

Network Work group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 26, 2019

N. Kumar, Ed.  
C. Pignataro, Ed.  
F. Iqbal  
Z. Ali  
Cisco  
October 23, 2018

Label Switched Path (LSP) Ping/Traceroute for Segment Routing SIDs with  
MPLS Data-plane  
draft-nainar-mpls-spring-lsp-ping-sids-00

Abstract

RFC8402 introduces Segment Routing architecture that leverages source routing and tunneling paradigms and can be directly applied to the Multi Protocol Label Switching (MPLS) data plane. A node steers a packet through a controlled set of instructions called segments, by prepending the packet with Segment Routing header. SR architecture defines different types of segments with different forwarding semantics associated.

RFC8287 defines the extensions to MPLS LSP Ping and Traceroute for Segment Routing IGP-Prefix and IGP-Adjacency Segment Identifier (SIDs) with an MPLS data plane. RFC8287 defines the Target FEC Stack Sub-TLVs and the procedures to apply RFC8029 on SR architecture with MPLS data plane.

This document defines the Target FEC Stack Sub-TLVs and the extension required for other SR Segments.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 26, 2019.

## Copyright Notice

Copyright (c) 2018 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction . . . . .	3
2. Requirements notation . . . . .	3
3. Terminology . . . . .	3
4. Segment ID sub-TLV . . . . .	4
4.1. BGP Prefix Segment ID . . . . .	4
4.2. BGP Peering Segment - Peer-Node-SID . . . . .	4
4.3. BGP Peering Segment - Peer-Adj-SID . . . . .	5
4.4. BGP Peering Segment - Peer-Set-SID . . . . .	7
4.4.1. Peer Set Sub-TLV . . . . .	8
4.5. Path Binding SID . . . . .	9
4.6. Multicast Replication . . . . .	10
5. Procedures . . . . .	10
5.1. BGP Prefix SID . . . . .	10
5.2. BGP Peering Segment Sub-TLVs . . . . .	10
5.2.1. Initiator Node Procedures . . . . .	10
5.2.2. Responder Node Procedures . . . . .	11
5.3. Path Binding SID . . . . .	11
5.3.1. Initiator Node Procedures . . . . .	11
5.3.2. Responder Node Procedures . . . . .	11
6. IANA Considerations . . . . .	11
7. Security Considerations . . . . .	11
8. Acknowledgement . . . . .	12
9. Contributors . . . . .	12
10. References . . . . .	12
10.1. Normative References . . . . .	12
10.2. Informative References . . . . .	13
Authors' Addresses . . . . .	14

## 1. Introduction

[RFC8402] introduces and describes a Segment Routing architecture that leverages the source routing and tunneling paradigms. A node steers a packet through a controlled set of instructions called segments, by prepending the packet with Segment Routing header. A detailed definition of the Segment Routing architecture is available in [RFC8402]

As described in [RFC8402] and [I-D.ietf-spring-segment-routing-mpls], the Segment Routing architecture can be directly applied to an MPLS data plane, the Segment identifier (Segment ID) will be of 20-bits size and the Segment Routing header is the label stack.

[RFC8287] defines the mechanism to perform LSP Ping and Traceroute for Segment Routing with MPLS data plane. [RFC8287] defines the Target FEC Stack Sub-TLVs for IGP-Prefix Segment ID and IGP-Adjacency Segment ID.

There are various other Segment IDs proposed by different documents that are applicable for SR architecture. [I-D.ietf-idr-bgp-prefix-sid] defines BGP Prefix Segment ID, [I-D.ietf-idr-bgpls-segment-routing-epe] defines BGP Peering Segment ID such as Peer Node SID, Peer Adj SID and Peer Set SID. [I-D.sivabalan-pce-binding-label-sid] defines Path Binding Segment ID.

As above Segment IDs get deployed in the field, operators require corresponding MPLS OAM procedures for the SIDs. This document describes the target FEC Stack Sub-TLVs and the procedure to use LSP Ping and Traceroute for the above defined Segment IDs to support path validation and fault isolation.

