

Network Working Group
Internet Draft
Intended status: Standards Track
Expires: September 2009

Sami Boutros (Ed.)
Siva Sivabalan (Ed.)
George Swallow
David Ward
Stewart Bryant
Cisco Systems, Inc.

March 9, 2009

Connection verification for MPLS Transport Profile LSP
draft-boutros-mpls-tp-cv-01.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/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 September 9, 2009.

Abstract

This document specifies method for verifying the connection of an MPLS Transport Profile(MPLS-TP) Label Switched Path (LSP) for management purpose. The proposed extension is based on MPLS Operation, Administration, and Maintenance (OAM). The goal is to verify that an MPLS-TP is properly setup in both control and data

planes, as well as to record the identities of all the LSRs along the path of MPLS-TP LSP.

Table of Contents

1.	Introduction.....	2
2.	Terminology.....	4
3.	MPLS-TP Connection Verification Mechanism.....	5
4.	MPLS-OAM Connection Verification Message.....	5
4.1.	In-band Message Identification.....	5
4.2.	Out-of-band Message Identification.....	6
4.3.	MPLS-TP CV Message Format.....	6
4.4.	MPLS-TP Connection Verification Record Route TLV.....	10
4.5.	Network Management System.....	10
5.	Operation.....	10
6.	Security Considerations.....	12
7.	IANA Considerations.....	12
8.	References.....	13
8.1.	Normative References.....	13
8.2.	Informative References.....	13
	Author's Addresses.....	14
	Full Copyright Statement.....	15
	Intellectual Property Statement.....	15

[1.](#) Introduction

In traditional transport networks, circuits are provisioned on multiple switches. Service Providers (SP) need to verify that the circuits are provisioned correctly in both control and data plane for management purpose. MPLS-TP bidirectional LSPs emulating traditional transport circuits need to provide the same connection verification capability. In this document, an MPLS-TP LSP as defined in [\[5\]](#) is based on the MPLS-TE, pseudowire (PW) or Multisegment PW [\[8\]](#).

An MPLS-OAM Connection Verification (CV) message originates at a Maintenance End Point (MEP) but can be directed to by any Maintenance Intermediate Point (MIP) along the path of MPLS-TP LSP as well as the other MEP. Therefore, the proposed mechanism addresses the verification of the full or partial path of an MPLS-TP LSP.

An MPLS-OAM CV message is intercepted at any MIP based on MPLS TTL expiry, and at MEP simply because it is the end of the LSP (i.e., regardless the value of the TTL).

In response to the MPLS-OAM CV request, each LSR along the path of the MPLS-TP LSP appends its ID using a newly defined TLV called Record Route TLV.

A Record Route TLV appended by a given LSR contains:

- . The LSR address which is represented by the format defined in [6].
- . Local Labels allocated by the LSR for both directions of the MPLS-TP LSP.

To describe the connection verification functionality, let us assume an MPLS-TP LSP between LSR-1 and LSR-5 passing through LSR-2, LSR-3, and LSR-4. Thus, LSR-1 and LSR-5 are MEPs whereas LSR-2, LSR-3, and LSR-4 are MIPs. The objective is to verify (both in control and data planes) the MPLS-TP LSP from LSR-1 to LSR-5 (end-to-end), and record all the IDs of the LSRs along the path. This could be accomplished using a conventional traceroute operation in which LSR-1 interrogates each LSR-2 through LSR-5 (using appropriate TTL value) in turn using a new message and response, and compiles the result. This approach requires 8 messages; a request and a response between LSR-1 and each of the other LSRs. On the other hand, the mechanism that we describe below can accomplish the goal with only 5 messages.

It is possible that the path of an MPLS-TP LSP contains loop(s) due to misconfiguration. Such mistakes are possible with manual configuration. For example, assume that MPLS-TP LSP under discussion is misconfigured such LSR-4 connects to LSR-2 instead of LSR-5. This results in a loop. In this case, the MPLS-OAM CV packets self limit when the MTU is reached, and when it happens, it is good practice to silently drop those packets.

