Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: April 6, 2021

H. Chen
R. Li
Futurewei
Y. Yang
IBM
A. Kumar S N
RtBrick
Y. Fan
Casa Systems
N. So

V. Liu

M. Toy
Verizon
L. Liu
Fujitsu
K. Makhijani
Futurewei
October 3, 2020

IS-IS Topology-Transparent Zone
draft-ietf-lsr-isis-ttz-01.txt

Abstract

This document specifies a topology-transparent zone in an area. A zone is a subset (block/piece) of an area, which comprises a group of routers and a number of circuits connecting them. It is abstracted as a virtual entity such as a single virtual node or zone edges mesh. Any router outside of the zone is not aware of the zone. The information about the circuits and routers inside the zone is not distributed to any router outside of the zone.

Table of Contents

1. Introduction

[ISO10589] describes two levels of areas in IS-IS, level 1 and level 2 areas. There are scalability issues in using areas as the number of routers in a network becomes larger and larger.

Through splitting the network into multiple level 1 areas connected by level 2, we may extend the network further. However, dividing a network from one area into multiple areas or from a number of existing areas to even more areas can be a challenging and time consuming task since it involves significant network architecture changes.

These issues can be resolved by using topology-transparent zone (TTZ), which abstracts a zone (i.e., a subset of an area) as a single virtual node or zone edges' mesh with minimum efforts and minimum service interruption. Note that a zone can be an entire area.

This document presents a topology-transparent zone and specifies extensions to IS-IS that support topology-transparent zones.

1.1. Requirements Language

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] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

TTZ: A Topology-Transparent Zone.

Zone: A subset (block or piece) of an area. In a special case, a zone is an entire area.

Zone External Node: A node outside of a zone.

Zone Internal Node: A node within a zone without any connection to a node outside of the zone.

Zone Edge/Border Node: A node that is part of a zone connecting to a node outside of the zone.

Zone Node: A zone internal node or a zone edge/border node (i.e., a node that is part of a zone).

Zone Link: A link connecting zone nodes (i.e., a link that is part of a zone).

Zone Neighbor Node: A node outside of a zone that is a neighbor of a zone edge/border node.

CLI: Command Line Interface.

LSP: A Link State Protocol Data Unit (PDU) in IS-IS. An LSP contains link state information. In general, a router/node originates multiple LSPs, distinguished by LSP fragment number, to carry the link state information about it and the links attached to it.

LS: Link State. In general, the LS for a node is all the LSPs that the node originates. The LS for a zone is the set of LSPs that all the nodes in the zone originate to carry the information about them and the links attached to them inside the zone.

2. Requirements

A Topology-Transparent Zone (TTZ) may be deployed to resolve some critical issues such as scalability in existing networks and future networks. The requirements for a TTZ are as follows:

- o TTZ MUST be backward compatible. When a TTZ is deployed on a set of routers in a network, the routers outside of the TTZ in the network do not need to know or support the TTZ.

- o TTZ MUST support at least one more levels of network hierarchy, in addition to the hierarchies supported by existing IS-IS.

- o Abstracting a zone as a TTZ virtual entity, which is a single virtual node or zone edges' mesh, SHOULD be smooth with minimum service interruption.

- o De-abstracting (or say rolling back) a TTZ virtual entity to a zone SHOULD be smooth with minimum service interruption.

- o Users SHOULD be able to easily set up an end-to-end service crossing TTZs.

- The configuration for a TTZ in a network SHOULD be minimal.

- The changes on the existing protocols to support TTZ SHOULD be minimal.

3. Zone Abstraction

A zone can be abstracted as a single virtual node or the zone edges' full mesh.

When a zone is abstracted as a single virtual node, this node appears to be connected to all the neighbors of the zone, and to be in the same area as those neighbors.

When a zone is abstracted as its edges' full mesh, there is a full mesh connections among the edges and each edge is also connected to its neighbors outside of the zone.

