

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: June 25, 2016

H. Chen
R. Li
Huawei Technologies
A. Retana
Y. Yang
Cisco Systems, Inc.
V. Liu
China Mobile
M. Toy
Comcast
December 23, 2015

OSPF Topology-Transparent Zone
draft-ietf-ospf-ttz-02.txt

Abstract

This document presents a topology-transparent zone in a domain. 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 inside the zone is not distributed to any router outside of the zone. Any link state change such as a link down inside the zone is not seen by 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 June 25, 2016.

Copyright Notice

Copyright (c) 2015 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	3
2.	Terminology	4
3.	Conventions Used in This Document	4
4.	Requirements	4
5.	Topology-Transparent Zone	5
5.1.	Overview of Topology-Transparent Zone	5
5.2.	Some Details on TTZ	5
6.	Extensions to OSPF Protocols	6
7.	Constructing LSAs for TTZ	10
8.	Establishing Adjacencies	11
8.1.	Discover TTZ Neighbor over Normal Adjacency	11
8.2.	Establishing TTZ Adjacencies	12
8.3.	Adjacency between TTZ Edge and TTZ External Router	12
9.	Distribution of LSAs	13
9.1.	Distribution of LSAs within TTZ	13
9.2.	Distribution of LSAs through TTZ	13
10.	Computation of Routing Table	13
11.	Operations	14
11.1.	Configuring TTZ	14
11.2.	Smooth Migration to TTZ	14
11.3.	Adding a Router into TTZ	15
12.	Prototype Implementation	15
12.1.	What are Implemented and Tested	16
12.2.	Implementation Experience	17
13.	Security Considerations	17
14.	IANA Considerations	17
15.	Contributors and Other Authors	18
16.	Acknowledgement	19
17.	References	19
17.1.	Normative References	19
17.2.	Informative References	20
Appendix A.	Constants for LSA Distribution	20
	Authors' Addresses	20

1. Introduction

The number of routers in a network becomes larger and larger as the Internet traffic keeps growing. Through splitting the network into multiple areas, we can extend the network further. However, there are a number of issues when a network is split further into more areas.

At first, dividing a network from one area into multiple areas or from a number of existing areas to even more areas is a very challenging and time consuming task since it is involved in 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 split 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 between the non backbone and the backbone area. There is not any direct connection between any two non backbone areas. Each area border router summarizes the topology of its attached non backbone area for transmission on the backbone area, and hence to all other area border routers.

Secondly, the services carried by the network may be interrupted while the network is being split from one area into multiple areas or from a number of existing areas into even more areas.

Furthermore, it is complex for a Multi-Protocol Label Switching (MPLS) Traffic Engineering (TE) Label Switching Path (LSP) crossing multiple areas to be setup. In one option, a TE path crossing multiple areas is computed by using collaborating Path Computation Elements (PCEs) [[RFC5441](#)] through the PCE Communication Protocol (PCEP)[[RFC5440](#)], which is not easy to configure by operators since the manual configuration of the sequence of domains is required. Although this issue can be addressed by using the Hierarchical PCE, this solution may further increase the complexity of network design. Especially, the current PCE standard method may not guarantee that the path found is optimal.

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

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 inside the zone is not distributed to any router outside of the zone. Any link state change such as a link down inside the zone is not seen by any router outside of the zone.

2. Terminology

TTZ internal link: a link between two TTZ adjacent routers in a same TTZ.

TTZ external link: a link between a TTZ edge router and a router outside of the TTZ.

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

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

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.


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 (TTZ) may be deployed for resolving 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 in an area.
- o A TTZ MUST hide the topology of the TTZ from any router outside of the TTZ.
- o Users SHOULD be able to easily set up an end to end service crossing TTZs.
- o The configuration for a TTZ in a network SHOULD be minimum.
- o The changes on the existing protocols for supporting TTZ SHOULD be minimum.



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 links connecting them.

There are two types of routers in a TTZ: TTZ internal and TTZ edge routers. TTZ 600 has four TTZ edge routers R61, R63, R65 and R67. Each TTZ edge router is connected to at least one router outside of TTZ 600. For instance, router R61 is a TTZ edge router since it is connected to router R15, which is outside of TTZ 600.