If a MIP does not understand the MPLS-OAM CV message, it must silently drop the packet. To trap this condition as well as to trap the looping condition, an ingress MEP that initiates connection verification starts a timer when it sends an MPLS-OAM CV message. If the timer expires before the response arrives, the MEP assumes one of the following conditions:

- . the MPLS-TP LSP is incomplete.

- . an LSR (either MIP or MEP) does not understand the MPLS-OAM CV message.
- . there is a loop.

The ingress MEP then examines the MPLS-TP LSP by using the classic one-hop at a time, direct response traceroute.

In case not all hops on the path of the MPLS-TP LSP are MIPs, an ingress MEP can send conventional trace route with incrementing TTL 1, 2, 3,, to all MIPs and to the egress MEP along the path. Some of those requests will be sent to non MIP/MEP LSRs and will be dropped silently. When the MIPs and egress MEP receive the request, they will respond with an MPLS-OAM CV response message. The TTL value of the response SHOULD be large to ensure the response message reaches the ingress MEP without being intercepted at any MIP. Optionally, the TTL value of the response MAY be set to 1 so that each MIP can verify its ID included in the response message as the response travels towards the ingress MEP.

The proposed mechanism is based on a set of new TLVs which can be transported using one of the following methods:

1. Using in-band MPLS Connection Verification (CV) messages which are forwarded as MPLS packets (Non-IP routing and forwarding based).
2. Using in-band LSP-Ping extensions defined in [2] where IP/UDP packets are used (IP-Based routing and forwarding). The LSP-Ping messages may be sent in-band using the codepoint defined in [3].

Method (1) and (2) are referred to as "in-band option" and "LSP-Ping option" respectively in the rest of the document.

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](#) Error!
Reference source not found.

[2.](#) Terminology

ACH: Associated Channel Header

CV: Connection Verification

GAL: Generalized Alert Label

Boutros

Expires September 9, 2009

[Page 4]

Internet-Draft

[draft-boutros-mpls-tp-cv-01.txt](#)

March 2009

LSR: Label Switching Router

MEP: Maintenance End Point

MIP: Maintenance Intermediate Point

MPLS-OAM: MPLS Operations, Administration and Maintenance

MPLS-TP: MPLS Transport Profile

MPLS-TP LSP: Bidirectional Label Switch Path representing a circuit

MS-PW: Mult-Segment PseudoWire

NMS: Network Management System

PW: PseudoWire

RR: Record Route

TLV: Type Length Value

TTL: Time To Live

[3.](#) MPLS-TP Connection Verification Mechanism

For the in-band option, the proposed mechanism uses a new code point in the Generic Associated Channel Header (G-ACH) described in [\[7\]](#). The LSP-Ping option will be in compliance to specifications [\[2\]](#) and [\[3\]](#).

Moreover, the proposed mechanism requires Record Route TLV (defined in this document). Also, Authentication TLV defined in [4] is also required for this mechanism.

4. MPLS-OAM Connection Verification Message

4.1. In-band Message Identification

In the in-band option, under MPLS label stack of the MPLS-TP LSP, the ACH with "MPLS-TP Connection Verification (CV)" code point indicates that the message is an MPLS-TP CV message.

Boutros

Expires September 9, 2009

[Page 5]

Internet-Draft

[draft-boutros-mpls-tp-cv-01.txt](#)

March 2009

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 0 0 1|Version|      Flags      | 0xHH MPLS-TP CV Code Point |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 1: ACH Indication of MPLS-TP Connection Verification

The first nibble (0001b) indicates the ACH. The version and the reserved values are both set to 0 as specified in [1]. MPLS-TP Connection Verification code point = 0xHH. [HH to be assigned by IANA from the PW Associated Channel Type registry.]