4. Topology-Transparent Zone

A Topology-Transparent Zone (TTZ) comprises an Identifier (ID) and a subset (piece/block) of an area such as a Level 2 area in IS-IS. It is abstracted as a single virtual node or its edges' full mesh. TTZ and zone as well as node and router will be used interchangeably below.

4.1. Zone as a Single Node

After a zone is abstracted as a single virtual node having a virtual node ID, every node outside of the zone sees a number of links connected to this single node. Each of these links connects a zone neighbor. The link states inside the zone are not advertised to any node outside of the zone. The virtual node ID may be derived from the zone ID.

4.1.1. An Example of Zone as a Single Node

The figure below shows an example of an area containing a TTZ: TTZ 600.

```
                 TTZ 600
                   \
                    \ ^~^~^~^~^~^~^~^~^~^~^~^~
                     (                        )
    ===[R15]========(==[R61]------------[R63]==)======[R29]===
        ||         (   |    \          /    |   )       ||
        ||         (   |     \        /     |   )       ||
        ||         (   |      \      /      |   )       ||
        ||         (   |    ___\    /       |   )       ||
        ||         (   |   /   [R71]        |   )       ||
        ||         (   | [R73] /    \       |   )       ||
        ||         (   |      /      \      |   )       ||
        ||         (   |     /        \     |   )       ||
        ||         (   |    /          \    |   )       ||
    ===[R17]========(==[R65]------------[R67]==)======[R31]===
         \\          (//                 \\)       //
          ||         //v~v~v~v~v~v~v~v~v~v~v~\\      ||
          ||        //                       \\     ||
          ||       //                         \\    ||
           \\     //                           \\  //
       ======[R23]==============================[R25]=====
             //                                    \\
            //                                      \\
```

Figure 1: An Example of TTZ 600

The area comprises routers R15, R17, R23, R25, R29 and R31. It also contains TTZ 600, which comprises routers R61, R63, R65, R67, R71 and R73, and the circuits connecting them.

There are two types of routers in a TTZ: TTZ internal routers and TTZ edge/border routers. A TTZ internal router is a router inside the TTZ and its adjacent routers are inside the TTZ. A TTZ edge/border router is a router inside the TTZ and has at least one adjacent router that is outside of the TTZ.

The TTZ in the figure above comprises four TTZ edge/border routers R61, R63, R65 and R67. Each TTZ edge/border router is connected to at least one router outside of the TTZ. For instance, router R61 is a TTZ edge/border router since it is connected to router R15, which is outside of the TTZ.

In addition, the TTZ comprises two TTZ internal routers R71 and R73. A TTZ internal router is not connected to any router outside of the TTZ. For instance, router R71 is a TTZ internal router since it is not connected to any router outside of the TTZ. It is just connected to routers R61, R63, R65, R67 and R73 inside the TTZ.

A TTZ MUST hide the information inside the TTZ from the outside. It MUST NOT directly distribute any internal information about the TTZ to a router outside of the TTZ.

For instance, the TTZ in the figure above MUST NOT send the information about TTZ internal router R71 to any router outside of the TTZ in the routing domain; it MUST NOT send the information about the circuit between TTZ router R61 and R65 to any router outside of the TTZ.

From a router outside of the TTZ, a TTZ is seen as a single node (refer to the Figure below). For instance, router R15, which is outside of TTZ 600, sees TTZ 600 as a single node Rz, which has normal connections to R15, R29, R17 and R23, R25 and R31.

```
              TTZ 600
                \
                 \ ^~^~^~^~^~^~^~^~^~^~^~^~
                  (                        )
 ===[R15]========(                          )======[R29]===
     ||         (                            )       ||
     ||         (                            )       ||
     ||         (                            )       ||
     ||         (                            )       ||
     ||         (             Rz             )       ||
     ||         (                            )       ||
     ||         (                            )       ||
     ||         (                            )       ||
     ||         (                            )       ||
 ===[R17]========(                          )======[R31]===
      \\          (                        )       //
       ||         //v~v~v~v~v~v~v~v~v~v~v~\\      ||
       ||        //                        \\     ||
       ||       //                          \\    ||
        \\     //                            \\  //
    ======[R23]==============================[R25]=====
          //                                     \\
         //                                       \\
```

Figure 2: TTZ 600 as Single Node Rz

4.1.2. Zone Leader Election

A node in a zone is elected as a leader for the zone, which is the node with the highest priority (and the highest node ID when there are more than one nodes having the same highest priority) in the zone. The leader election mechanism described in

[I-D.ietf-lsr-dynamic-flooding] may be used to elect the leader for the zone.

