Internet Draft
Expires : April, 2001
<draft-kini-rsvp-lsp-restoration-00.txt>

Sriganesh Kini
Murali Kodialam
T.V.Lakshman
- Bell Labs

Curtis Villamizar
- Avici Systems

ReSerVation Protocol with Traffic Engineering extensions. Extension for Label Switched Path restoration

draft-kini-rsvp-lsp-restoration-00.txt

Status of this Memo

Abstract

Traffic engineering using MPLS involves the setting up of label switched paths (LSP) possibly with explicit routing and with bandwidth guarantees. The reliability of these LSPs can be increased by providing a backup LSP onto which traffic can be switched upon failure of an element in the path of the active LSP. Backup LSPs can be routed in a way that bandwidth can be shared between backup links of more than one active path while still guaranteeing recoverability for a set of failures. This sharing greatly increases the network efficiency thereby increasing the number of LSPs that can be carried while maintaining guarantees. Algorithms which can route such recoverable LSPs while using only aggregate network usage information are being developed. Keeping these algorithms as the primary motivation this document describes a mechanism to signal shared backup LSPs using RSVP.

1. Introduction

The concept of sharing links along backup paths is explained in [3]. RSVP (as described in [1]) with the extensions (as described in [2]) is used to signal a LSP. This document proposes further extensions RSVP. These extensions are useful to signal shared backup LSPs.

2. Extensions to RSVP-TE

2.1 Backup LSP Information Object

A new type of object "Backup LSP Information Object" is defined.

Four new C-Types are defined

2.1.1 Type of LSP C-Type = 1

0 = Active path
1 = Backup path

This is a one byte length field.

When the status changes from backup to active path the node should allocate bandwidth as if this path is an active path.

2.1.2 Route of Active Path C-Type = 2

The route of the active LSP is specified in the Path message during the setup of the backup LSP using the Route of Active Path C-Type = 2.

The format is same as ERO Class = 20 and C-Type = 1.

Processing :
This object should be forwarded unchanged to the next hop.
Changes in Route of Active Path C-Type = 2 :
A change in this object SHOULD result in deallocation of previously allocated bandwith for backup path and then reallocate bandwidth for this backup LSP.

2.1.3 Failure Entity C-Type = 3

The entity(ies) whose failure has to be protected against, is listed using the sub objects given below. If this information is conveyed in Path messages of the backup path then they can be used to allocate bandwidth for backup paths by sharing it between different backup paths. A failure entity set is used to denote a grouping for which a backup route can be specified. This object only needs to be sent in the path messages of the active path. The first node in the ERO for failure entity should setup the backup path for the failure entity set. If no ERO is given for the failure entity set the node adjacent to the first

object of the failure entity (or the first object of the failure entity set) should setup the backup LSP.

Five types of subobjects are currently defined.

2.1.3.1 Subobject 1: Link wih IPv4 address

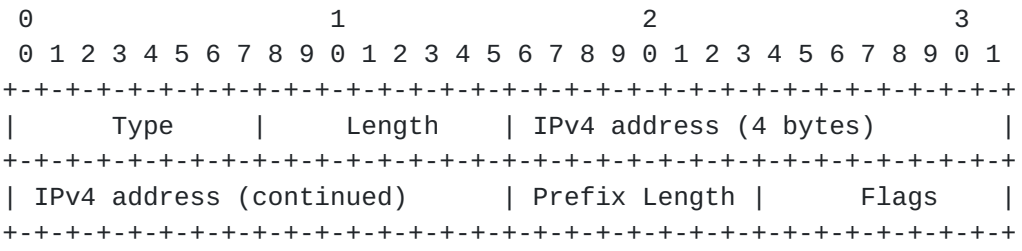

Type

0x01 IPv4 address of the link

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. A length of 2 denotes that every link has to be protected. Otherwise the length is 8.

IPv4 address

A 32-bit unicast, host address. Any network-reachable interface address is allowed here. Illegal addresses, such as certain loopback addresses, SHOULD NOT be used.

Prefix length

A prefix length of 32 denotes a specific link.

