

Network Working Group
Internet Draft
Category: Standards Track

D.Ceccarelli
D.Caviglia
Ericsson
Fatai Zhang
Dan Li
Huawei

Expires: April 2010

October 16, 2009

**Traffic Engineering Extensions to OSPF for Generalized MPLS (GMPLS)
Control of Evolutive G.709 OTN Networks**

[draft-ceccarelli-ccamp-gmpls-ospf-g709-00.txt](http://www.ietf.org/drafts/ceccarelli-ccamp-gmpls-ospf-g709-00.txt)

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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/1id-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 24, 2010.

Abstract

This document describes OSPF routing protocol extensions to support the evolutive Optical Transport Networks (OTN) under the control of Generalized MPLS (GMPLS).

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 [\[RFC2119\]](#).

Table of Contents

1.	Introduction.....	2
2.	Terminology.....	3
3.	Overview of the Evolutive G.709.....	3
4.	G.709 Digital Layer TE Information.....	5
4.1.	Tributary Slot type.....	5
4.2.	TE link type.....	6
4.3.	LO ODU signal type.....	6
4.4.	TE link Total Bandwidth.....	7
4.5.	TE link Unreserved Bandwidth.....	7
4.6.	Maximum LSP Bandwidth.....	7
5.	OSPF extensions.....	8
5.1.	Extensions to Interface Switching Capability Descriptor..	8
6.	Compatibility Considerations.....	11
7.	Example.....	11
8.	Security Considerations.....	13
9.	IANA Considerations.....	13
10.	Acknowledgments.....	13
11.	References.....	13
11.1.	Normative References.....	13
11.2.	Informative References.....	14
12.	Authors' Addresses.....	14
13.	Contributors.....	15

[1.](#) Introduction

Per OSPF, an Opaque LSA (Link State Advertisements) carrying application-specific information can be generated and advertised to other nodes following the flooding mechanism defined in [\[RFC 2370\]](#). Three types of opaque LSA are defined, i.e. type 9 - link-local flooding scope, type 10 - area-local flooding scope, type 11 - AS flooding scope.

Traffic Engineering (TE) LSA using type 10 of opaque LSA is defined in [\[RFC 3630\]](#) for TE purpose. This type of LSA is composed of a standard LSA header and a payload including one top-level TLV (Type/Length/Value triplet) and possible several nested sub-TLVs.

RFC [3630] defines two top-level TLVs: Router Address TLV and Link TLV; and nine possible sub-TLVs for the Link TLV, used to carry link related TE information.

The Link type sub-TLVs are enhanced by [RFC 4203] in order to support GMPLS networks and related specific link information.

In GMPLS networks, each node generates TE LSAs to advertise its TE information and capabilities (link-specific, or node-specific), through the network. The TE information carried in the LSAs are collected by the other nodes of the networks and stored into their local Traffic Engineering Databases (TED).

In the GMPLS based G.709 Optical Transport Networks (OTNs), in order to automatically establish ODUK connections through GMPLS RSVP-TE signaling, routing is the foundation.

OTN networks provide flexible and various multiplexing relationships (e.g., ODU_j multiplexed into ODU_k ($j < k$)), two different tributary slots for ODU_k ($k=1, 2, 3$) and ODUflex signal type, which is being standardized in ITU-T. In order to present this information in the routing process, the OSPF protocol needs to be extended.

This document describes OSPF routing protocol extensions to support the evolutive OTNs under the control of GMPLS.

Note that the routing information for Optical Channel Layer (OCh) (i.e., wavelength) is out of the scope. Please refer to [WSON-Frame] for further information.

2. Terminology

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].

3. Overview of the Evolutive G.709

The traditional OTN specification [G709] describes the Optical Transport Hierarchy (OTH) and introduces three types of ODU (Optical Channel Data Unit) signal (i.e., ODU₁, ODU₂ and ODU₃). The ODU_j can be mapped into one or more Tributary Slots (with granularity of 2.5Gbps) of OPU_k (Optical Channel Payload Unit- k) where $j < k$. The ODU_j can also be mapped into OTU_j (Optical Channel Transport Unit- j , $j=1, 2$ or 3) directly.

Recent revisions of ITU-T Recommendation G.709 have introduced new features for OTNs: ODU0, ODU4, ODU2e, ODU3e1, ODU3e2 and ODUflex. The new features for the evolutive OTNs are described in separate ITU-T documents. ODU0, ODU2e and ODU4 ODUflex are described in [G709-V3]. ODU3e1 and ODU3e2 are described in [Gsup43].