4.1.3. LS Generation for Zone as a Single Node

The leader for the zone originates the LS (i.e., set of LSPs) for the zone as a single virtual node and sends it to its neighbors.

This LS comprises all the adjacencies between the virtual node and the zone neighbors. The System ID of each LSP ID is the ID of the virtual node for the zone. The Source ID or Advertising Node/Router ID is the ID of the virtual node.

In addition, this LS may contain the IP prefixes such as the loopback IP addresses inside the zone to be accessed by zone external nodes (i.e., nodes outside of the zone). These IP prefixes are included in the IP internal reachability TLV.

4.1.4. Adjacency Establishment

A zone edge node X, acting as a proxy for the single virtual node for the zone, forms an adjacency between the virtual node and a node Y that is outside of the zone and in node X's area as described below.

For a new adjacency (i.e., no adjacency exists between X and Y):

Every IS-IS protocol packet, such as Hello, that edge node X originates and sends node Y, uses the virtual node ID as Source ID.

When node X synchronizes its link state database (LSDB) with node Y, it sends Y all the link state information except for the link state belonging to the zone that is hidden from the nodes outside of the zone.

At the end of the LSDB synchronization, the LS for the zone as a single virtual node is originated by the zone leader and distributed to node Y. This LS contains the adjacencies between the virtual node and all the zone neighbors, including this newly formed zone neighbor Y.

For an existing adjacency (i.e., an adjacency already exists between X and Y):

At first, edge node X acting as a proxy for the virtual node creates a new adjacency between the virtual node for the zone and node Y in a normal way. It sends Hellos and other packets containing the virtual node ID as Source ID to node Y. Node Y establishes an adjacency with the virtual node in the normal way.

Then, node X terminates the existing adjacency between node X and node Y after the zone has transitioned to be the virtual node. It stops sending Hellos for the adjacency to node Y. Without receiving Hellos from node X for a given time such as hold-timer interval, node Y removes the adjacency to node X. Even though this adjacency terminates, node X keeps the link to node Y in its LS.

In the case where node Y is not in node X's area, is in the backbone and connected to node X, node X, acting as a proxy for the virtual node, creates a new adjacency between the virtual node and node Y in a normal way and sends the LS for the virtual node to node Y if the zone includes all the nodes in its area.

4.1.5. Computation of Routes

After a zone is transferred/migrated to a single virtual node, every zone node computes the routes (i.e., shortest paths to the destinations) using the zone topology, the connections between each zone edge and its zone neighbor, and the topology outside of the zone without the virtual node. The metric of a link outside of the zone is one order of magnitude larger than the metric of a link inside the zone.

4.2. Extensions to Protocols

The following TLV is defined in IS-IS.

o Zone ID TLV: containing a zone ID, a flags field and optional sub-TLVs.

