

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: September 11, 2016

H. Chen
R. Li
Huawei Technologies
A. Retana
Y. Yang
Cisco Systems, Inc.
V. Liu
China Mobile
M. Toy
Comcast
March 10, 2016

OSPF Topology-Transparent Zone
draft-ietf-ospf-ttz-03.txt

Abstract

This document presents a topology-transparent zone in an OSPF area. A topology-transparent zone comprises a group of routers and a number of links connecting these routers. Any router outside of the zone is not aware of the zone. The information about the links and routers such as a link down inside the zone is not advertised to any router outside of the zone.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://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 September 11, 2016.

Copyright Notice

Copyright (c) 2016 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

(<http://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	4
2.	Terminology	4
3.	Conventions Used in This Document	5
4.	Requirements	5
5.	Topology-Transparent Zone	5
5.1.	Overview of Topology-Transparent Zone	5
5.2.	Some Details on TTZ	6
6.	Extensions to OSPF Protocols	7
6.1.	General Format of TTZ LSA	8
6.2.	TTZ ID TLV	8
6.3.	TTZ Router TLV	9
6.4.	TTZ Options TLV	10
6.5.	Link Scope TTZ LSA	11
7.	Constructing LSAs for TTZ	11
8.	Establishing Adjacencies	13
8.1.	Discover TTZ Neighbors	13
8.2.	Adjacency between TTZ Edge and TTZ External Router	14
9.	Advertisement of LSAs	14
9.1.	Advertisement of LSAs within TTZ	15
9.2.	Advertisement of LSAs through TTZ	15
10.	Computation of Routing Table	15
11.	Operations	16
11.1.	Configuring TTZ	16
11.2.	Smooth Migration to TTZ	16
11.3.	Adding a Router into TTZ	18
12.	Prototype Implementation	18
12.1.	What are Implemented and Tested	18
12.2.	Implementation Experience	20
13.	Security Considerations	20
14.	IANA Considerations	20
15.	Contributors and Other Authors	21
16.	Acknowledgement	22
17.	References	22
17.1.	Normative References	22
17.2.	Informative References	22
Appendix A.	Constants for LSA Advertisement	23

Authors' Addresses	23
------------------------------	--------------------

1. Introduction

The number of routers in a network becomes larger and larger as the network expands. For scalability and manageability, the network is reorganized into more areas when it becomes too big. However, this causes a number of issues.

At first, reorganizing a network from one area into multiple areas or from a number of existing areas into even more areas is a very challenging and time consuming task since it involves significant network architecture changes. Considering the one area case, originally the network has only one area, which is the backbone. This original backbone area will be reorganized into a new backbone and a number of non-backbone areas. In general, each of the non-backbone areas is connected to the new backbone area through the Area Border Routers (ABRs) between the non-backbone and the backbone area (refer to [RFC 2328 section 3](#)).

Secondly, the services carried by the network may be interrupted while the network is being reorganized from one area into multiple areas or from a number of existing areas into even more areas since every OSPF interface with an area change is going down with its old area and then up with a new area.

This document presents a topology-transparent zone (TTZ) in an OSPF area and describes extensions to OSPF for supporting the topology-transparent zone, which is scalable and resolves the issues above.

2. Terminology

TTZ internal link: a link between two TTZ adjacent routers in the same TTZ. A TTZ internal link is called a TTZ link in general.

TTZ internal router: a router in a TTZ whose adjacent routers are all in the same TTZ.

TTZ external router: a router outside of a TTZ without any TTZ internal link.

TTZ external link: a link between a TTZ edge router and a TTZ external router.

TTZ edge router: a router in a TTZ that has one (or more) adjacent routers which belong to the same TTZ, and one (or more) adjacent routers which do not belong to the TTZ.

TTZ router: a router in a TTZ, i.e., a TTZ internal router or a TTZ

edge router.

3. 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](#).

4. Requirements

Topology-Transparent Zone may be deployed to resolve some critical issues in existing networks and future networks. The requirements for TTZ are listed as follows:

- o Routers outside a TTZ MUST NOT require any changes to operate with the TTZ.
- o A TTZ MUST be enclosed in a single area.
- o A TTZ MUST hide the topology of the TTZ from any router outside of the TTZ.