4.2. Out-of-band Message Identification

[To be added]

4.3. MPLS-TP CV Message Format

The format of an MPLS-TP CV Message is shown below.

```

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

```

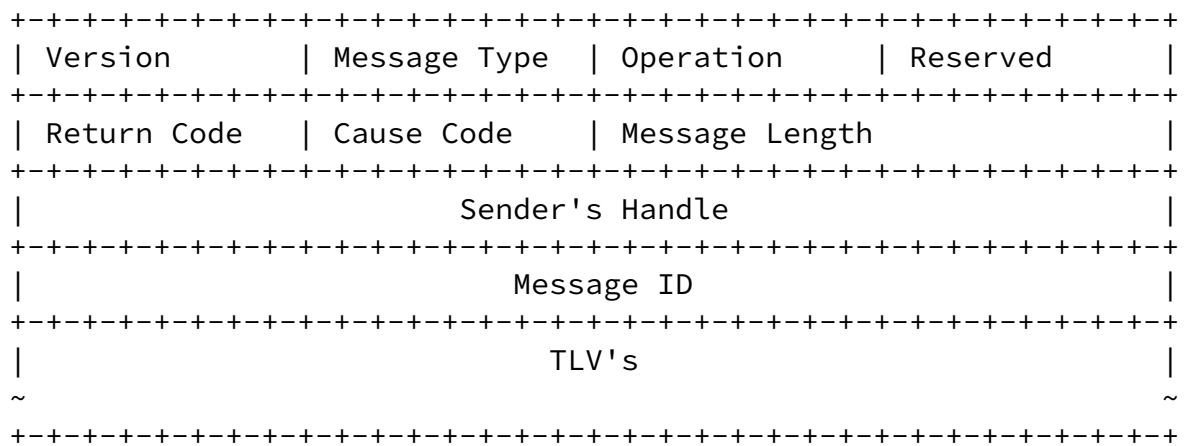


Figure 2: MPLS-TP CV Message Format

Version

Boutros

Expires September 9, 2009

[Page 6]

Internet-Draft

[draft-boutros-mpls-tp-cv-01.txt](#)

March 2009

The Version Number is currently 1.

Message Type

The following two message types are defined:

Message Type	Description
-----	-----
0x0	CV Request
0x1	CV Reply

Operation

The following two operations are defined:

Operation	Description
-----	-----
0x0	Only verify MPLS-TP LSP
0x1	Verify MPLS-TP LSP and record route
0x2	Verify MPLS-TP LSP, record route, and verify LSR ID in the record route before forwarding response

First, all operation codes are meaningful only in the MPLS-TP CV request message, and this field currently ignored in the MPLS-TP CV response message.

The operation code 0x0 is used to instruct a MIP or the receiver MEP to simply verify the MPLS-TP LSP associated with the MPLS-TP CV request message. In this case, after successfully processing the request message, an LSR should simply forward the message without appending Record Route TLV.

The operation code 0x1 is used to instruct a MIP or the receiver MEP to not only verify the MPLS-TP but also append a Record Route TLV to the request message if the message is successfully processed. Also, if the receiver needs to send a positive response back to the sender, it MUST include all the Record Route TLVs appended to the message by itself and all the upstream LSRs. Note that if a negative response is to be sent, Record Route TLVs are not appended to the response.

The operation code 0x2 is used to instruct a MIP receiving an MPLS-TP CV request message to verify the connection and append Record Route TLV. Additionally, it also instructs the LSR originating the response (MIP or MEP) to set the TTL value in the response such that the response will be intercepted by each upstream LSR. The intention is to let each upstream LSR to verify that the Record Route TLV that it appended to the request message exists in the response as well. Note that such verification is required only when positive response is sent. To facilitate such verification, the originator of the response as well as each LSR intercepting the response MUST set the TTL value to 1 in the response.

Return code

Value	Meaning
-----	-----
0	Success
1	Failure

Cause code

Value	Meaning
-----	-----
0	No cause code
1	Fail to find MPLS-TP LSP
2	Malformed CV message received
3	Received unknown TLV
4	Authentication failed
5	MPLS-TP LSP not setup in downstream direction
6	MPLS-TP LSP not setup in upstream direction
7	MPLS-TP LSP not setup in both directions
8	LSR-ID is missing in the record route of positive response

In the case of cause code 3, the unknown TLVs can be optionally sent in the response message. Use cases of the above cause codes are explained in the operation section below.

When MPLS-TP CV response travels back to the sender, a MIP intercepting the message could check if the Record Route TLV that it appended to the request exists in the response. As

such, the cause code 8 is meaningful only in the response message.

