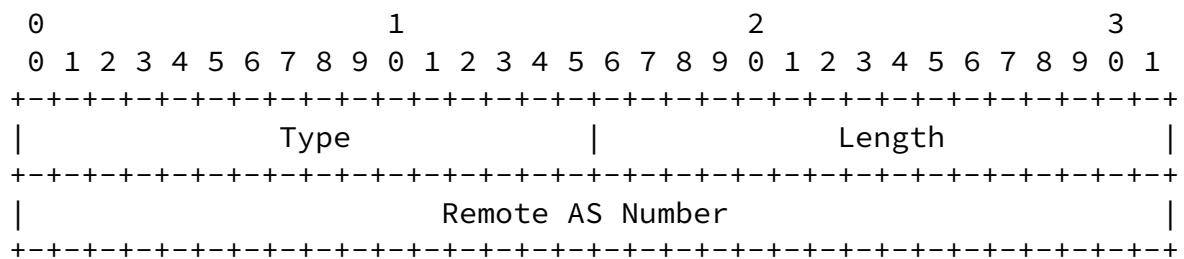
In this document, two new sub-TLVs are added to the extended IS reachability TLV to carry the information about the neighboring AS number and the remote ASBR ID of an inter-AS link. A new TLV, which

is referred to as inter-AS reachability TLV, is defined to flood the information about the neighboring AS and the remote ASBR ID within a whole AS.

3.1. Remote AS Number Sub-TLV

As described in [[ISIS-TE](#)], the Extended IS reachability TLV describes a single link and consists of a set of sub-TLVs. A new sub-TLV, the Remote AS Number sub-TLV is added to the extended IS reachability TLV when advertising inter-AS links. The Remote AS Number sub-TLV specifies the AS number of the neighboring AS to which the advertised link connects.

The Remote AS number sub-TLV is TLV type 23 (which needs to be confirmed by IANA), and is four octets in length. The format is as follows:

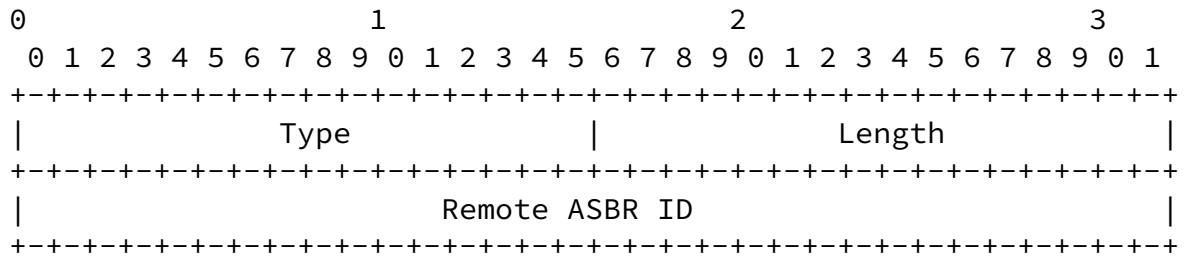


The Remote AS number field has 4 octets. When only two octets are used for the AS number, as in current deployments, the left (high-order) two octets MUST be set to zero. The Remote AS Number Sub-TLV MUST be included when a router advertises an inter-AS TE link.