5. Topology-Transparent Zone

5.1. Overview of Topology-Transparent Zone

A Topology-Transparent Zone is identified by an Identifier (ID), and it consists of a group of routers and a number of links connecting the routers. A TTZ MUST be contained within an OSPF area.


The ID of a TTZ or TTZ ID is a 32-bit number that is unique for identifying a TTZ. The ID SHOULD NOT be 0.

In addition to having the functions of an OSPF area, an OSPF TTZ makes some improvements on an OSPF area, which include:

- o An OSPF TTZ is virtualized as the TTZ edges connected each other.
- o An OSPF TTZ receives the link state information about the topology outside of the TTZ, stores the information in the TTZ and floods the information through the TTZ to the routers outside of the TTZ.

5.2. Some Details on TTZ

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



All the routers in the figure are in area X. Routers with T (i.e., T61, T63, T65, T67, T71, T73, T75, T77, T79 and T81) are also in TTZ 600, which contains the TTZ internal links connecting them. To create a TTZ, we need configure it (refer to [section 11](#)).

There are two types of routers in a TTZ: TTZ internal and TTZ edge routers. TTZ 600 has four TTZ edge routers T61, T63, T65 and T67. Each of them has at least one adjacent router in TTZ 600 and one adjacent router outside of TTZ 600. For instance, router T61 is a TTZ edge router since it has an adjacent router R15 outside of TTZ 600 and three adjacent routers T75, T71 and T81 in TTZ 600.


In addition, TTZ 600 comprises six TTZ internal routers T71, T73, T75, T77, T79 and T81. Each of them has all its adjacent routers in TTZ 600. For instance, router T71 is a TTZ internal router since its adjacent routers T61, T63, T65, T67 and T73 are all in TTZ 600. Note that none of the TTZ internal routers can be an ABR.

A TTZ hides the internal topology of the TTZ from the outside. It

TTZ. Some control information regarding a TTZ can also be contained and advertised in opaque LSAs within the TTZ. These opaque LSAs are called TTZ opaque LSAs or TTZ LSAs for short.

6.1. General Format of TTZ LSA

The following is the general format of a TTZ LSA. It has an LS Type = 10/9 and TTZ-LSA-Type, and contains a number of TLVs.



Where TTZ-LSA-Type is 9, the exact number is to be assigned by IANA.

There are three top-level TLVs defined: TTZ ID TLV, TTZ Router TLV, and TTZ Options TLV. A TTZ LSA of LS Type 10 contains a mandatory TTZ ID TLV, which is followed by a number of other top-level TLVs.

A TTZ LSA having a optional TTZ Router TLV is called a TTZ Router LSA. A TTZ LSA containing a TTZ Options TLV is called a TTZ Control LSA.

6.2. TTZ ID TLV

A TTZ ID TLV has the following format. It contains a TTZ ID and some flags.


```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   TTZ-ID-TLV-Type (1)   |   TLV-Length   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               TTZ ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Reserved (MUST be zero)   |E|Z|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
E = 1: Indicating a router is a TTZ Edge router
Z = 1: Indicating a router has migrated to TTZ

```

When a TTZ router originates a TTZ LSA containing a TTZ ID TLV, it sets flag E to 1 in the TTZ ID TLV if it is a TTZ edge router, and to 0 if it is a TTZ internal router. It sets flag Z to 1 after it has migrated to TTZ.

6.3. TTZ Router TLV

The format of a TTZ Router TLV is as follows. It contains the contents of a router LSA.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   TTZ-RT-TLV-Type (2)   |   TLV-Length   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   0   |V|E|B|           0   |   # links   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link Data               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   |   # TOS   |   metric   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               ...                               ~

```

For a router link, the existing eight bit Type field for a router link is split into two fields as follows:


```

    0   1   2   3   4   5   6   7
+---+---+---+---+---+---+---+---+
| I |           Type-1           |
+---+---+---+---+---+---+---+---+

```

I bit flag:

1: Router link is a TTZ internal link.

0: Router link is a TTZ external link.

Type-1: The kind of the link. The values for Type-1 are the same as those for Type defined in [RFC 2328 section 12.4.1](#).

6.4. TTZ Options TLV

The format of a TTZ Options TLV is as follows.