## 2. Requirements notation

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. Terminology

This document uses the terminologies defined in [RFC8402], [RFC8029], readers are expected to be familiar with it.

The term "BGP EPE node" is used to refer to node assigning and advertising BGP Peering Segment SIDs to steer traffic towards a BGP peer, as described in [I-D.ietf-spring-segment-routing-central-epe].

#### 4. Segment ID sub-TLV

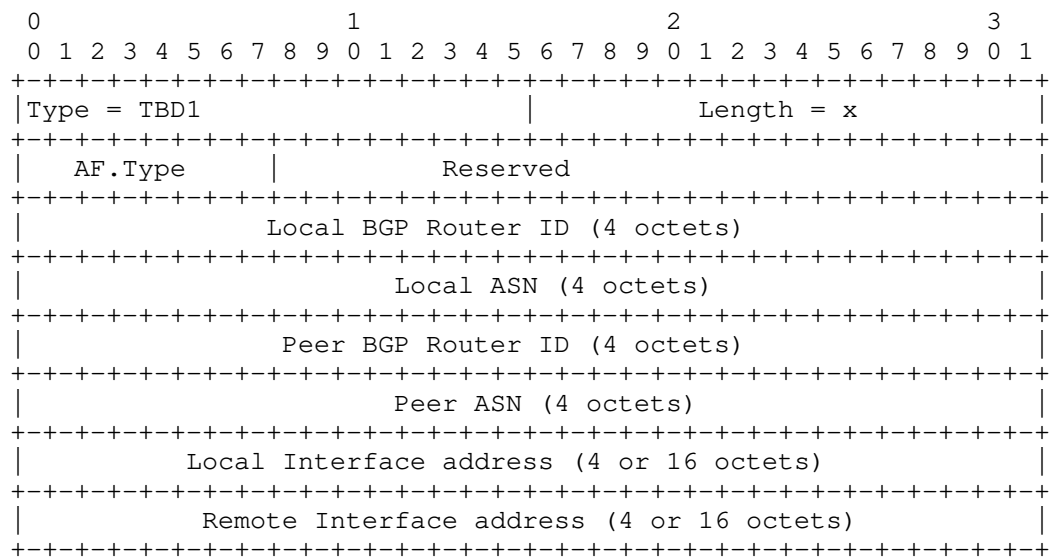
As defined in Section 5 of [RFC8287], the format of the following Segment ID sub-TLVs defined in this document follows the philosophy of Target FEC Stack TLV carrying FECs corresponding to each label in the label stack.

##### 4.1. BGP Prefix Segment ID

Section 3.2.13 and 3.2.14 of [RFC8029] defines the Sub-TLV for BGP labeled IPv4 and IPv6 prefix respectively. This document proposes the use of the same Sub-TLV for IPv4 and IPv6 BGP Prefix SID without any change.

##### 4.2. BGP Peering Segment - Peer-Node-SID

Peer-Node-SID identifies the peer node in the BGP Peering Segment. The sub-TLV format for Peer-Node-SID of BGP Peering Segment MUST be set as shown in the below TLV format:



#### AF.Type

Set to 4 if the address in Local/Remote Interface address field is IPv4 and set to 6 if the address in Local/Remote Interface address field is IPv6.

Reserved

MUST be set to 0 on send and MUST be ignored on receipt.

Local BGP Router ID

4-octet BGP Router ID of the node that assigns the Peer-Node-SID.

Local ASN

4-octet local ASN number of the node that assigns the Peer-Node-SID.

Peer BGP Router ID

4-octet BGP Router ID of the peer node.

Peer ASN

4-octet ASN number of the peer node.

