Network Working Group Kireeti Kompella
Internet Draft Juniper Networks
Expiration Date: March 2001 Yakov Rekhter
Cisco Systems
Alan Kullberg
NetPlane Systems

Signalling Unnumbered Links in CR-LDP

draft-kompella-mpls-crldp-unnum-00.txt

1. Status of this Memo

2. Abstract

Current signalling used by MPLS TE doesn't provide support for unnumbered links. This document defines procedures and extensions to CR-LDP, one of the MPLS TE signalling protocols, that are needed in order to support unnumbered links.

3. Overview

Supporting MPLS TE over unnumbered links (i.e., links that do not have IP addresses) involves two components: (a) the ability to carry (TE) information about unnumbered links in IGP TE extensions (ISIS or OSPF), and (b) the ability to specify unnumbered links in MPLS TE signalling. The former is covered in [ISIS-TE, OSPF-TE]. The focus of this document is on the latter.

Current signalling used by MPLS TE doesn't provide support for unnumbered links because the current signalling doesn't provide a way to indicate an unnumbered link in its Explicit Route Objects. This document proposes simple procedures and extensions that allow CR-LDP [CR-LDP] signalling to be used with unnumbered links.

4. Interface Identifiers

Since unnumbered links are not identified by an IP address, then for the purpose of MPLS TE they need some other identifier. We assume that each unnumbered link on a Label Switched Router (LSR) is given a unique 16-bit identifier. The scope of this identifier is the LSR to which the link belongs; moreover, the IS-IS and/or OSPF and CR-LDP modules on an LSR must agree on interface identifiers.

Note that links are directed, i.e., a link l is from some LSR A to some other LSR B. LSR A chooses the interface identifier for link l. To be completely clear, we call this the "outgoing interface identifier from LSR A's point of view". If there is a reverse link from LSR B to LSR A (for example, a point-to-point SONET interface connecting LSRs A and B would be represented as two links, one from A to B, and another from B to A), B chooses the outgoing interface identifier for the reverse link. There is no a priori relationship between the two interface identifiers.

5. Unnumbered Forwarding Adjacencies

If an LSR that originates an LSP advertises this LSP as an unnumbered Forwarding Adjacency in IS-IS or OSPF [LSP-HIER], the LSR MUST allocate an interface ID to that Forwarding Adjacency. Moreover, the Local CR-LSP ID in the LSPID TLV of the Request Message for the LSP MUST be set to that interface ID, and the Ingress LSR Router ID in the LSPID TLV of the LSP MUST be set to the Router ID of the LSR that originates the LSP.

If the LSP is bidirectional, and the tail-end LSR (of the forward LSP) advertises the reverse LSP as an unnumbered Forwarding

Adjacency, the tail-end LSR MUST allocate an interface ID to the reverse Forwarding Adjacency. Furthermore, it MUST set the "Reverse Interface ID" field in the Reverse Interface ID TLV in the MAPPING message to the reverse FA's interface ID. The Reverse Interface ID's format is shown in Figure 1:

Figure 1: Reverse Interface ID 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0|0|             Type          |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             MUST be zero      | Reverse Interface ID (16 bits)|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Type (Reverse Interface ID) is to be determined by IETF consensus and the Length is 4.

6. Signalling Unnumbered Links in EROs

A new subobject of the Explicit Route Object (ERO) is used to specify unnumbered links. This subobject has the following format:

Figure 2: Unnumbered Interface ID Subobject

```
 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|             Type          |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|           MUST be zero      |     Interface ID (16 bits)    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

This subobject MUST be strict (i.e., the L bit MUST be 0). The Type is 0x0805 (Unnumbered Interface ID) and the Length is 4.

An LSR sending a Request message that includes an Unnumbered Interface ID subobject as the first subobject in the ERO MUST also include a PHOP TLV, specifying the Router ID of the sending LSR. This TLV is depicted in Figure 3.

The Type (PHOP TLV) is to be determined by IETF consensus and the Length is 4.

Figure 3: PHOP 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0|0|            Type           |            Length             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Sending LSR's Router ID                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

6.1. Interpreting the Unnumbered Interface ID Subobject

The Interface ID is the outgoing interface identifier with respect to the previous node in the path (i.e., the PHOP). If the Request message contains an Unnumbered Interface ID subobject as the first subobject in the ERO, then the PHOP object in the message must contain the router ID of the previous node.

6.2. Processing the Unnumbered Interface ID Subobject

A node that receives a Request message with an Unnumbered Interface ID as the first subobject in the ERO carried by the message MUST check whether the tuple <PHOP, Interface ID> matches the tuple <Ingress LSR Router ID, CR-LSP ID> of any of the LSPs for which the node is a tail-end. If a match is found, the match identifies the Forwarding Adjacency for which the node has to perform label allocation.

Otherwise, the node MUST check whether the tuple <PHOP, Interface ID> matches the tuple <Ingress LSR Router ID, Reverse Interface ID> of any of the bidirectional LSPs for which the node is the head-end. If a match is found, the match identifies the Forwarding Adjacency for which the node has to perform label allocation, namely, the reverse Forwarding Adjacency for the LSP identified by the match.

Otherwise, it is assumed that the node has to perform label allocation for the link over which the Request message was received. In this case the receiving node MAY validate that it received the Request Message correctly. To do so, the node must maintain a database of Traffic Engineering information distributed by IS-IS and/or OSPF.

To validate that it received the Request message correctly, the node looks up in its Traffic Engineering database for the node corresponding to the router ID of the sender of the Request message. It then checks that there is a link from the previous node to itself that carries the same Interface ID as the one in the ERO subobject.

If this is not the case, the receiving node has received the message in error and SHOULD return a "Bad Initial ER-Hop" error. Otherwise, the receiving node removes the first subobject, and continues processing the ERO.

6.3. Selecting the Next Hop

If, after processing and removing all initial subobjects in the ERO that refer to itself, the receiving node finds a subobject of type Unnumbered Interface ID, it determines the next hop as follows. The Interface ID MUST refer to an outgoing interface identifier that this node allocated; if not, the node SHOULD return a "Bad Strict Node" error. The next hop is the node at the other end of the link that the Interface ID refers to.

Furthermore, when sending a Request message to the next hop, the ERO to be used is the current ERO (starting with the Unnumbered Interface ID subobject).

7. Security Considerations

This document raises no new security concerns for CR-LDP.

8. Acknowledgments

Thanks to Rahul Aggarwal for his comments on the text.

9. References

[CR-LDP] Jamoussi, B., editor, "Constraint-Based LSP Setup using LDP", draft-ietf-mpls-cr-ldp-04.txt (work in progress)

[ISIS-TE] Smit, H., and Li, T., "IS-IS extensions for Traffic Engineering", draft-ietf-isis-traffic-02.txt (work in progress)

[LSP-HIER] Kompella, K., and Rekhter, Y., "LSP Hierarchy with MPLS TE", draft-ietf-mpls-lsp-hierarchy-01.txt (work in progress)

[OSPF-TE] Katz, D., and Yeung, D., "Traffic Engineering Extensions to OSPF", draft-katz-yeung-ospf-traffic-02.txt (work in progress)

10. Author Information

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
e-mail: kireeti@juniper.net

Yakov Rekhter
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
e-mail: yakov@cisco.com

Alan Kullberg
NetPlane Systems, Inc.
888 Washington St.
Dedham, MA 02026
e-mail: akullber@netplane.com