```

    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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      TTZ-OP-TLV-Type (3)      |      TLV-Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| T | M | N | R |               Reserved (MUST be zero)               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

T = 1: Advertising TTZ Topology Information for Migration

M = 1: Migrating to TTZ

N = 1: Advertising Normal Topology Information for Rollback

R = 1: Rolling back from TTZ

Flags T, M, N and R are exclusive. When one of them is set to 1, the others MUST be set to 0.

After a user configures a TTZ router to advertise TTZ topology information, the TTZ router originates a TTZ Control LSA having a TTZ Options TLV with flag T set to 1. It also originates its TTZ router LSA. When another TTZ router receives the LSA with T = 1, it originates its TTZ router LSA as needed.

After a user configures a TTZ router to migrate to TTZ, the TTZ router originates a TTZ Control LSA having a TTZ Options TLV with flag M set to 1 and migrates to TTZ. When another TTZ router receives the LSA with M = 1, it also migrates to TTZ.

After a user configures a TTZ router to advertise normal topology information, the TTZ router originates a TTZ Control LSA having a TTZ Options TLV with flag N set to 1. It also advertises its normal LSAs such as its normal router LSA and stops advertising its other TTZ LSAs. When another TTZ router receives the LSA with N = 1, it advertises its normal LSAs and stops advertising its TTZ LSAs.

After a user configures a TTZ router to roll back from TTZ, the TTZ router originates a TTZ Control LSA having a TTZ Options TLV with flag R set to 1 and rolls back from TTZ. When another TTZ router receives the LSA with R = 1, it also rolls back from TTZ.

After a TTZ router originates a TTZ control LSA in response to a configuration described above to control TTZ, it updates the TTZ control LSA accordingly if another configuration to control TTZ is issued on it.

6.5. Link Scope TTZ LSA

A TTZ LSA of LS Type 9 has the following format.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS age               |   Options   | LS Type = 9 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TTZ-LSA-Type(9)|               Instance ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Advertising Router               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS Sequence Number               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS checksum               |               Length               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               ~               TTZ ID TLV               ~               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|               ~               (Options TLV)               ~               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

It contains a mandatory TTZ ID TLV, which may be followed by a optional Options TLV. It is used to discover a TTZ neighbor.

7. Constructing LSAs for TTZ

The LSAs for representing a TTZ include TTZ router LSAs and normal router LSAs for virtualizing the TTZ.

A TTZ router LSA generated by a TTZ edge router has a TTZ ID TLV and a TTZ Router TLV. The former includes the ID of the TTZ to which the router belongs and flag E set to 1, which indicates the originator of the LSA is a TTZ Edge router. The latter contains the links attached

to the router.

A TTZ router LSA generated by a TTZ internal router has a TTZ ID TLV containing the ID of the TTZ to which the router belongs and flag E set to 0, which indicates the originator of the LSA is a TTZ internal router. The TTZ internal router generates the TTZ router LSA with just the TTZ ID TLV and its normal router LSA.

After receiving a trigger to migrate to TTZ such as a TTZ LSA with flag M = 1, a TTZ edge router originates its normal router LSA for virtualizing a TTZ, which comprises three groups of links in general.

The first group are the router links connecting the TTZ external routers. These router links are normal router links. There is a router link for every adjacency between this TTZ edge router and a TTZ external router.

The second group are the "virtual" router links connecting to the other TTZ edge routers. For each of the other TTZ edge routers, there is a corresponding point-to-point router link to it from this TTZ edge router. The cost of the link is the cost of the shortest path from this TTZ edge router to the other TTZ edge router within the TTZ.

In addition, the LSA may contain a third group of links, which are the stub links for the loopback addresses inside the TTZ to be accessed by nodes outside of the TTZ.

To migrate to a TTZ smoothly, a TTZ edge router virtualizes the TTZ in two steps. At first, the router updates its normal router LSA by adding a point-to-point link to each of the other edge routers of the TTZ and a stub link for each of the loopback addresses in the TTZ to be accessed outside of the TTZ into the LSA. And then it removes the links configured as TTZ links from its updated router LSA after sending its updated router LSA and receiving the updated router LSAs originated by the other TTZ edge routers for MaxLSAAdvTime or after sending its updated router LSA for MaxLSAGenAdvTime (refer to [Appendix A](#)).