Local Interface Address

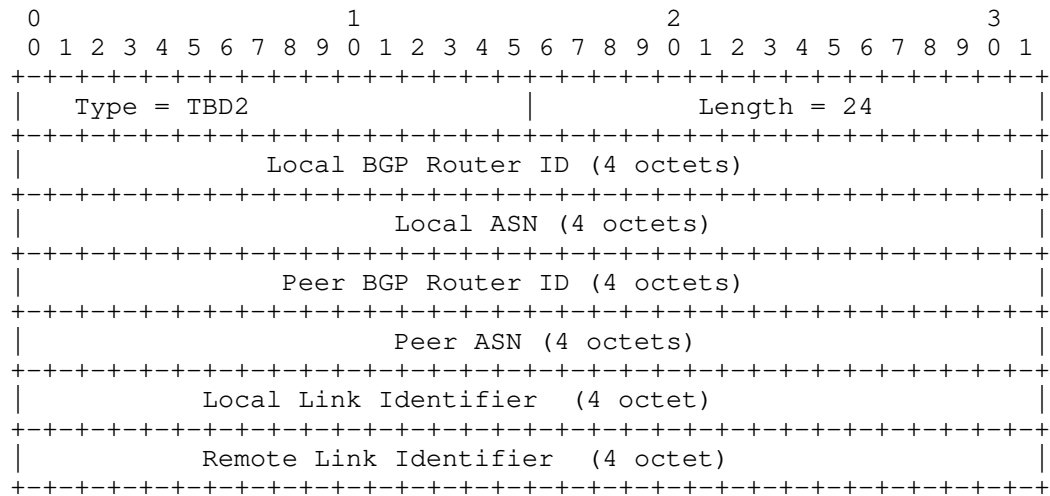
Set to the address used by the local node for BGP session peering. When AF.Type is set to 4, this address is 4-octet IPv4 address and when AF.Type is set to 6, this address is 16-octet IPv6 address.

Remote Interface Address

Set to the address used by the peer node for BGP session peering. When AF.Type is set to 4, this address is 4-octet IPv4 address and when AF.Type is set to 6, this address is 16-octet IPv6 address.

#### 4.3. BGP Peering Segment - Peer-Adj-SID

Peer-Adj-SID identifies the underlying link to the BGP peer node. The sub-TLV format for Peer-Adj-SID of BGP Peering Segment MUST be set as shown in the below TLV format:

**Local BGP Router ID**

4-octet BGP Router ID of the node that assigns the Peer-Node-SID.

**Local ASN**

4-octet local ASN number of the node that assigns the Peer-Node-SID.

**Peer BGP Router ID**

4-octet BGP Router ID of the peer node.

**Peer ASN**

4-octet ASN number of the peer node.

**Local Link Identifier**

Set to 4-octet link identifier of the local interface to which Peer-Adj-SID is assigned to.

**Remote Link Identifier**

Set to 4-octet link identifier of the peer interface to which Peer-Adj-SID is assigned to. Set to all-zeros when this identifier is unknown.

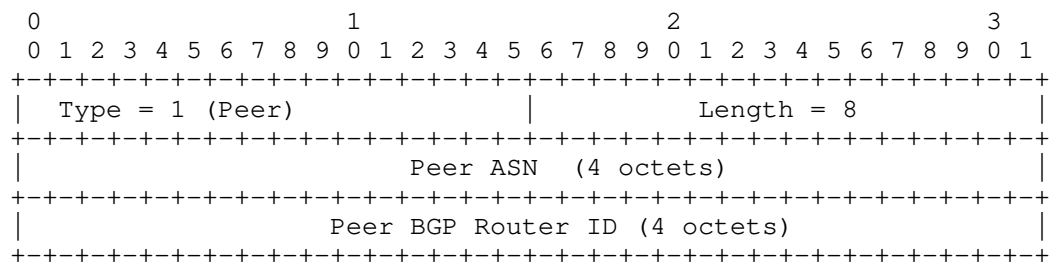


## 4.4.1. Peer Set Sub-TLV

As defined in section 5.3 of [I-D.ietf-idr-bgpls-segment-routing-epe], Peer-Set-SID can identify the set where the members can be Peer-Node or Peer-Adj from same or different ASN. The format of the Peer Set Sub-TLV will identify each such member.