In addition, TTZ 600 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 TTZ 600. It is just connected to routers R61, R63, R65, R67 and R73 in the TTZ.

A TTZ hides the internal topology of the TTZ from the outside. It does not directly distribute any internal information about the TTZ to a router outside of the TTZ.

For instance, TTZ 600 does not send the information about TTZ internal router R71 to any router outside of TTZ 600; it does not send the information about the link between TTZ router R61 and R65 to any router outside of TTZ 600.

In order to create a TTZ, we may assign the same TTZ ID to the TTZ edge routers and identify the TTZ internal links on them through CLI. In addition, we assign the TTZ ID to every TTZ internal router which indicates that every link of the router is a TTZ internal link.

From a router outside of the TTZ, a TTZ is seen as a group of routers fully connected. For instance, router R15 in the figure above, which is outside of TTZ 600, sees TTZ 600 as a group of TTZ edge routers: R61, R63, R65 and R67, which are fully connected.

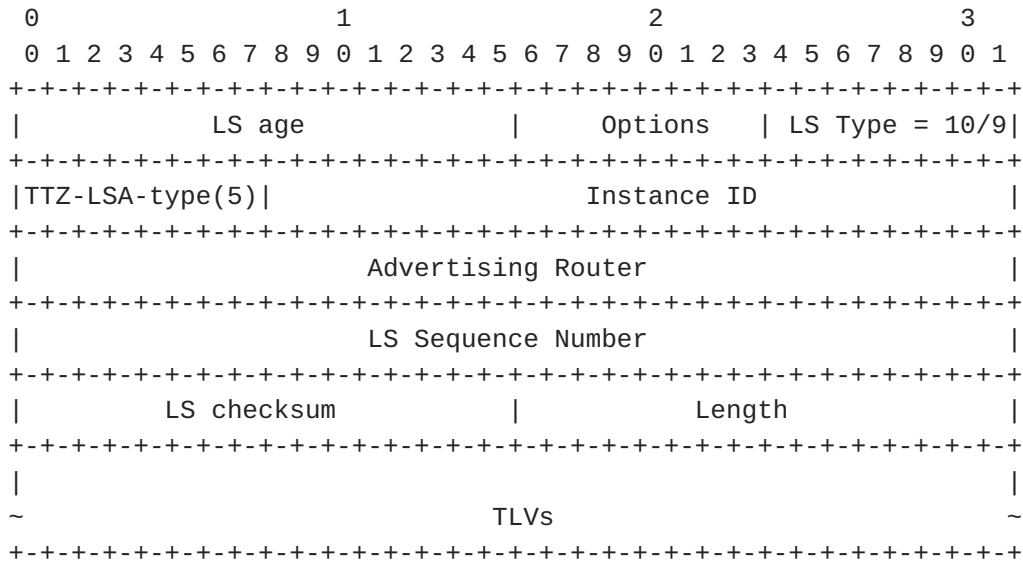
In addition, a router outside of the TTZ sees TTZ edge routers having normal connections to the routers outside of the TTZ. For example, router R15 sees four TTZ edge routers R61, R63, R65 and R67, which have the normal connections to R15, R29, R17 and R23, R25 and R31 respectively.

6. Extensions to OSPF Protocols

The link state information about a TTZ includes router LSAs and network LSAs. Router LSAs can be contained and distributed in opaque LSAs [[RFC5250](#)] within the TTZ. Some control information on a TTZ can also be contained and distributed in opaque LSAs within the TTZ.

These opaque LSAs are called TTZ opaque LSAs or TTZ LSAs for short.