To roll back from a TTZ smoothly after receiving a trigger to roll back from TTZ, a TTZ edge router updates its normal router LSA in the above two steps in a reverse way. At first, it updates its normal router LSA by adding the normal links for the links configured as TTZ links into the LSA. And then it removes the point-to-point links to the other edge routers of the TTZ for virtualizing the TTZ and the stub links for the loopback addresses from its updated router LSA after sending its updated router LSA and receiving the updated router LSAs originated by the other TTZ edge routers for MaxLSAAdvTime or

after sending its updated router LSA for MaxLSAGenAdvTime.

8. Establishing Adjacencies

This section describes the adjacencies in different cases.

8.1. Discover TTZ Neighbors

For two routers A and B connected by a P2P link and having a normal adjacency, they discover TTZ each other through a TTZ LSA of LS Type 9 with a TTZ ID TLV. We call this LSA D-LSA for short.

If two ends of the link have different TTZ IDs, no TTZ adjacency over the link will be "formed".

If two ends of the link have the same TTZ ID and Z flag value, A and B are TTZ neighbors. The following is a sequence of events related to TTZ for this case.

A		B
Configure TTZ		Configure TTZ
	D-LSA (TTZ-ID=100)	
	----->	Same TTZ ID and Z
		A is B's TTZ Neighbor
	D-LSA (TTZ-ID=100)	
Same TTZ ID and Z <-----		
B is A's TTZ Neighbor		

A sends B a D-LSA with TTZ-ID after the TTZ is configured on it. B sends A a D-LSA with TTZ-ID after the TTZ is configured on it.

When A receives the D-LSA from B and determines they have the same TTZ ID and Z flag value, B is A's TTZ neighbor. A also sends B all the TTZ LSAs it has and originates its TTZ router LSA if (Z==0 and there is a TTZ LSA with T=1) OR Z==1.

B is symmetric to A and acts similarly to A.

If two ends of the link have the same TTZ ID but Z flags are different, a TTZ adjacency over the link is "formed" in the following steps. Suppose that A has migrated to TTZ and B has not (i.e., flag Z in A's D-LSA is 1 and flag Z in B's D-LSA is 0).

When A receives the D-LSA from B and determines they have the same TTZ ID but its Z=1 and B's Z=0, A sends B all the TTZ LSAs it has and triggers B to migrate to TTZ. A updates and sends B its D-LSA by

adding an Options TLV with M=1 after sending B all the TTZ LSAs.

When B receives the D-LSA from A and determines they have the same TTZ ID but its Z=0 and A's Z=1, B sends A all the TTZ LSAs it has and starts to migrate to TTZ after receiving A's D-LSA with M=1. After migration to TTZ, B updates and advertises its LSAs as needed.

After receiving B's D-LSA with Z=1, A updates and sends B its D-LSA by removing the Options TLV. It also updates and advertises its LSAs as needed.

For a number of routers connected through a broadcast link and having normal adjacencies among them, they also discover TTZ each other through D-LSAs. The DR for the link "forms" TTZ adjacencies with the other routers if all the routers attached to the link have the same TTZ ID configured on the connections to the link. Otherwise, the DR does not "form" any TTZ adjacency with any router attached to the link.

For a number of routers connected through a broadcast link and having TTZ adjacencies among them, if a mis-configured router is introduced on the broadcast link, the DR for the link will not "form" any TTZ adjacency with this mis-configured router.

For routers connected via a link without any adjacency among them, they discover TTZ each other through D-LSAs in the same way as described above after they form a normal adjacency.

8.2. Adjacency between TTZ Edge and TTZ External Router

A TTZ edge router forms an adjacency with any TTZ external router to which it is connected.

When the TTZ edge router synchronizes its link state database with the TTZ external router, it sends the TTZ external router the information about all the LSAs except for the LSAs belonging to the TTZ that are hidden from any router outside of the TTZ.

At the end of the link state database synchronization, the TTZ edge router originates its own router LSA for virtualizing the TTZ and sends this LSA to its adjacent routers including the TTZ external router.

9. Advertisement of LSAs

LSAs can be divided into a couple of classes according to their Advertisements. The first class of LSAs is advertised within a TTZ.

The second is advertised through a TTZ.

9.1. Advertisement of LSAs within TTZ

Any LSA about a link state in a TTZ is advertised within the TTZ. It is not advertised to any router outside of the TTZ. For example, a router LSA generated for a router in a TTZ is advertised within the TTZ.

Any network LSA generated for a broadcast or NBMA network in a TTZ is advertised within the TTZ.