The ITU-T documents also define the new multiplexing hierarchy for the evolutive OTN. In this multiplexing hierarchy, LO (Lower Order) ODU_j can be mapped into an OTU_j, or multiplexed into an HO (Higher Order) ODU_k (where $j < k$) occupying several TSs (Tributary Slots).

In the case of LO ODU_j mapping into OTU_j, the following mappings are defined:

- o ODU1 into OTU1 mapping
- o ODU2 into OTU2 mapping
- o ODU3 into OTU3 mapping
- o ODU4 into OTU4 mapping
- o ODU2e into OTU2e mapping

In the case of LO ODU_j multiplexing into HO ODU_k, a new Tributary Slot granularity (i.e., 1.25Gbps) is introduced in [G709-V3]. For the evolutive OTN, the multiplexing of ODU_j ($j = 0, 1, 2, 2e, 3, flex$) into an ODU_k ($k > j$) signal can be depicted as follows:

- ODU0 into ODU1 multiplexing (with 1,25Gbps TS granularity)
- ODU0, ODU1, ODUflex into ODU2 multiplexing (with 1.25Gbps TS granularity)
- ODU1 into ODU2 multiplexing (with 2.5Gbps TS granularity)
- ODU0, ODU1, ODU2, ODU2e and ODUflex into ODU3 multiplexing (with 1.25Gbps TS granularity)
- ODU1, ODU2 into ODU3 multiplexing (with 2.5Gbps TS granularity)
- ODU0, ODU1, ODU2, ODU2e, ODU3 and ODUflex into ODU4 multiplexing (with 1.25Gbps TS granularity)
- ODU2e into ODU3e1 multiplexing (with 2.5Gbps TS granularity)
- ODU2e into ODU3e2 multiplexing (with 1.25Gbps TS granularity)

In order to be backward compatible with the 2.5Gbps TS defined in [G709-V3], both the 2.5Gbps TS and the 1.25Gbps TS can be used in the two cases listed below:

- o ODU1 into ODU2 multiplexing
- o ODU1 and ODU2 into ODU3 multiplexing

From the link perspective, it can only work under one TS type. For example, if the both ends (or interfaces) of the link can support 2.5Gbps TS and 1.25Gbps TS, then it will work under 2.5Gbps TS or 1.25Gbps TS. If one end can support 1.25Gbps TS, and another end can support 2.5Gbps TS, the end with 1.25Gbps TS MUST adopt a 2.5Gbps TS).

4. G.709 Digital Layer TE Information

This document only considers the TE information needed for LO ODU path computation. WSON TE information is out of scope. Please refer to [WSON-OSPF] for more information about WSON routing information.

From the perspective of [G709-V3], there are two types of LO ODU:

(1) A LO ODU_k mapped into an OTU_k. In this case, the server layer of this LO ODU is an OTU_k. For example, if a STM-16 signal is encapsulated into ODU₁, and then ODU₁ is mapped into OTU₁, the ODU₁ is LO ODU.

(2) A LO ODU_j multiplexed into a HO (Higher Order) ODU_k ($j < k$) occupying several TSs. In this case, the server layer of this LO ODU is a HO ODU_k. For example, if ODU₁ is multiplexed into ODU₂, and ODU₂ is mapped into OTU₂, the ODU₁ is LO ODU and ODU₂ is HO ODU.

In order to compute a suitable path the PCE (centralized or distributed) need a set of data that should be advertised by the routing protocol. In the following sections each type of data is listed and analyzed, while the possible values are shown in [section 5](#).

4.1. Tributary Slot type

There are two types of TS defined in the ITU-T recommendations, and from the link perspective, it can only work under one TS type. For example, if the both ends (or interfaces) of a link can support 2.5Gbps TS or 1.25Gbps TS, then it will work under 2.5Gbps TS or 1.25Gbps TS. If one end can support 1.25Gbps TS, and another end can support 2.5Gbps TS, the end with 1.25Gbps TS should adopt the 2.5Gbps TS).

In addition, the bandwidth accounting depends on the type of TS. Therefore, the type of the TS should be known during LO ODU path computation.

4.2. TE link type

The link type indicates the OTUk/H0 ODUk type of the TE link.

The TS bandwidth of different types of OTUk is different. The TS bandwidth in an OTUk is increasing along with the increasing of k (see [G709-V3]). The bandwidth of a TS in a TE link can be deduced from the TS type and link type of the TE link. For example, the bandwidth of a 1.25G TS without NJO (Negative Justification Opportunity) in an OTU2 is about 1.249409620 Gbps, while the bandwidth of a 1.25G TS without NJO in an OTU3 is about 1.254703729 Gbps.