4.2.1. Zone ID TLV

The format of IS-IS Zone ID TLV is illustrated below. It may be added into an LSP for a zone node.

```
 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 (TBD1)  |     Length    |           Zone ID             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Zone ID (Continue)                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  RESV                |E|  OP |           Sub TLVs              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               +
~                                                               ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: IS-IS Zone ID TLV

Type (1 byte): TBD1.

Length (1 byte): Its value is variable with a minimum of 8. A value larger than 8 means that sub-TLVs are present. If length is less than 8, the TLV MUST be ignored.

Zone ID (6 bytes): It is the identifier (ID) of a zone.

Flags field (16 bits): one flag bit E, OP of 3 bits, and a reserved subfield are as follows:

RESV: Reserved. MUST be send as zero and ignored on receipt.
E = 1: Indicating a node is a zone edge node
E = 0: Indicating a node is a zone internal node

When a Zone ID is configured on a zone node (refer to Section 6.1), the node updates its LSP by adding an IS-IS Zone ID TLV with the Zone ID. If it is a zone internal node, the TLV has its flag E = 0; otherwise (i.e., it is a zone edge node) the TLV has its flag E = 1 and may include a Zone ISN Sub TLV containing the zone links configured. Every link of a zone internal node is a zone link. If every link of a zone edge node is a zone link, the TLV with E = 1 does not include any Zone ISN Sub TLV; otherwise (i.e., some of its links are zone links), it includes the Zone ISN Sub TLV containing the zone links.

OP Value	Meaning (Operation)
0x001 (T):	Advertising Zone Topology Information for Migration
0x010 (M):	Migrating Zone to a Virtual Entity such as Virtual Node
0x011 (N):	Advertising Normal Topology Information for Rollback
0x100 (R):	Rolling Back from the Virtual Entity

The value of OP indicates one of the four operations above. When any of the other values is received, the TLV MUST be ignored.

When a zone node, such as the zone leader, receives a command via management, such as a CLI command, to advertise zone information for migration, or determines to advertise, it sets OP = T (i.e., 0x001) in the Zone ID TLV of its LSP. When a zone node receives a command to migrate zone to a virtual entity, or determines to migrate, it sets OP = M (i.e., 0x010) in the Zone ID TLV of its LSP. When a zone node receives a command to advertise Normal topology information for roll back, it sets OP = N (i.e., 0x011) in the Zone ID TLV of its LSP. When a zone node receives a command to roll back or determines to roll back, it sets OP = R (i.e., 0x100) in the Zone ID TLV of its LSP.

Two new sub-TLVs are defined, which may be added to an IS-IS Zone ID TLV. One is the Zone IS Neighbor sub-TLV, or Zone ISN sub-TLV for short. The other is the Zone ES Neighbor sub-TLV, or Zone ESN sub-TLV for short. A Zone ISN sub-TLV contains the information about a number of IS neighbors in the zone connected to a zone edge node. It has the format 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Type (TBD)   |    Length     |         Neighbor ID(i)        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               +
|                                                               |
+               +-----------------------------------------------+
|               |           Metric (i)                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                                                               ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: Zone ISN Sub TLV

A Zone ISN Sub TLV has 1 byte of Type, 1 byte of Length of n*(IDLength + 3), which is followed by n tuples of Neighbor ID and Metric.

A Zone ESN sub-TLV contains the information about a number of ES neighbors in the zone connected to a zone edge node. It has the format 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Type (TBD)   |    Length     |         Neighbor ID(i)        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               +
|                                                               |
+               +-----------------------------------------------+
|               |           Metric (i)                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                                                               ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Zone ESN Sub TLV

4.3. Zone as Edges Full Mesh

4.3.1. Updating LSPs for Zone as Edges' Mesh