Any opaque LSA generated for a TTZ internal TE link is advertised within the TTZ.

After migrating to TTZ, every edge router of a TTZ **MUST NOT** advertise any LSA about a link state in the TTZ to any router outside of the TTZ.

For any TTZ LSA originated by a router within the TTZ, every edge router of the TTZ **MUST NOT** advertise it to any router outside of the TTZ.

9.2. Advertisement of LSAs through TTZ

Any LSA about a link state outside of a TTZ received by an edge router of the TTZ is advertised using the TTZ as transit. For example, when an edge router of a TTZ receives an LSA from a router outside of the TTZ, it floods it to its neighboring routers both inside the TTZ and outside of the TTZ. This LSA may be any LSA such as a router LSA that is advertised within an OSPF area.

The routers in the TTZ continue to flood the LSA. When another edge router of the TTZ receives the LSA, it floods the LSA to its neighboring routers both outside of the TTZ and inside the TTZ.

10. Computation of Routing Table

After a router migrates to TTZ, the computation of the routing table on the router is the same as that described in [RFC 2328 section 16](#) with one exception. The router in a TTZ ignores the router LSAs generated by the TTZ edge routers for virtualizing the TTZ. This can be achieved by adding a flag into every link stored in LSDB and setting this flag to 1 in every link in these router LSAs, which indicates that the link is unusable. It computes routes within the TTZ topology and the topology outside of the TTZ without using any unusable links.

11. Operations

11.1. Configuring TTZ

This section proposes some options for configuring a TTZ.

1. Configuring TTZ on Every Link in TTZ

If every link in a TTZ is configured with a same TTZ ID as a TTZ link, the TTZ is determined. A router with some TTZ links and some normal links is a TTZ edge router. A router with only TTZ links is a TTZ internal router.

2. Configuring TTZ on Every Router in TTZ

A same TTZ ID is configured on every router in a TTZ, and on every TTZ edge router's links connecting to the routers in the TTZ.

A router configured with the TTZ ID on some of its links is a TTZ edge router. A router configured with the TTZ ID only is a TTZ internal router. All the links on a TTZ internal router are TTZ links. This option is simpler than the above one.

11.2. Smooth Migration to TTZ

For a group of routers and a number of links connecting the routers in an area, making them transfer to work as a TTZ without any service interruption takes a few of steps or stages.

At first, a user configures the TTZ feature on every router in the TTZ. In this stage, a router does not originate its TTZ router LSA. It will discover its TTZ neighbors.

Secondly, after configuring the TTZ, a user issues a CLI command on one router in the TTZ, which triggers every router in the TTZ to generate and advertise TTZ information among the routers in the TTZ. When the router receives the command, it originates a TTZ control LSA with T=1 (indicating TTZ information generation and advertisement for migration). It also originates its TTZ router LSA, and advertises the LSA to its TTZ neighbors. When another router in the TTZ receives the LSA with T=1, it originates its TTZ router LSA. In this stage, every router in the TTZ has dual roles. One is to function as a normal router. The other is to generate and advertise TTZ information.

Thirdly, a user checks whether a router in the TTZ is ready for migration to TTZ. A router in the TTZ is ready after it has received all the necessary information from all the routers in the TTZ. This

information may be displayed on a router through a CLI command.

And then a user activates the TTZ through using a CLI command such as migrate to TTZ on one router in the TTZ. The router migrates to TTZ, generates and advertises a TTZ control LSA with $M = 1$ (indicating Migrating to TTZ) after it receives the command. After another router in the TTZ receives the TTZ control LSA with $M = 1$, it also migrates to TTZ. Thus, activating the TTZ on one TTZ router propagates to every router in the TTZ, which migrates to TTZ.

For an edge router of the TTZ, migrating to work as a TTZ router comprises generating a router LSA to virtualize the TTZ and flooding this LSA to all its neighboring routers in two steps as described in [section 7](#).

In normal operations for migration to TTZ and rollback from TTZ, a user issues a series of commands according to certain procedures. In an abnormal case, for example two conflicting commands are issued on two TTZ routers in a TTZ at the same time, a TTZ router issues an error and logs the error when it detects a conflict.