## 4.4.1.1. Peer Node

The format for this sub-TLV MUST be set as below:



Peer ASN

4-octet ASN number of the peer node.

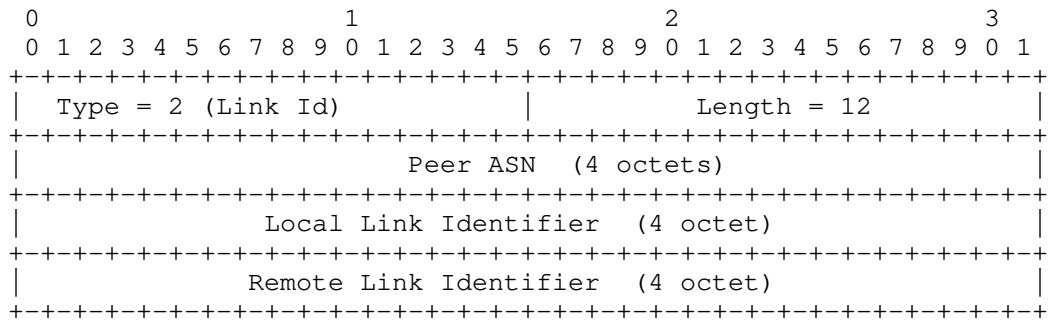
Peer Router ID

4-octet BGP Router ID of the peer node.

## 4.4.1.2. Link Identifier

The format for this sub-TLV is as below:





Peer ASN

4-octet ASN number of the peer node.

Local Link Identifier

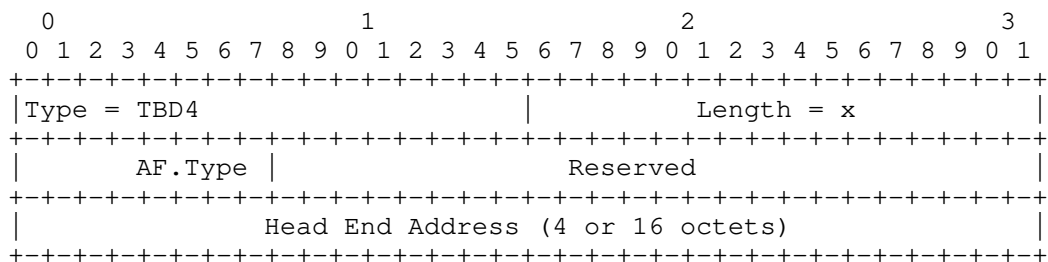
Set to 4-octet link identifier of the local interface to which Peer-Adj-SID is assigned to.

Remote Link Identifier

Set to 4-octet link identifier of the peer interface to which Peer-Adj-SID is assigned to. Set to all-zeros when this identifier is unknown.

#### 4.5. Path Binding SID

Path Binding SID identifies the Binding Segment Identifier associated with an RSVP-TE or SR-TE path. The format for this sub-TLV is as below:



AF.Type

Set to 4 if the address in Head End Address field is IPv4 and set to 6 if the address in Head End address field is IPv6.

Reserved

MUST be set to 0 on send and MUST be ignored on receipt.

Head End Address

Set to the address of the head end node to which the policy is assigned. When AF.Type is 4, this address is IPv4 and when AF.Type is 6, it is IPv6.

#### 4.6. Multicast Replication

[I-D.voyer-spring-sr-p2mp-policy] describes Segment Routing Multicast Replication Policy and introduces the notion of Tree SID to achieve this. A future version of this document will describe LSP Ping and Traceroute Target FEC Stack sub-TLV and procedures for Tree SID validation.

### 5. Procedures

This section describes the aspects of LSP Ping and Traceroute operations that require further considerations beyond [RFC8029] and [RFC8287].