For every zone edge node, it updates its LSP in three steps and floods the LSP to all its neighbors.

At first, it adds each of the other zone edge nodes as an IS neighbor into the Intermediate System Neighbors TLV in its LSP after it receives an LSP containing an IS-IS Zone ID TLV with OP = M or a command activating migration zone to an abstracted entity. The metric to the neighbor is the metric of the shortest path to the edge node within the zone.

In addition, it adds an IP internal reachability TLV into its LSP. The TLV contains a number of IP prefixes in the zone to be reachable from outside of the zone.

And then it removes the IS neighbors corresponding to the IS neighbors in the Zone ID TLV (i.e., in the Zone ISN sub TLV) from Intermediate System Neighbors TLV in its LSP, and the ES neighbors corresponding to the ES neighbors in the Zone ID TLV (i.e., in the Zone ESN sub TLV) from End System Neighbors TLV in the LSP. This SHOULD be done after it receives the LSPs for virtualizing zone from the other zone edges for a given time.

4.4. Advertisement of LSPs

LSPs can be divided into a couple of classes according to their Advertisements. The first class of LSPs is advertised within a zone. The second is advertised through a zone.

4.4.1. Advertisement of LSPs within Zone

Any LSP about a link state in a zone is advertised only within the zone. It is not advertised to any router outside of the zone. For example, a LSP generated for a zone internal node is advertised only within the zone.

Any LSP generated for a broadcast network in a zone is advertised only within the zone. It is not advertised outside of the zone.

After migrating to a zone as a single virtual node or edges' full mesh, every zone edge MUST NOT advertise any LSP belonging to the zone or any information in a LSP belonging to the zone to any node outside of the zone. The zone edge determines whether an LSP is about a zone internal link state by checking if the originating node of the LSP is a zone internal node.

For any LSP originated by a node within the zone, every zone edge node MUST NOT advertise it to any node outside of the zone.

4.4.2. Advertisement of LSPs through Zone

Any LSP about a link state outside of a zone received by a zone edge is advertised using the zone as transit. For example, when a zone edge node receives an LSP from a node outside of the zone, it floods the LSP to its neighbors both inside and outside of the zone.

The nodes in the zone continue to flood the LSP. When another zone edge receives the LSP, it floods the LSP to its neighbors both inside and outside of the zone.

5. Seamless Migration

This section presents the seamless migration between a zone and its single virtual node.

5.1. Transfer Zone to a Single Node

Transferring a zone to a single virtual node smoothly takes a few steps or stages.

At first, a user configures the zone on every node of the zone (refer to Section 6.1). Every zone node updates its LSP by including a Zone ID TLV. For a zone edge node, the TLV has the Zone ID configured, its flag E = 1 and a Zone ISN Sub TLV containing the zone links configured. For a zone internal node, the TLV has the Zone ID configured and its flag E = 0.

Second, after finishing the configuration of the zone, a user may issue a command, such as a CLI command, on a zone node, such as the zone leader, to trigger transferring the zone to the single virtual node. When the node receives the command, it updates its LSP by setting OP = T in its Zone ID TLV, which is distributed to every zone node. After receiving the Zone ID TLV with OP = T, every zone edge node, acting as a proxy of the virtual node, establishes a new adjacency between the virtual node and each of its zone neighbor nodes.

The command may be replaced by the determination made by a zone node, such as the zone leader. After determining that the configuration of the zone is finished for a given time such as 10 seconds, it updates its LSP by setting OP = T in its Zone ID TLV. The configuration is complete if every zone link configured is bidirectional. For every zone internal node configured with the Zone ID, there is an LSP containing its Zone ID TLV with E = 0 in the LSDB, which indicates

that each link from the node (one direction) is a zone link. For every zone edge node, each of its zone links configured from the edge node (one direction) is included in its LSP containing its Zone ID TLV with E = 1 and Zone ISN Sub TLV in the LSDB.

Third, after receiving the updated LSPs from all the zone neighbor nodes, the zone leader checks if all the new adjacencies between the virtual node and the zone neighbor nodes have been established. If so, it originates an LS for the virtual node and updates its LSP (i.e., the LSP for itself zone leader) by setting OP = M in its Zone ID TLV, which is distributed to every zone node.