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

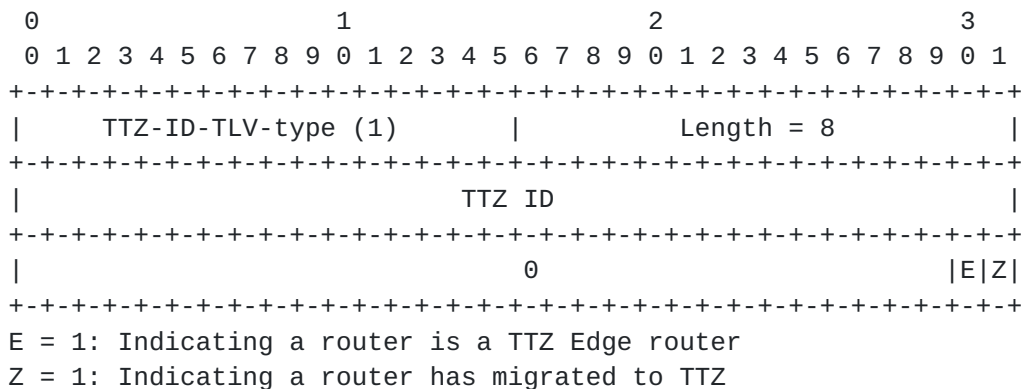


Where TTZ-LSA-type is 5, the exact number for it is to be assigned by IANA.

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 one of the other top-level TLVs.

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

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



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.

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.

For a link inside a TTZ, I bit flag is set to one, indicating that this link is a TTZ internal link. For a link connecting to a router outside of a TTZ from a TTZ edge router, I bit flag is set to zero, indicating that this link is a TTZ external link.

The value of Type-1 is 1, 2, 3, or 4, which indicates that the kind of a link being described is a point-to-point connection to another router, a connection to a transit network, a connection to a stub network, or a virtual link respectively.

The format of 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)      |      Length = 4      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|T|M|N|R|      0      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
T = 1: Distributing TTZ Topology Information for Migration
M = 1: Migrating to TTZ
N = 1: Distributing 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 triggers a TTZ router in a way such as issuing a CLI command to distribute 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 asks a TTZ router in a way such as a CLI command 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 triggers a TTZ router to distribute normal topology information, the TTZ router originates a TTZ Control LSA having a TTZ Options TLV with flag N set to 1. It also distributes its normal LSAs such as its normal router LSA and stops distributing its other TTZ LSAs. When another TTZ router receives the LSA with N = 1, it distributes its normal LSAs and stops distributing its TTZ LSAs.

After a user asks 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.

If two CLI trigger commands are issued on a TTZ router, two TTZ Control LSAs having a same instance ID are originated by the TTZ router.