#### 5.1. BGP Prefix SID

The procedures described in [RFC8029] are sufficient for MPLS Ping and Traceroute operations for BGP Prefix SID using the FEC definitions from Section 3.2.13 and 3.2.14 of [RFC8029].

#### 5.2. BGP Peering Segment Sub-TLVs

BGP Peering Segment sub-TLVs (BGP-Node-SID, BGP-Adj-SID, Peer-Set-SID) are assigned by BGP EPE node for a particular BGP neighbor, and advertised to the peer nodes. Any LSP Ping and Traceroute operation MUST be performed on the BGP EPE node, and not the remote neighbor node, as only the BGP EPE node can validate the contents of BGP Peering Segment sub-TLVs. Additionally, leaking the echo packet to the peer node may not be desirable for network operators.

##### 5.2.1. Initiator Node Procedures

If the bottom-most label in the label stack is BGP Peer Segment label, the initiating node MUST set the TTL of the bottom-most label to 1 to ensure that MPLS TTL expires at the BGP EPE node, and the

echo packet does not leak to the BGP peer node. Echo packet MUST include one of BGP-Node-SID, BGP-Adj-SID, or Peer-Set-SID sub-TLV in the Target FEC Stack TLV corresponding to the BGP Peer Segment label. Operator MAY push one or more transport labels on top of the BGP Peer Segment label to forward the echo packet to the BGP EPE node.

#### 5.2.2. Responder Node Procedures

In addition to procedures defined in [RFC8029], the responding node, upon TTL expiry of the echo packet, MUST process the incoming BGP Peer Segment sub-TLV of the Target FEC Stack. It MUST validate that contents of the sub-TLV and ensure the incoming label is advertised for the processed BGP Peer Segment sub-TLV.

#### 5.3. Path Binding SID

##### 5.3.1. Initiator Node Procedures

Similar to BGP Peering Segment sub-TLVs, Path Binding SID sub-TLV MUST be validated at the node assigning and advertising the Binding SID, instead of the endpoint of the path associated with the Binding SID. The initiating node MUST set the TTL of the Binding SID label to 1 and include the associated Path Binding SID TLV in the Target FEC Stack TLV of the echo request. Operator MAY push one or more transport labels on top of Binding SID label to forward echo packet from initiating node to the assigning node.

##### 5.3.2. Responder Node Procedures

In addition to procedures defined in [RFC8029], the responding node, upon TTL expiry of the echo packet, MUST process the incoming Path Binding SID sub-TLV of the Target FEC Stack. The responding node MUST ensure that it is the advertising node specified in the Path Binding SID sub-TLV, and the incoming Binding SID label matches the advertised label value.

#### 6. IANA Considerations

To be Updated.

#### 7. Security Considerations

To be Updated

## 8. Acknowledgement

TBD

## 9. Contributors

TBD

## 10. References

### 10.1. Normative References

[I-D.ietf-idr-bgp-prefix-sid]

Previdi, S., Filsfils, C., Lindem, A., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix SID extensions for BGP", draft-ietf-idr-bgp-prefix-sid-27 (work in progress), June 2018.

[I-D.ietf-idr-bgpls-segment-routing-epe]

Previdi, S., Filsfils, C., Patel, K., Ray, S., and J. Dong, "BGP-LS extensions for Segment Routing BGP Egress Peer Engineering", draft-ietf-idr-bgpls-segment-routing-epe-15 (work in progress), March 2018.

[I-D.sivabalan-pce-binding-label-sid]

Sivabalan, S., Tantsura, J., Filsfils, C., Previdi, S., Hardwick, J., and D. Dhody, "Carrying Binding Label/Segment-ID in PCE-based Networks.", draft-sivabalan-pce-binding-label-sid-04 (work in progress), March 2018.