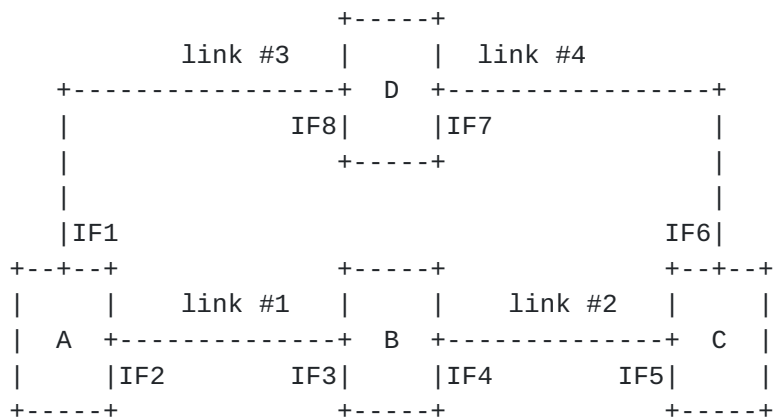
The actual TS bandwidth of a TE link is useful to determine the TS number needed by an ODUFlex service. And the actual TS bandwidth of a TE link can be deduced by the TE link type and TS type.

In the case of link bundling, only the component links with the same TS type and TE link type can be bundled together.

4.3. LO ODU signal type

It is possible that some equipment can not always support all the LO ODU signal types. When a path computation procedure for a LO ODU is performed, it needs to check whether a link has the capability to carry a specific type of LO ODU or not. If a link can not carry this type of LO ODU, it should be excluded during the path computation. Only the links with the capability of carrying this type of LO ODU can be the candidates.

For example, in the following figure, the interfaces IF1, IF2, IF8, IF7, IF5 and IF6 can support ODUFlex signals, while the interfaces IF3 and IF4 can not support ODUFlex signals. In this case, if one ODUFlex connection from A to C is requested, link #1 and #2 are excluded, link #3 and link #4 are the candidates (the possible path could be A-D-C through link #3 and link #4).



Therefore, it is necessary to advertise the LO ODU types that the OTU or HO ODU TE link can support.

4.4. TE link Total Bandwidth

For an unbundled link, the Total Bandwidth is dependent on the type of the OTU or HO ODU, while in case of link bundling, the Total Bandwidth is the sum of the total bandwidth of all the physical links composing the TE link.

For a specific OTUK or HO ODUK TE link, the Total Bandwidth can be deduced through the total number of the Tributary Slots. For example, an unbundled OTU4 link has 80 Tributary Slots with 1.25Gbps. If two component OTU4 links are bundled, the TE link has 160 Tributary Slots with 1.25Gbps.

4.5. TE link Unreserved Bandwidth

In the GMPLS based OTN networks, the Unreserved Bandwidth of a TE link is the sum of the unreserved bandwidths of all the component links associated with the bundled link.

The unreserved bandwidth can be accounted through the unallocated Tributary Slots of the TE link.

4.6. Maximum LSP Bandwidth

The Maximum Bandwidth that a LSP can occupy in a TE link is determined by the component link with the maximum unreserved bandwidth in this TE link.

For example, if two OTU3 component links are bundled to a TE link, the unreserved bandwidth of the first component link is $20 \times 1.25G$ TSs, and the unreserved bandwidth of the second component link is $24 \times 1.25G$

TSs. Then the unreserved bandwidth of this TE link is $44 \times 1.25G$ TSs, but the maximum TSs that a LSP can occupy in this TE link is 24, not 44.

5. OSPF extensions

In terms of GMPLS based OTN networks, each OTUk/H0 ODUk can be viewed as a component link, and each component link can carry one or more types of LO ODU. One or more component links with the same attributes can be bundled as a TE link (see [RFC 4201] for link bundling).

Each TE LSA can carry a top-level TLV with several nested sub-TLVs to describe different attributes of a TE link. Two top-level TLVs are defined in [RFC 3630]. (1) The Router Address TLV (referred to as the Node TLV) and (2) the TE link TLV. One or more sub-TLVs can be nested into the two top-level TLVs. The sub-TLV set for the two top-level TLVs are also defined in [RFC 3630] and [RFC 4203].