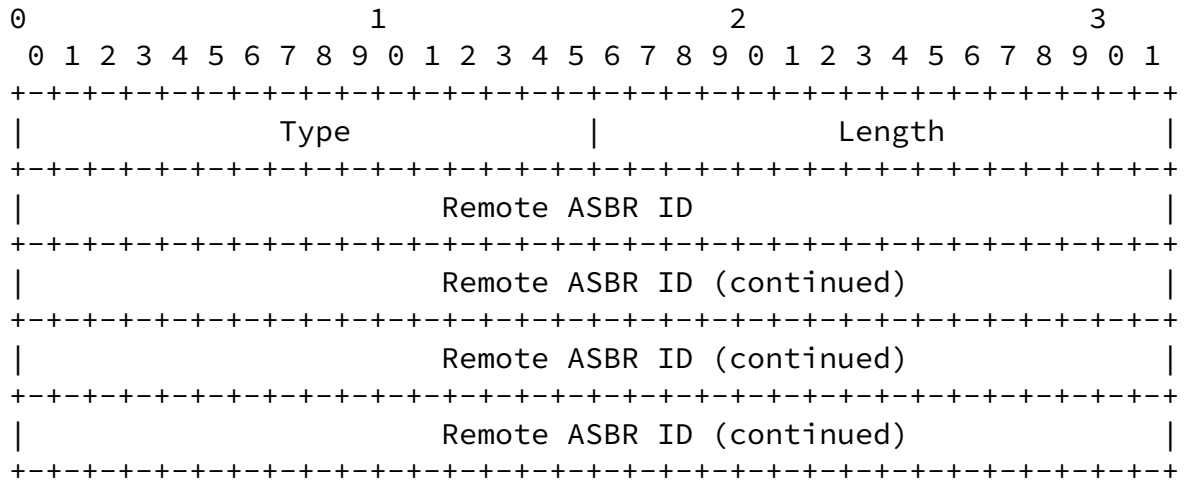
3.2. Remote ASBR ID Sub-TLV

A new sub-TLV, which is referred to as the Remote ASBR ID sub-TLV, is added to the extended IS reachability TLV when advertising inter-AS links. The remote ASBR ID sub-TLV specifies the identifier of the remote ASBR to which the advertised inter-AS link connects, which could be any stable and routable address of the remote ASBR (e.g., the Router ID, TE Router ID or interface address). The TE Router ID is RECOMMENDED.

The Remote ASBR ID sub-TLV is TLV type 24 (which needs to be confirmed by IANA), and is four or sixteen octets in length. The format of the remote ASBR ID sub-TLV is as follows:



or



If a router implements traffic engineering for IPv4, the length of the remote ASBR ID is 4. If a router implements traffic engineering for IPv6, the length of the remote ASBR ID is 16. The Remote ASBR ID Sub-TLV MUST be included when a router advertises an inter-AS TE link .

3.3. Inter-AS Reachability TLV

The inter-AS reachability TLV has type 141 (which needs to be confirmed by IANA). This is an optional TLV, when needed, it is used to flood the reachability information of the inter-AS links within a whole AS. And such reachability information SHOULD include the neighboring AS number and the remote ASBR ID to which an inter-AS link connects. The inter-AS reachability TLV contains a data

structure consisting of:

- 6 octets of System ID
- 1 octet of Pseudonode Number
- 1 octet flags
 - 1 bit of up/down information
 - 1 bit indicating the presence of sub-TLVs
 - 6 bits reserved
- 1 octet of length of sub-TLVs
- 0-246 octets of sub-TLVs
 - where each sub-TLV consists of a sequence of:
 - 1 octet of sub-type
 - 1 octet of length of the value field of the sub-TLV
 - 0-244 octets of value

In this document, two sub-TLVs are defined for the inter-AS Reachability TLV, they are:

Sub-TLV type	Length	Name
-----	-----	-----
23	4	Remote AS number
24	4or16	Remote ASBR Identifier

These two sub-TLVs have the same format and semantics as defined in [Section 3.1](#) and [section 3.2](#) of this memo.

4. Procedure for Inter-AS TE Links

When TE is enabled on an inter-AS link and the link is up, the ASBR SHOULD advertise this link using the normal procedures for ISIS-TE [[ISIS-TE](#)]. When either the link is down or TE is disabled on the link, the ASBR SHOULD withdraw the advertisement. When there are changes to the TE parameters for the link (for example, when the available bandwidth changes) the ASBR SHOULD re-advertise the link, but the ASBR MUST take precautions against excessive re-advertisements.

Helloes MUST NOT be exchanged (and consequently, an ISIS adjacency MUST NOT be formed) over the inter-AS link.

The information advertised comes from the ASBR's knowledge of the TE capabilities of the link, the ASBR's knowledge of the current status and usage of the link, and configuration at the ASBR of the remote AS number and remote ASBR TE Router ID.

When the inter-AS reachability information needs to reach all routers (including area border routers, ASBRs, and PCEs) in the AS, the ASBR SHOULD carry the Remote AS sub-TLV and Remote ASBR ID sub-TLV in the inter-AS reachability TLV. As defined in Section 4.1 of [[ISIS-TE](#)], the inter-AS reachability TLV also defines an up/down bit to facilitate the redistribution of inter-AS reachability information freely between level 1 and level 2. The semantics of the up/down bit in the new inter-AS reachability TLV are identical to the semantics of the up/down bit defined in [[ISIS-TE](#)]. That is, the up/down bit SHALL be set to 0 when the inter-AS reachability information first injected into ISIS [[ISIS](#)], and the up/down bit SHALL be set to 1 if the inter-AS reachability information needs to be advertised from high level to low level.

Legacy routers receiving an advertisement for an inter-AS TE link are able to ignore it because they do not know the new TLV and sub-TLVs that are defined in [Section 3](#) in this document. They will continue to flood the LSP, but will not attempt to use the information received as if the link were an intra-AS TE link.

Since there is no ISIS adjacency running on the inter-AS link, the local ASBR SHOULD do a "proxy" advertisement for the backward direction of an inter-AS TE link, which facilitates a path computation entity to do a 2-way check before including the link in a path computation. As the objective of such a "proxy" advertisement is to avoid using an inter-AS TE link when the backward direction of the inter-AS TE link is unavailable or unsuitable, only some mandatory or essential TE information needs to be advertised, i.e. the Link ID, the Link Type, and the Remote AS number of an inter-AS TE link.