After receiving the Zone ID TLV with OP = M, every zone node migrates to zone as virtual node. Every zone edge node does not send any LS inside the zone to any zone neighbors. It advertises its LSP without any zone links to the nodes outside of the zone or purges its LSP outside of the zone, terminates its adjacency to each of its zone neighbors, but contains the adjacency in its LSP that is distributed within the zone. Every zone node computes the routes according to Section 4.1.5.

5.2. Roll Back from Zone as a Single Node

After abstracting a zone to a single virtual node, we may want to roll back the node to the zone smoothly in some cases. The process of rolling back has a few steps or stages.

At first, a user issues a command, such as a CLI command, on a zone node, such as the zone leader, to start (or prepare) for roll back. When receiving the command, the node updates its LSP by setting OP = N in its Zone ID TLV, which will be distributed to every node in the zone. After receiving the Zone ID TLV with OP = N, every zone edge node establishes a normal adjacency between the edge node and each of its zone neighbor nodes, and advertises the link state of the zone over the adjacency if it crosses the adjacency, but holds off its LSP containing the normal adjacency.

Second, a user may issue a command, such as a CLI command, on a zone node, such as the zone leader, to roll back from the virtual node to the zone if the following conditions are met.

Condition 1: All the normal adjacencies between every zone edge node and each of its zone neighbor nodes have been established.

Condition 2: All the link state about the zone that is supposed to be advertised outside of the zone has been advertised.

After receiving the command, the node updates its LSP by setting OP = R in its Zone ID TLV, which is distributed to every zone node. After receiving the Zone ID TLV with OP = R,

- o every zone edge node, acting as a proxy of the virtual node, terminates the adjacency between the virtual node and each of its zone neighbor nodes and advertises its LSP containing the normal adjacencies between it and each of its zone neighbor nodes;
- o The zone leader purges the LS for the virtual node abstracted from the zone; and
- o Every zone node rolls back to normal.

The command may be replaced by the determination made by a zone node, such as the zone leader. After determining that all the conditions are met, it updates its LSP by setting OP = R in its Zone ID TLV, which is distributed to every zone node.

Condition 1 is met if it has its LSDB containing the link from each zone neighbor node to its zone edge node. That is that for every link from a zone neighbor node to the virtual node in the LSDB, there is a corresponding link from the zone neighbor to a zone edge node.

Condition 2 is met after Condition 1 has been met for a given time, such as maximum LSP advertisement time (MaxLSPAdvTime) crossing a network. We may assume that MaxLSPAdvTime is 5 seconds.

6. Operations

6.1. Configuring Zone

In general, a zone is a subset of an area and has a zone ID. It consists of some zone internal nodes and zone edge nodes. To configure it, a user configures this zone ID on every zone internal node and on every zone link of each zone edge node.

A node configured with the zone ID has all its links to be the zone links. The zone internal nodes and all their links plus the zone edge nodes and their zone links constitute the zone.

In a special case, a zone is an entire area and has a zone ID. All the links in the area are the zone links of the zone. To configure this zone, a user configures the zone ID on every zone node.

6.2. Transferring Zone to Node

Transferring a zone to a single virtual node smoothly may take a few steps or stages.

At first, a user configures the zone on every node of the zone.

After finishing the configuration of the zone, the user may issue a command, such as a CLI command, on a zone node, such as the zone leader, to trigger transferring the zone to the node. When receiving the command, the node distributes it to every zone node. After receiving it, every zone edge node, acting as a proxy of the virtual node, establishes a new adjacency between the virtual node and each of its zone neighbor nodes.

If automatic transferring zone to node is enabled, the user does not need to issue the command. A zone node, such as the zone leader, will distribute the "command" to every zone node after determining that the configuration of the zone has been finished.