A conflicting command may be detected on a router on which the command is issued. Thus some abnormal cases may be prevented. When a command for migration/rollback is issued on a router, the router checks whether it is in a correct sequence of commands for migration/rollback through using the information it has. For migrating a part of an area to a TTZ, the correct sequence of commands is as follows in general:

- 1) configure TTZ on every router in the part of the area to be migrated to TTZ;
- 2) configure on one router in the TTZ to trigger every router in the TTZ to generate and advertise TTZ information for migration; and
- 3) configure on one router in the TTZ to trigger every router in the TTZ to migrate to TTZ.

For rolling back from TTZ, it is similar.

After receiving a command on a router to migrate to TTZ, which is for 3), the router checks whether 2) is performed through checking if it has received/originated TTZ LSAs. If it has not, it issues an error to an operator (generation and advertisement of TTZ information for migration to TTZ is not done yet) and rejects the command at this time.

After a router receives a TTZ LSA with $M=1$ for 3) from another

router, it checks whether 2) is performed through checking if it has received/originated TTZ LSAs. If it has not, it issues an error and logs the error. The operation for migration will continue.

After receiving a command on a router to generate and advertise TTZ information, which is for 2), the router checks whether 1) is performed through checking if TTZ is configured on it. If it is not, it issues an error to an operator (TTZ is not configured on it yet) and rejects the command at this time.

11.3. Adding a Router into TTZ

When a non TTZ router (say R1) is connected via a P2P link to a TTZ router (say T1) working as TTZ and there is a normal adjacency between them over the link, a user can configure TTZ on two ends of the link to add R1 into the TTZ to which T1 belongs. They discover TTZ each other with the TTZ as described in [section 8](#).

When a number of non TTZ routers are connected via a broadcast link to a TTZ router (say T1) working as TTZ and there are normal adjacencies among them, a user configures TTZ on the connection to the link on every router to add the non TTZ routers into the TTZ to which T1 belongs. The DR for the link "forms" TTZ adjacencies with the other routers connected to the link if they all have the same TTZ ID configured for the link. This is determined through the discovery process described in [section 8](#).

When a router (say R1) is connected via a P2P link to a TTZ router (say T1) and there is not any adjacency between them over the link, a user can configure TTZ on two ends of the link to add R1 into the TTZ to which T1 belongs. R1 and T1 will form an adjacency in the same way as described in [section 8](#).

12. Prototype Implementation

12.1. What are Implemented and Tested

1. CLI Commands for TTZ

The CLIs implemented and tested include:

- o the CLIs of the simpler option for configuring TTZ, and
- o the CLIs for controlling migration to TTZ.

2. Extensions to OSPF Protocols for TTZ

All the extensions defined in section "Extensions to OSPF Protocols" are implemented and tested except for rolling back from TTZ. The testing results illustrate:

- o A TTZ is virtualized to outside as its edge routers connected each other. Any router outside of the TTZ sees the edge routers (as normal routers) connecting each other and to some other routers.
- o The link state information about the routers and links inside the TTZ is contained within the TTZ. It is not advertised to any router outside of the TTZ.
- o TTZ is transparent. From a router inside a TTZ, it sees the topology (link state) outside of the TTZ. From a router outside of the TTZ, it sees the topology beyond the TTZ. The link state information outside of the TTZ is advertised through the TTZ.
- o TTZ is backward compatible. Any router outside of a TTZ does not need to support or know TTZ.

3. Smooth Migration to TTZ

The procedures and related protocol extensions for smooth migration to TTZ are implemented and tested. The testing results show:

- o A part of an OSPF area is smoothly migrated to a TTZ without any routing disruptions. The routes on every router are stable while the part of the area is being migrated to the TTZ.
- o Migration to TTZ is very easy to operate.

4. Add a Router to TTZ

Adding a router into TTZ is implemented and tested. The testing results illustrate:

- o A router can be easily added into a TTZ and becomes a TTZ router.
- o The router added into the TTZ is not seen on any router outside of the TTZ, but it is a part of the TTZ.

5. Leak TTZ Loopbacks Outside

Leaking loopback addresses in a TTZ to routers outside of the TTZ is implemented and tested. The testing results illustrate:

- o The loopback addresses inside the TTZ are advertised to the routers outside of the TTZ.
- o The loopback addresses are accessible from a router outside of the TTZ.