[I-D.voyer-spring-sr-p2mp-policy]

daniel.voyer@bell.ca, d., Hassen, C., Gillis, K., Filsfils, C., Parekh, R., and H. Bidgoli, "SR Replication Policy for P2MP Service Delivery", draft-voyer-spring-sr-p2mp-policy-01 (work in progress), October 2018.

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

[RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, DOI 10.17487/RFC3443, January 2003, <<https://www.rfc-editor.org/info/rfc3443>>.

- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <<https://www.rfc-editor.org/info/rfc4203>>.
- [RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, DOI 10.17487/RFC5307, October 2008, <<https://www.rfc-editor.org/info/rfc5307>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.
- [RFC8287] Kumar, N., Ed., Pignataro, C., Ed., Swallow, G., Akiya, N., Kini, S., and M. Chen, "Label Switched Path (LSP) Ping/Traceroute for Segment Routing (SR) IGP-Prefix and IGP-Adjacency Segment Identifiers (SIDs) with MPLS Data Planes", RFC 8287, DOI 10.17487/RFC8287, December 2017, <<https://www.rfc-editor.org/info/rfc8287>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

## 10.2. Informative References

- [I-D.ietf-isis-segment-routing-extensions]  
Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-19 (work in progress), July 2018.
- [I-D.ietf-ospf-ospfv3-segment-routing-extensions]  
Psenak, P., Filsfils, C., Previdi, S., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPFv3 Extensions for Segment Routing", draft-ietf-ospf-ospfv3-segment-routing-extensions-15 (work in progress), August 2018.

- [I-D.ietf-ospf-segment-routing-extensions]  
Psenak, P., Previdi, S., Filsfils, C., Gredler, H.,  
Shakir, R., Henderickx, W., and J. Tantsura, "OSPF  
Extensions for Segment Routing", draft-ietf-ospf-segment-  
routing-extensions-25 (work in progress), April 2018.
- [I-D.ietf-spring-segment-routing-central-epe]  
Filsfils, C., Previdi, S., Dawra, G., Aries, E., and D.  
Afanasiev, "Segment Routing Centralized BGP Egress Peer  
Engineering", draft-ietf-spring-segment-routing-central-  
epe-10 (work in progress), December 2017.
- [I-D.ietf-spring-segment-routing-ldp-interop]  
Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., and  
S. Litkowski, "Segment Routing interworking with LDP",  
draft-ietf-spring-segment-routing-ldp-interop-15 (work in  
progress), September 2018.
- [I-D.ietf-spring-segment-routing-mpls]  
Bashandy, A., Filsfils, C., Previdi, S., Decraene, B.,  
Litkowski, S., and R. Shakir, "Segment Routing with MPLS  
data plane", draft-ietf-spring-segment-routing-mpls-14  
(work in progress), June 2018.
- [IANA-MPLS-LSP-PING]  
IANA, "Multi-Protocol Label Switching (MPLS) Label  
Switched Paths (LSPs) Ping Parameters",  
<[http://www.iana.org/assignments/mpls-lsp-ping-parameters/  
mpls-lsp-ping-parameters.xhtml](http://www.iana.org/assignments/mpls-lsp-ping-parameters/mpls-lsp-ping-parameters.xhtml)>.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5,  
RFC 792, DOI 10.17487/RFC0792, September 1981,  
<<https://www.rfc-editor.org/info/rfc792>>.

#### Authors' Addresses

Nagendra Kumar (editor)  
Cisco Systems, Inc.  
7200-12 Kit Creek Road  
Research Triangle Park, NC 27709-4987  
US

Email: [naikumar@cisco.com](mailto:naikumar@cisco.com)

Carlos Pignataro (editor)  
Cisco Systems, Inc.  
7200-11 Kit Creek Road  
Research Triangle Park, NC 27709-4987  
US

Email: cpignata@cisco.com

Faisal Iqbal  
Cisco Systems, Inc.

Email: faiqbal@cisco.com

Zafar Ali  
Cisco Systems, Inc.

Email: zali@cisco.com