Sender's Handle

The Sender's Handle is filled in by the sender, and remains in tact as the CV request message travels. Also, this handle MUST be returned unchanged in all CV response messages. There are no semantics associated with this handle, although a sender may find this useful for matching up request with replies.

Message Length

The total length of any included TLVs.

Message ID

The Message ID is set by the sender of the MPLS-TP CV request message. It MUST be copied unchanged by any MEP or other MIP both in the CV request and response message. A sender SHOULD increment this value on each new message. A retransmitted message SHOULD leave the value unchanged.

An MPLS-TP CV request message MUST contain a LSPI TLV to identify the MPLS-TP LSP being verified, Source Address TLV identifying the sender of the message, and Destination Address TLV identifying the target recipient of the message. Note that in the successful case, the MPLS-TP CV response message MUST be originated from the target recipient of the request, and the target recipient can be MIP or a MEP. However, in the case of negative response, the LSR that fails to process the message generates the response message. When sending a response, the Source Address TLV identifies the LSR originating the response and the Destination Address TLV identifies the intended recipient of the message (which is the source of the request message). Format of these TLVs are specified in [6]. Furthermore, an Authentication TLV defined in [4] can be optionally included in the request message as well.

[4.4.](#) MPLS-TP Connection Verification Record Route TLV

```

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 = TBD   |                               Length = variable   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Upstream Label                        |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Downstream Label                      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               LSR Address                          |
~                               ~
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 3: MPLS-TP CV Record Route TLV format

The Record Route TLV includes the LSR address sub-TLV defined in [6] as well as the upstream and downstream labels (allocated by the LSR for both directions of the LSP). The upstream label is the label allocated by the LSR for the direction of the connection verification request message. The label value of 0 means that a label is not allocated for the respective direction.

Note that recording route is meaningful only if the connection verification operation is successful. As such, a receiver MUST examine any Record Route TLV only if the return code is 0 (success) in the connection verification response message.

[4.5](#). Network Management System

An operator should be able to provision any given LSR to send MPLS-OAM CV Request packets from a MEP and notify NMS when MPLS-OAM CV Response arrives.

[More description is to added]

[5](#). Operation

Consider an MPLS-TP LSP LSR-1 <--> LSR-2 <--> LSR-3 <--> LSR-4 <--> LSR-5. LSR-1 and LSR-5 are ingress and egress LSR for the respective direction. LSR-1 and LSR-5 are MEPs, and LSR-2 through LSR-4 are MIPs.

The proposed mechanism operates as follows:

1. LSR-1 sends an MPLS-TP CV Request message where the Source Address TLV, Destination Address TLV, and LSPI TLV represent LSR-1, LSR-5, and the LSP being verified respectively. An Authentication TLV may also be included.
2. The MPLS-TP CV Request message is intercepted at LSR-2 (MIP) because of TTL expiry. LSR-2 then verifies the request and:
 1. if the MPLS-TP LSP cannot be located, it sends a response with return code 1 and cause code 1.

2. if the message is malformed, it sends a response with return code 1 and cause code 2.
3. if any of the TLV is not known, it sends a response with return code 1 and cause code 3. It may also include the unknown TLVs.
4. if message authentication fails, it sends a response with return code 4 and cause code 4. This step is valid only if an Authentication TLV is present in the request.
5. if the MPLS-TP LSP is not setup in downstream direction, it sends a response with return code 1 and cause code 5.
6. if the MPLS-TP LSP is not setup in upstream direction, it sends a response with return code 1 and cause code 6.
7. if the MPLS-TP LSP is not setup in both directions, it sends a response with return code 1 and cause code 7.

Note that MPLS TTL value is set to 255 in the response message. In the response message, Source LSR address TLV is filled with the address of LSR-2.