A general Interface Switching Capability Descriptor (ISCD) sub-TLV is defined in [RFC 4203]. The bandwidth accounting is encoded in a 4 octets field in the IEEE floating point format. Max LSP Bandwidth is accounted at each priority X (0~7).

To keep the minor changes on the existing routing protocol, the routing procedures and the TLVs defined in the existing standard documents, such as [RFC 3630] and [RFC 4203] should be reused as possible.

Therefore, OTN specific routing information can be simply carried in the "Switching Capability-specific information" of ISCD defined in [RFC 4203].

The routing procedures are the same as [RFC 3630] and [RFC 4203].

5.1. Extensions to Interface Switching Capability Descriptor

The Interface Switching Capability Descriptor is defined as follows.

A new type of Switching Cap is introduced to distinguish the evolutive OTN and the existing TDM networks (e.g., SDH networks). The new Switching Cap type is defined in the following:

When the Switching Capability field is evolved OTN, the Max LSP Bandwidth at priority x is counted by TS number.

[Page 9]

[illegible]

The type of this sub-TLV is TBD. The length of this TLV is eight octets.

- o T bits (2 bits): Indicates the type of the Tributary Slot of this TE link, value 0 means the TS type is 1.25Gbps, value 1 means the TS type is 2.5Gbps.
- o OD(T)Uk (4 bits): Indicates the type of the TE link, i.e. the server layer signal that the LO ODUs can be mapped or multiplexed into. The following values are defined:

0: Reserved (for future use)

```
1: OTU1/H0 ODU1
```

2: OTU2/H0 ODU2

3: OTU3/H0 ODU3

4: OTU4/H0 ODU4

5: OTU2e/H0 ODU2e

6: H0 ODU3e1

7: H0 ODU3e2

8-15: Reserved (for future use)

Note that an ODU3e1/ODU3e2 must carry several LO ODUs, non-OTN signals can not be mapped into ODU3e1/ODU3e2 directly. So ODU3e1/ODU3e2 can not be a LO ODU, and OTU3e1/OTU3e2 can not be a server layer for LO ODU.

The bandwidth of a TS in this TE link can be deduced from T bit and OD(T)Uk field.

- o Signal Flags (16 bits): This field indicates the LO ODU type supported by the TE link. A flag set to 1 indicates that the TE link supports the corresponding LO ODU signal. Currently the following flags are defined:

Flag A: indicates whether LO ODU0 is supported.

Flag B: indicates whether LO ODU1 is supported.

Flag C: indicates whether LO ODU2 is supported.

Flag D: indicates whether LO ODU3 is supported.

Flag E: indicates whether LO ODU4 is supported.

Flag F: indicates whether LO ODU2e is supported.

Flag G: indicates whether LO ODUFlex is supported.

Other bits are reserved and must be set to zero when sent and should be ignored when received.

- o Total TS (12 bits): Indicates the total TS number of the TE link.
- o Unreserved TS (12 bits): Indicates the number unreserved TS in the TE link.

When a node receives a LSA with Switching Capability type 110, the Maximum Bandwidth sub-TLV and Unreserved Bandwidth sub-TLV in this LSA SHOULD be ignored if they exist, the bandwidth information MUST be deduced from ISCD sub-TLV (i.e., Total TS, Unreserved TS).

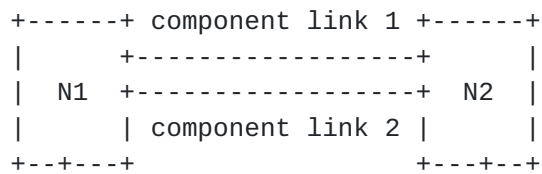
All the reserved fields must be set to zero and should be ignored when received.

6. Compatibility Considerations

The legacy nodes that do not implement the extensions defined in this document are able to ignore the LSA containing an ISCD sub-TLV with the Switching Cap type 110, because it is an unknown value. They will continue to flood the LSA to other neighbors, but will not use the information carried in this LSA.

7. Example

Based on the sub-TLVs defined in [[RFC 3630](#)], [[RFC 4203](#)] and this document, a G.709 digital TE link can be described as follows.



This picture shows a simple example of an OTN network. Both the link type of the two component links are OTU3 and both of the two component links have the capability of carrying ODU0, ODU1, ODU2 and ODUFlex client signals. The TS type is 1.25Gbps, the number of the total Tributary Slots is 32.

The two component links have the same nature, so the two component links can be bundled as a TE link. It is also feasible that each component link can be regarded as a TE link separately.