Then, all the zone nodes, including the zone leader, zone edge nodes and zone internal nodes, work together to make the zone to appear as a single virtual node smoothly in a couple of steps.

6.3. Rolling back Node to Zone

After abstracting a zone to a single virtual node, we may want to roll back the node to the zone smoothly in some cases. The process of rolling back has a few steps or stages.

At first, a user issues a command, such as a CLI command, on a zone node, such as the zone leader, to start (or prepare) for roll back. When receiving the command, the node distributes it to every node in the zone. After receiving it, every zone edge node establishes a normal adjacency between the edge node and each of its zone neighbor nodes, and advertises the link state of the zone over the adjacency if it crosses the adjacency, but holds off its LSP containing the normal adjacency.

Second, a user may issue a command, such as a CLI command, on a zone node, such as the zone leader, to roll back from the virtual node to the zone if it is ready for roll back.

After receiving the command, the node distributes it to every node in the zone. After receiving it, all the zone nodes work together to roll back from the virtual node to the zone.

If automatic roll back Node to Zone is enabled, the user does not need to issue the command. A zone node, such as the zone leader, will distribute the "command" to every zone node after determining that it is ready for roll back.

7. Security Considerations

The mechanism described in this document does not raise any new security issues for the IS-IS protocols.

8. IANA Considerations

Under the registry name "IS-IS TLV Codepoints", IANA is requested to assign a new registry type for Zone ID as follows:

TLV Type	TLV Name	reference
TBD1	Zone ID	This document

IANA is requested to create a new sub-registry "Adjacent Node ID Sub-TLVs" on the IANA IS-IS TLV Codepoints web page as follows:

Type	Name	reference
0	Reserved	
1	Zone ISN	This document
2	Zone ESN	This document
3 - 255	Unassigned	

9. Contributors

Alvaro Retana
Futurewei
Raleigh, NC
USA

Email: alvaro.retana@futurewei.com

10. Acknowledgement

The authors would like to thank Acee Lindem, Abhay Roy, Christian Hopps, Dean Cheng, Russ White, Tony Przygienda, Wenhu Lu, Lin Han, Donald Eastlake, Tony Li, Robert Raszuk, Padmadevi Pillay Esnault, and Yang Yu for their valuable comments on TTZ.

11. References

11.1. Normative References

[I-D.ietf-lsr-dynamic-flooding]
Li, T., Psenak, P., Ginsberg, L., Chen, H., Przygienda, T., Cooper, D., Jalil, L., Dontula, S., and G. Mishra, "Dynamic Flooding on Dense Graphs", draft-ietf-lsr-dynamic-flooding-07 (work in progress), June 2020.

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-15 (work in progress), January 2018.

[ISO10589]
International Organization for Standardization, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Nov. 2002.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <https://www.rfc-editor.org/info/rfc1195>.

[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>.

[RFC5029] Vasseur, JP. and S. Previdi, "Definition of an IS-IS Link Attribute Sub-TLV", RFC 5029, DOI 10.17487/RFC5029, September 2007, <https://www.rfc-editor.org/info/rfc5029>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC7142] Shand, M. and L. Ginsberg, "Reclassification of RFC 1142 to Historic", RFC 7142, DOI 10.17487/RFC7142, February 2014, <https://www.rfc-editor.org/info/rfc7142>.

[RFC8099] Chen, H., Li, R., Retana, A., Yang, Y., and Z. Liu, "OSPF Topology-Transparent Zone", RFC 8099, DOI 10.17487/RFC8099, February 2017, <https://www.rfc-editor.org/info/rfc8099>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

11.2. Informative References

[Clos] Clos, C., "A Study of Non-Blocking Switching Networks", The Bell System Technical Journal Vol. 32(2), DOI 10.1002/j.1538-7305.1953.tb01433.x, March 1953, <http://dx.doi.org/10.1002/j.1538-7305.1953.tb01433.x>.

[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>.

Authors' Addresses

Huaimo Chen
Futurewei
Boston, MA
USA

Email: huaimo.chen@futurewei.com

Richard Li
Futurewei
2330 Central expressway
Santa Clara, CA
USA

Email: richard.li@futurewei.com

Yi Yang
IBM
Cary, NC
United States of America

Email: yyietf@gmail.com

Anil Kumar S N
RtBrick
Bangalore
India

Email: anil.ietf@gmail.com

Yanhe Fan
Casa Systems
USA

Email: yfan@casa-systems.com

Ning So
Plano, TX 75082
USA

Email: ningso01@gmail.com

Vic Liu
USA

Email: liu.cmri@gmail.com

Mehmet Toy
Verizon
USA

Email: mehmet.toy@verizon.com

Lei Liu
Fujitsu
USA

Email: liulei.kddi@gmail.com

Kiran Makhijani
Futurewei
USA

Email: kiranm@futurewei.com