12.2. Implementation Experience

The implementation of TTZ is relatively easy compared to other features of OSPF. Re-using the existing OSPF code along with additional simple logic does the work. A couple of engineers started to work on implementing the TTZ from the middle of June, 2014 and finished coding it just before IETF 90. After some testing and bug fixes, it works as expected.

In our implementation, the link state information in a TTZ opaque LSA is stored in the same link state database as the link state information in a normal LSA. For each TTZ link in the TTZ opaque LSA, there is an additional flag, which is used to differentiate between a TTZ link and a Normal link.

Before migration to TTZ, every router in the TTZ computes its routing table using the normal links. After migration to TTZ, every router in the TTZ computes its routing table using the TTZ links and normal links. In the case where both the TTZ link and the normal link exist, the TTZ link is used.

13. Security Considerations

The mechanism described in this document does not raise any new security issues for the OSPF protocols since a TTZ is enclosed in a single area.

14. IANA Considerations

Under Registry Name: Opaque Link-State Advertisements (LSA) Option Types [[RFC5250](#)], IANA is requested to assign a new Opaque type registry value for Topology-Transparent Zone (TTZ) LSA as follows:

+=====+=====+=====+			
Registry Value	Opaque Type	reference	
+=====+=====+=====+			
IANA TBD	TTZ LSA	This document	
(9 Suggested)			
+-----+-----+-----+			

IANA is requested to assign Types for new TLVs in the new TTZ LSA as follows:

Type	Name	Allowed in
1	TTZ ID TLV	TTZ LSA of LS Type 10 and 9
2	TTZ Router TLV	TTZ LSA of LS Type 10
3	TTZ Options TLV	TTZ LSA of LS Type 10 and 9

15. Contributors and Other Authors

1. Other Authors

Gregory Cauchie
FRANCE
Email: greg.cauchie@gmail.com

Anil Kumar S N
Huawei Technologies
Bangalore
India
Email: anil.sn@huawei.com

Ning So
Tata Communications
2613 Fairbourne Cir.
Plano, TX 75082
USA
Email: ningso01@gmail.com

Lei Liu
Fujitsu
USA
Email: lliu@us.fujitsu.com

2. Contributors

Veerendranatha Reddy Vallem
Huawei Technologies
Bangalore
India

Email: veerendranatharv@huawei.com

William McCall
Rightside Co.
Kirkland, WA
USA
will.mccall@rightside.co

16. Acknowledgement

The authors would like to thank Acee Lindem, Abhay Roy, Dean Cheng, Russ White, Tony Przygienda, Wenhui Lu, Lin Han, Kiran Makhijani, Padmadevi Pillay Esnault and Yang Yu for their valuable comments on this draft.

17. References

17.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/[RFC2119](#), March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/[RFC2328](#), April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.

17.2. Informative References

- [RFC5441] Vasseur, JP., Ed., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", [RFC 5441](#), DOI 10.17487/RFC5441, April 2009, <<http://www.rfc-editor.org/info/rfc5441>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#),

DOI 10.17487/RFC5440, March 2009,
<<http://www.rfc-editor.org/info/rfc5440>>.

Appendix A. Constants for LSA Advertisement

MaxLSAAdvTime: The maximum time for an LSA to be advertised to all the routers in an area. The value of MaxLSAAdvTime is set to 0.1 second.

MaxLSAGenAdvTime: The maximum time for all TTZ router LSAs to be generated by all TTZ edge routers and advertised to all the routers in an area after a first TTZ router LSA is generated. The value of MaxLSAGenAdvTime is set to 0.3 second.

Authors' Addresses

Huaimo Chen
Huawei Technologies
Boston, MA
USA

Email: huaimo.chen@huawei.com

Renwei Li
Huawei Technologies
2330 Central expressway
Santa Clara, CA
USA

Email: renwei.li@huawei.com

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Rd.
Raleigh, NC 27709
USA

Email: aretana@cisco.com

Yi Yang
Cisco Systems, Inc.
7025 Kit Creek Rd.
Raleigh, NC 27709
USA

Email: yiya@cisco.com

Vic Liu
China Mobile
No.32 Xuanwumen West Street, Xicheng District
Beijing, 100053
China

Email: liuzhiheng@chinamobile.com

Mehmet Toy
Comcast
1800 Bishops Gate Blvd.
Mount Laurel, NJ 08054
USA

Email: mehmet_toy@comcast.com