If the two component links are bundled together, N1 and N2 should assign a link local ID to the TE link. Then N1 can get the link remote ID automatically or manually.

N1 can generate an LSA to describe the above attributes of the TE link. Suppose the link IDs are unnumbered, the LSA should carry a link TLV with the following nested minimal sub-TLVs:

```
< G.709 Digital Link > ::= < Link Type > < Link ID > < Link
Local/Remote Identifiers > < Interface Switching Capability
Descriptor >
```

- o Link Type sub-TLV: Defined in [[RFC 3630](#)], G.709 digital links are always type 1 - Point-to-point link.
- o Link ID sub-TLV: Defined in [[RFC 3630](#)], for point-to-point link, indicates the remote router ID.
- o Link Local/Remote Identifiers sub-TLV: Defined in [[RFC 4203](#)], indicates the local link ID and the remote link ID.
- o Interface Switching Capability Descriptor sub-TLV: Defined in this document, carries the characteristic of this G.709 digital TE link.

Suppose the unreserved TS number of component link 1 is 20, and the unreserved TS number of component link 2 is 28. Also Suppose the Max LSP Bandwidth at priority 0 of component link 1 is 10 TSs, and the Max LSP Bandwidth at priority 0 of component link 2 is 18 TSs.

The ISCD sub-TLV in this LSA will carry the above attributes, i.e. the value of T bits is 1, OD(T)Uk bits is 3, total TS field is 64, Unreserved TS field is 48, Max LSP Bandwidth at priority 0 field is 18 TSs, and A/B/C/G bits are set to indicate the types of LO ODU that this TE link can carry.

When another node receives this LSA, it can determine that this TE link can carry ODU0, ODU1, ODU2, ODUFlex signals. The TS type is 1.25G, and the real bandwidth of TS is about 1.254703729 Gbps.

An ODUFlex signal at priority 0 can occupy 18 TS at most in this TE link. If an ODU0 (priority 0) path computation is performed, this TE link is a candidate. If a 40G ODUFlex (priority 0) path computation is performed, this TE link should be excluded, because this ODUFlex signal needs more than 18 TSs.

8. Security Considerations

TBD.

9. IANA Considerations

TBD.

10. Acknowledgments

TBD.

11. References

11.1. Normative References

- [RFC 2370] R. Coltun, "The OSPF Opaque LSA Option", [RFC 2370](#), July 1998.
- [RFC 3630] D. Katz, K. Kompella, D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.
- [RFC 4201] K. Kompella, Y. Rekhter, L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", [RFC 4201](#), October 2005.
- [RFC 4203] K. Kompella, Y. Rekhter, "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4203](#), October 2005.

[WSN-Frame] G. Bernstein, Y. Lee, W. Imajuku, "Framework for GMPLS and PCE Control of Wavelength Switched Optical Networks", work in progress: [draft-ietf-ccamp-rwa-WSN-Framework-02.txt](#), March 2009.

[WSN-OSPF] Fatai Zhang, Y. Lee, etc., " OSPF Extensions in Support of Routing and Wavelength Assignment (RWA) in Wavelength Switched Optical Networks (WSNs) ", work in progress: [draft-zhang-ccamp-rwa-wson-routing-ospf-01.txt](#), July 2009.

11.2. Informative References

[Gsup43] ITU-T, "Proposed revision of G.sup43 (for agreement)", December 2008.

[G709] ITU-T G.709 Recommendation, Interface for the Optical Transport Network (OTN) ITU-T, March 2003.

[G709-v3] ITU-T, "Draft revised G.709, version 3", consented by ITU-T on Oct 2009.

12. Authors' Addresses

Daniele Ceccarelli
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: daniele.ceccarelli@ericsson.com

Diego Caviglia
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: diego.caviglia@ericsson.com

Fatai Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base

Bantian, Longgang District
Shenzhen 518129 P.R.China
Phone: +86-755-28972912
Email: zhangfatai@huawei.com

Dan Li
Huawei Technologies Co., Ltd.
F3-5-B R&D Center, Huawei Base,
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972910
Email: danli@huawei.com

13. Contributors

Xiaobing Zi
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China
Phone: +86-755-28973229
Email: zixiaobing@huawei.com

Intellectual Property

The IETF Trust 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 any IETF 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.

Copies of Intellectual Property 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>

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 any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions.

For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of [RFC 5378](#). No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under [RFC 5378](#), shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

Disclaimer of Validity

All IETF Documents and the information contained therein 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 THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Full Copyright Statement

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