A TTZ LSA of LS Type 9 with a TTZ ID TLV below is used to discover a TTZ neighbor.


```

      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(5)|          Instance ID          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Advertising Router          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          LS Sequence Number          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          LS checksum          |          Length          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
~                      TTZ ID TLV                      ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

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 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. This LSA implies that all the links connected to the router are TTZ internal links. In other words, all the router links in the normal router LSA generated by the TTZ internal router are TTZ internal links. Thus a TTZ internal router can choose one of two options: 1) generating a TTZ router LSA with a TTZ ID TLV and a TTZ Router TLV and 2) generating its normal router LSA and a TTZ router LSA with just a TTZ ID TLV. These two options are compatible.

After receiving a trigger to migrate to TTZ such as a TTZ Control LSA with flag M = 1, a TTZ edge router originates a 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. For each of the other TTZ edge routers, there is a point-to-point router link to it. The cost of the link is the cost of the shortest path from this TTZ edge router to it within the TTZ.

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

To migrate to 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 leaked. 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 MaxLSADistTime or after sending its updated router LSA for MaxLSAGenDistTime.

8. Establishing Adjacencies

This section describes the adjacencies in some different cases.

8.1. Discover TTZ Neighbor over Normal Adjacency

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 the same TTZ ID, A and B are TTZ neighbors. The following is a sequence of events related to TTZ.

A		B
Configure TTZ		Configure TTZ
	D-LSA (TTZ-ID=100)	
	----->	Same TTZ ID
		A is B's TTZ Neighbor
	D-LSA (TTZ-ID=100)	
Same TTZ ID <-----		
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, B is A's TTZ neighbor. When B receives the D-LSA from A and

determines they have the same TTZ ID, A is B's TTZ neighbor.

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 adjacency with each of 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.

8.2. Establishing TTZ Adjacencies

When a router (say A) is connected via a P2P link to another router (say B) and there is not any adjacency between them over the link, a user configures TTZ on two ends of the link to form a TTZ adjacency. They start to discover TTZ each other through D-LSAs in the same way as described above after the normal adjacency is formed.

For a number of routers connected through a broadcast link and having no adjacency among them, they start to form TTZ adjacencies after TTZ is configured on the link. While forming adjacencies, they discover TTZ each other through D-LSAs in the same way as described above after the normal adjacency is formed.

8.3. Adjacency between TTZ Edge and TTZ External Router

For an edge router in a TTZ, it forms an adjacency with any TTZ external router that has a connection with it.

When the 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 edge router originates its own router LSA for virtualizing the TTZ and sends this LSA to the TTZ external router.

From the point of view of the TTZ external router, which is outside of the TTZ, it sees the other end as a normal router and forms the adjacency in the same way as a normal router. It is not aware of anything about its neighboring TTZ. From the LSAs related to the TTZ edge router in the other end, it knows that the TTZ edge router is connected to each of the other TTZ edge routers and some routers outside of the TTZ.

9. Distribution of LSAs

LSAs can be divided into a couple of classes according to their distributions. The first class of LSAs is distributed within a TTZ. The second is distributed through a TTZ.

9.1. Distribution of LSAs within TTZ

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

Any network LSA generated for a broadcast or NBMA network in a TTZ is distributed in the TTZ and not sent to a router outside of the TTZ.

Any opaque LSA generated for a TTZ internal TE link is distributed within the TTZ and not distributed to any router outside of the TTZ.

After migrating to TTZ, every edge router of a TTZ MUST NOT distribute 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 distribute it to any router outside of the TTZ.

9.2. Distribution of LSAs through TTZ

Any LSA about a link state outside of a TTZ received by an edge router of the TTZ is distributed through the TTZ. 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 distributed in a domain.

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

When a router migrates to TTZ, the computation of the routing table on the router is the same as that described in [RFC 2328](#), with one exception. The router in a TTZ ignores the router LSAs generated by the TTZ edge routers for virtualizing the TTZ. It computes routes through using the TTZ topology and the topology outside of the TTZ.

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

We may configure a same TTZ ID on every router in the TTZ, and on every 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 distribute 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 distribution for migration). It also originates its TTZ router LSA, and distributes the LSA to its TTZ neighbors. When a 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 distribute TTZ information.

Thirdly, a user checks whether every 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 distributes a TTZ control LSA with M = 1 (indicating Migrating to TTZ) after it receives the command.

After a 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 makes every router in the TTZ migrate to TTZ, which computes routes through using the TTZ topology and the topology outside of the 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 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 conflict commands are issued on two TTZ routers in a TTZ at the same time, a TTZ router gives a warning when it detects the abnormal case.

[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 in the same way 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 adjacency with each of the other routers if all the routers have the same TTZ ID configured on the connections to the link.

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 fully connected. 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 distributed 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 distributed 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 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 distributed 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 stored, 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 that there are one TTZ link and one normal link to select, 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.

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
5	TTZ LSA	This document

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

15. Contributors and Other Authors

1. Other Authors

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

Anil Kumar S N
Huawei Technologies
Banglore
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
Banglore
India
Email: veerendranatharv@huawei.com

William McCall
Cisco Systems, Inc.
Bellevue, WA
USA
wimccall@cisco.com

16. Acknowledgement

The author would like to thank Acee Lindem, Abhay Roy, Dean Cheng, Russ White, Tony Przygienda, Wenhui Lu, Lin Han 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>>.
- [RFC5613] Zinin, A., Roy, A., Nguyen, L., Friedman, B., and D. Yeung, "OSPF Link-Local Signaling", [RFC 5613](#), DOI 10.17487/RFC5613, August 2009, <<http://www.rfc-editor.org/info/rfc5613>>.

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 Distribution

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

MaxLSAGenDistTime: The maximum time for all TTZ router LSAs to be generated by all TTZ edge routers and distributed to all the routers in an area after a first TTZ router LSA is generated. The value of MaxLSAGenDistTime 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@cable.comcast.com