Flags

0x01 Start of failure entity set
0x02 End of failure entity set

2.1.3.2 Subobject 2: Link with IPv6 address

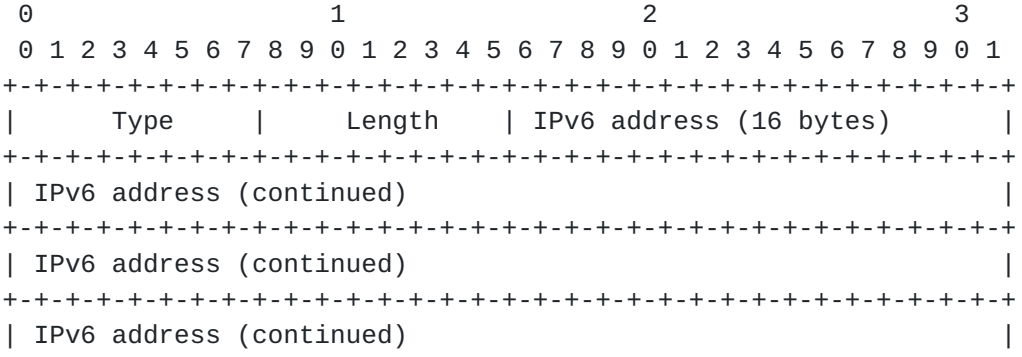

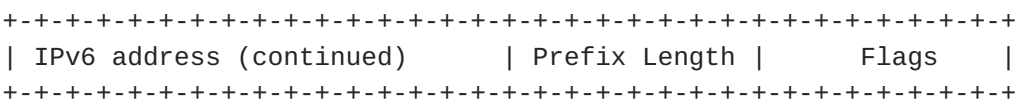
| IPv6 address (continued) | Prefix Length | Flags |

Type

0x02 IPv6 address

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. A length of 2 denotes that every link has to be protected. Otherwise the length is 20.

IPv6 address

A 128-bit unicast host address.

Prefix length

A prefix length of 128 denotes a specific link.

Flags

0x01 Start of failure entity set
0x02 End of failure entity set

2.1.3.3 Subobject 3: Node wih IPv4 address

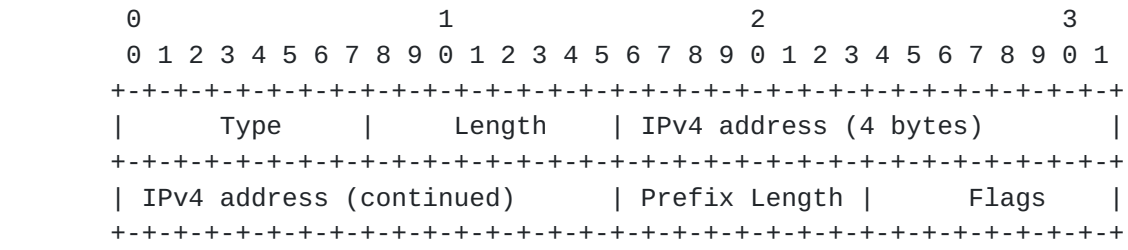

Type

0x03 IPv4 address (any address of the node)

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. A length of 2 denotes that every node has to be protected. Otherwise the length is 8.

IPv4 address

A 32-bit unicast, host address. Any network-reachable interface address is allowed here. Illegal addresses, such as certain loopback addresses, SHOULD NOT be used.

Prefix length

A prefix length of 32 denotes a specific node.