Routers or PCEs that are capable of processing advertisements of inter-AS TE links SHOULD NOT use such links to compute paths that exit an AS to a remote ASBR and then immediately re-enter the AS through another TE link. Such paths would constitute extremely rare occurrences and SHOULD NOT be allowed except as the result of specific policy configurations at the router or PCE computing the path.

5. Security Considerations

The protocol extensions defined in this document are relatively minor and can be secured within the AS in which they are used by the existing ISIS security mechanisms.

There is no exchange of information between ASes, and no change to the ISIS security relationship between the ASes. In particular, since no ISIS adjacency is formed on the inter-AS links, there is no requirement for ISIS security between the ASes.

It should be noted, however, that some of the information included in these new advertisements (the remote AS number and the remote ASBR ID) are obtained from a neighboring administration and cannot be verified in anyway. Since the means of delivery of this information is likely to be part of a commercial relationship, the source of the information should be carefully checked before it is entered as configuration information at the ASBR responsible for advertising the inter-AS TE links.

6. IANA Considerations

IANA is requested to make the following allocations from registries under its control.

6.1. Inter-AS Reachability TLV

This document defines the following new ISIS TLV type that needs to be reflected in the ISIS TLV code-point registry as described in [Section 3.3](#):

Type	Description	IIH	LSP	SNP
-----	-----	---	---	---
141	Inter-AS reachability information	n	y	n

6.2. Sub-TLVs for the Inter-AS Reachability TLV

This document defines the following new sub-TLV types of top-level TLV 141 that need to be reflected in the ISIS sub-TLV registry for TLV 141 as described in Sections [3.3](#) and [6.1](#):

Internet-Draft ISIS extensions for Inter-AS TE September 2007

Type	Description	Length
----	-----	-----
23	Remote AS number	4
24	Remote ASBR Identifier	4 or 16

[6.3.](#) Sub-TLVs for the Extended IS Reachability TLV

This document also defines the following new sub-TLV types of top-level TLV 22 that need to be reflected in the ISIS sub-TLV registry for TLV 22 as described in [Section 3.1](#):

Type	Description	Length
----	-----	-----
23	Remote AS number	4
24	Remote ASBR Identifier	4 or 16

[7.](#) Acknowledgments

The authors would like to thank Adrian Farrel, Jean-Louis Le Roux, Christian Hopps, and Les Ginsberg for their review and comments to this document.

[8.](#) References

[8.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [ISIS] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), December 1990.

[8.2.](#) Informative References

- [INTER-AS-TE-REQ] Zhang and Vasseur, "MPLS Inter-AS Traffic

- [PD-PATH] Ayyangar, A., Vasseur, JP., and Zhang, R., "A Per-domain path computation method for establishing Inter-domain", [draft-ietf-ccamp-inter-domain-pd-path-comp](#), (work in progress).
- [BRPC] JP. Vasseur, Ed., R. Zhang, N. Bitar, JL. Le Roux, "A Backward Recursive PCE-based Computation (BRPC) procedure to compute shortest inter-domain Traffic Engineering Label Switched Paths ", [draft-ietf-pce-brpc](#), (work in progress)
- [PCE] Farrel, A., Vasseur, JP., and Ash, J., "A Path Computation Element (PCE)-Based Architecture", [RFC4655](#), August 2006.
- [ISIS-TE-V3] Harrison, J., Berger, J., and Bartlett, M., "IPv6 Traffic Engineering in IS-IS", [draft-ietf-isis-ipv6-te](#), {work in progress}.
- [ISIS-TE] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", [RFC 3784](#), June 2004.
- [GMPLS-TE] K.Kompella and Y.Rekhter, "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching", [RFC 4205](#), October 2005.

Internet-Draft

ISIS extensions for Inter-AS TE

September 2007

Authors' Addresses

Mach Chen
Huawei Technologies Co.,Ltd
KuiKe Building, No.9 Xinxu Rd.,
Hai-Dian District
Beijing, 100085
P.R. China

Email: mach@huawei.com

Renhai Zhang
Huawei Technologies Co.,Ltd
KuiKe Building, No.9 Xinxu Rd.,
Hai-Dian District
Beijing, 100085
P.R. China

Email: zhangrenhai@huawei.com

Duanxiao Dong
China Mobile
53A,Xibianmennei Ave,Xunwu District
Beijing, China

Email: duanxiaodong@chinamobile.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to

pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

Mach, Renhai & Duan

Expires March 6, 2008

[Page 14]

Internet-Draft

ISIS extensions for Inter-AS TE

September 2007

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