When LSR-1 receives the MPLS-TP CV Response, the Destination Address TLV indicates that it is the intended recipient of the message. Furthermore, it learns that connection verification for the MPLS-TP LSP in question failed at LSR-2 by examining the LSPI and Source Address TLVs respectively in the message.

3. If LSR-2 is able to successfully process the MPLS-TP CV Request message, and if the MPLS-TP LSP is setup in both upstream and downstream directions, and if the destination address in CV request does not match LSR-2 address, it forwards the message to LSR-3 with TTL equals 1. LSR-2 appends its address as well as the upstream and

downstream labels to the message if the operation code is 1 or 2. Otherwise, LSR-2 simply forwards the message to LSR-3.

4. LSR-3 repeats the steps (2) or (3). In the absence of error, the messages progresses towards LSR-5 with each LSR adding its own ID and the local labels (for operation code 1 or 2).

5. Upon getting the MPLS-TP CV message, LSR-5 verifies the request. If an MPLS-TP LSP represented by LSPI TLV in the message is found, and if that MPLS-TP LSP is fully setup, LSR-5 checks the destination address in the CV request and if the destination address matches its address it sends an MPLS-TP CV response with return code 0 (success) back to the LSR-1. If the operation code in the request message is 1 or 2, LSR-5 appends all the Record Route TLVs received from upstream LSRs. Otherwise, the response does not include the Record Route TLVs received from the upstream LSRs. The TTL value in the response can set as follows:
 1. If the operation code in the request is 1, the TTL value is set to 255 so that the response message reaches LSR-1 without further interception at any other LSR.
 2. If the operation code in the MPLS-TP CV request message is 2, LSR-5 sends the response down the return path with TTL value equals 1 so that an LSR intercepting the message can verify its address and labels included in the response.
6. In case LSR-4 receives the response message, it checks if its address and labels are included in the record route. If the check fails, it sends an MPLS-TP CV response with return code 1 (error) with cause code 8 back to LSR-1, and in this case the address of LSR-4 is included in the Source Address TLV of the response. If the check succeeds, LSR-4 simply forwards the message to LSR-3.
7. When LSR-1 receives a response with a record route, it learns the address and the distance (in terms of hop count) of each LSR on the path of the MPLS-TP LSP.

6. Security Considerations

The security considerations for the authentication TLV need further study.

7. IANA Considerations

To be added.

8. References

8.1. Normative References

- [1] Bradner. S, "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March, 1997.
- [2] K. Kompella, G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), February 2006.
- [3] T. Nadeau, et. al, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires ", [RFC 5085](#), December 2007.
- [4] S. Boutros, et. al., "Operating MPLS Transport Profile LSP in Loopback Mode ", [draft-boutros-mpls-tp-loopback-01.txt](#), Work in Progress, December 2008.

8.2. Informative References

- [5] M. Bocci, et. al., "A Framework for MPLS in Transport Networks", [draft-ietf-mpls-tp-framework-00.txt](#), Work in Progress, November 2008.
- [6] S. Boutros, et. al., "Definition of ACH TLV Structure", [draft-bryant-mpls-tp-ach-tlv-00.txt](#), Work in Progress, January 2009.
- [7] M. Bocci, et. al., "MPLS Generic Associated Channel", [draft-ietf-mpls-tp-gach-gal-02.txt](#), work in progress, January 6, 2009.
- [8] Nabil Bitar, et. al, "Requirements for Multi-Segment Pseudowire Emulation Edge-to-Edge (PWE3)", [RFC5254](#), October 2008.

Author's Addresses

Sami Boutros
Cisco Systems, Inc.
3750 Cisco Way
San Jose, California 95134
USA
Email: sboutros@cisco.com

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario, K2K 3E8
Canada
Email: msiva@cisco.com

George Swallow
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough , MASSACHUSETTS 01719
United States
Email: swallow@cisco.com

David Ward
Cisco Systems, Inc.
3750 Cisco Way
San Jose, California 95134
USA
Email: wardd@cisco.com

Stewart Bryant
Cisco Systems, Inc.
250, Longwater, Green Park,
Reading RG2 6GB, UK
UK
Email: stbryant@cisco.com

Internet-Draft [draft-boutros-mpls-tp-cv-01.txt](#)

March 2009

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.

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.

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.

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.

Boutros

Expires September 9, 2009

[Page 15]

Internet-Draft

[draft-boutros-mpls-tp-cv-01.txt](#)

March 2009

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.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

Boutros

Expires September 9, 2009

[Page 16]