Flags

0x01 Start of failure entity set
0x02 End of failure entity set

2.1.3.4 Subobject 4: Node with IPv6 address

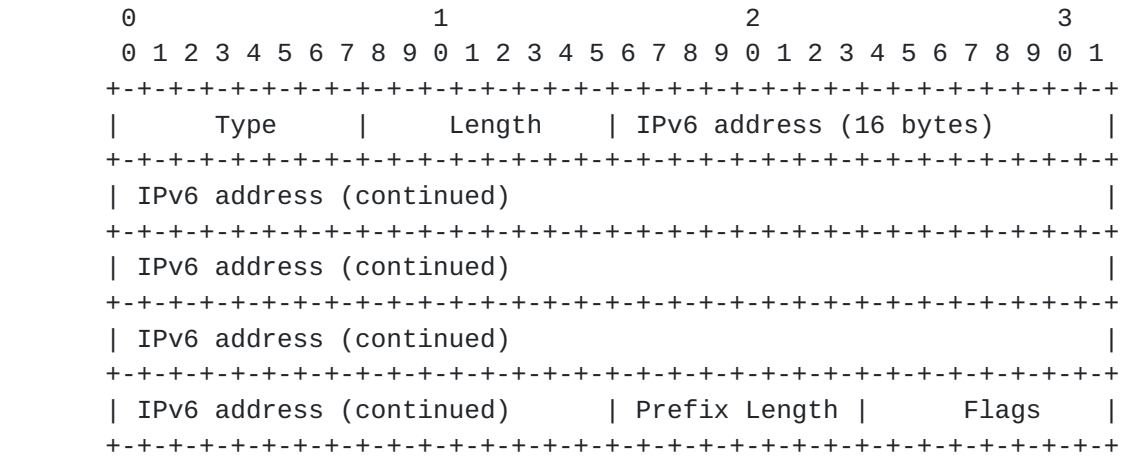

Type

0x04 IPv6 address (any address of the node)

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. A length of 2 denotes that every node has to be protected. Otherwise the length is 20.

IPv6 address

A 128-bit unicast host address.

Prefix length

A prefix length of 128 denotes a specific node.

Flags

0x01 Start of failure entity set
0x02 End of failure entity set

2.1.3.5 Subobject 5 : SRLG

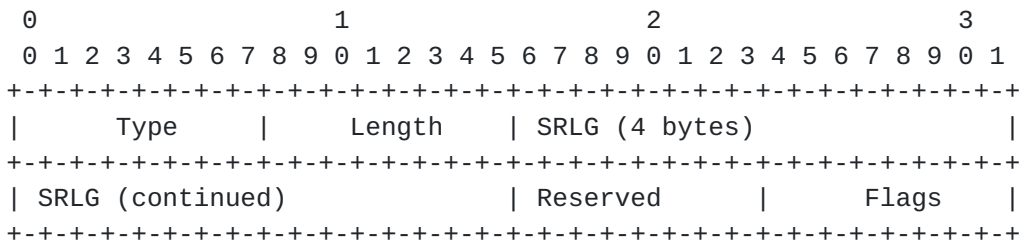

Type

0x05 SRLG

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. A length of 2 denotes that every SRLG has to be protected. Otherwise the length is 8.

SRLG

A 32-bit SRLG.

Flags

0x01 Start of failure entity set
0x02 End of failure entity set

2.1.3.6 Subobject 6 : ERO for that failure entity set

Denotes the Explicit route for that failure entity set. Should immediately follow the failure entity set. Format same as that described in section 4.3 of [2].

2.1.4 Active path session id C-Type = 5

Same as the session object as described in 4.6.1 of [2]. Should be included in all messages signaling the backup path.

2.1.5 Active path sender Template C-Type = 6

Same as the sender template object as described in 4.6.2 of [2]. Should be included in all messages signaling the backup path.

3. Security Considerations

This document raises no new security issues for RSVP.

4. IANA Considerations

The value of the new object "Backup LSP Information Object" should be

assigned by IANA. The value should be assigned so that the object is forwarded unchanged by a node which does not understand the object.

5. Acknowledgments

The authors would like to thank Vishal Sharma and Roch Guerin for their comments on this work.

6. References

[1] Braden, R., Zhang, L., Berson, S., Herzog, S., "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[2] Awduche, D.O, et al, "RSVP-TE: Extensions to RSVP for LSP Tunnesl", Internet Draft <draft-ietf-mpls-rsvp-lsp-tunnel-06.txt>, July 2000.

[3] Kini. S, et al, "Shared backup Label Switched Path restoration", Work in progress, Internet Draft <draft-kini-restoration-shared-backup-00.txt>. November 2000.

7. Author's Addresses

Sriganesh Kini
Lucent Technologies, Bell Labs
Room 4C-526, 101 Crawfords Corner Road
Holmdel, NJ 07733-3030
Phone : 732 949 6418
Email : kini@dnrc.bell-labs.com

Murali Kodialam
Lucent Technologies, Bell Labs
Room 4D-525, 101 Crawfords Corner Road
Holmdel, NJ 07733-3030
Phone : 732 949 6296
Email : muralik@dnrc.bell-labs.com

T.V.Lakshman
Lucent Technologies, Bell Labs
Room 4D-531, 101 Crawfords Corner Road
Holmdel, NJ 07733-3030
Phone : 732 949 4778
Email : lakshman@dnrc.bell-labs.com

Curtis Villamizar
Avici Systems
Email : curtis@avici.com

Full Copyright Statement