

Internet Engineering Task Force (IETF)
Request for Comments: 8099
Category: Experimental
ISSN: 2070-1721

H. Chen
R. Li
Huawei Technologies
A. Retana
Cisco Systems, Inc.
Y. Yang
Sokrate
Z. Liu
China Mobile
February 2017

OSPF Topology-Transparent Zone

Abstract

This document presents a Topology-Transparent Zone (TTZ) in an OSPF area. A TTZ 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. A TTZ hides the internal topology of the TTZ from the outside. It does not directly advertise any internal information about the TTZ to a router outside of the TTZ. The information about the links and routers such as a link down inside the TTZ is not advertised to any router outside of the TTZ.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8099>.

Copyright Notice

Copyright (c) 2017 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. TTZ Example	6
6. Extensions to OSPF Protocols	8
6.1. General Format of TTZ LSA	8
6.2. TTZ ID TLV	9
6.3. TTZ Router TLV	9
6.4. TTZ Options TLV	10
6.5. Link Scope TTZ LSA	12
7. Constructing LSAs for TTZ	12
7.1. TTZ Migration Process	13
8. Establishing Adjacencies	14
8.1. Discovery of TTZ Neighbors	14
8.2. Adjacency between TTZ Edge and TTZ-External Router	17
9. Advertisement of LSAs	17
9.1. Advertisement of LSAs within TTZ	17
9.2. Advertisement of LSAs through TTZ	18
10. Computation of Routing Table	18
11. Operations	18
11.1. Configuring TTZ	18
11.2. Migration to TTZ	19
11.3. Adding a Router into TTZ	21
12. Manageability Considerations	22
13. Security Considerations	22
14. IANA Considerations	22
15. References	23
15.1. Normative References	23
15.2. Informative References	23
Appendix A. Prototype Implementation	24
A.1. What Is Implemented and Tested	24
A.2. Implementation Experience	25
Acknowledgements	26
Contributors	26
Authors' Addresses	27

1. Introduction

Networks expand as business grows and traffic increases. For scalability and manageability, a hierarchical network architecture is usually deployed in OSPF networks by regrouping routers into areas, which is often challenging and causes service interruptions.

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). It demands careful redesigning of network topology in advance to guarantee backbone area continuity and non-backbone-area reachability, and it requires significant modifications of configurations on many routers to ensure consistent routing.

Second, 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 OSPFv2 for supporting the TTZ, which is scalable and resolves the issues above. A TTZ hides the internal topology of the TTZ from the outside. It does not directly advertise any internal information about the TTZ to any router outside of the TTZ.

2. Terminology

TTZ link or TTZ-internal link:

A link whose ends are within a single TTZ.

TTZ-internal router:

A router whose links are TTZ-internal links inside a single TTZ.

TTZ-external router:

A router outside of a TTZ that has no TTZ-internal links.

TTZ-external link:

A link not configured to be within a TTZ.

TTZ edge router:

A router is called a TTZ edge router if some, but not all, of its links are within a single TTZ.

TTZ router:

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 [RFC2119].

4. Requirements

A Topology-Transparent Zone may be deployed to resolve some critical issues in existing networks and future networks. The requirements for a 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 a TTZ 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.

A TTZ ID is a 32-bit number that is unique for identifying a TTZ. The TTZ ID SHOULD NOT be 0 in order to avoid confusion with Area 0. The same TTZ ID MUST be configured on the routers and/or links that make up a specific instance of a TTZ. All TTZ instances in an OSPF area MUST be unique.

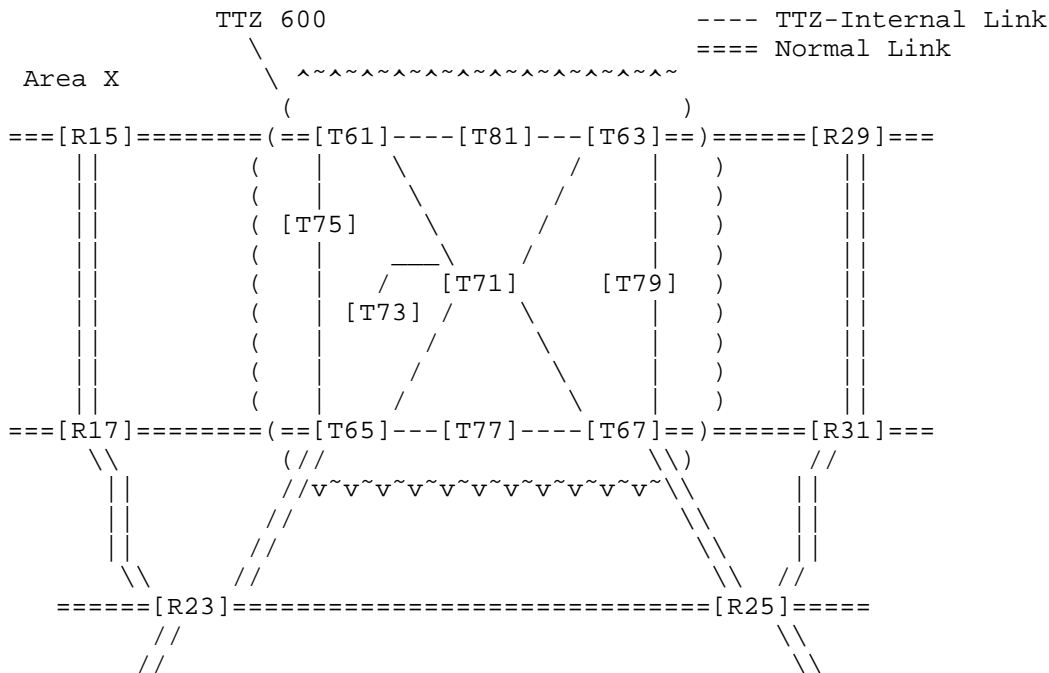
In addition to having similar functions of an OSPF area, an OSPF TTZ makes some improvements on an OSPF area, which include:

- o An OSPF TTZ represents a set of TTZ edge routers, connected by a full mesh of virtual connections between them.

- o Non-TTZ link-state information is handled as normal. TTZ routers receive the link-state information about the topology outside of the TTZ, store the information, and flood the information through the TTZ to the routers outside of the TTZ.

5.2. TTZ Example

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 to 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 T71, T75, and T81 in TTZ 600.

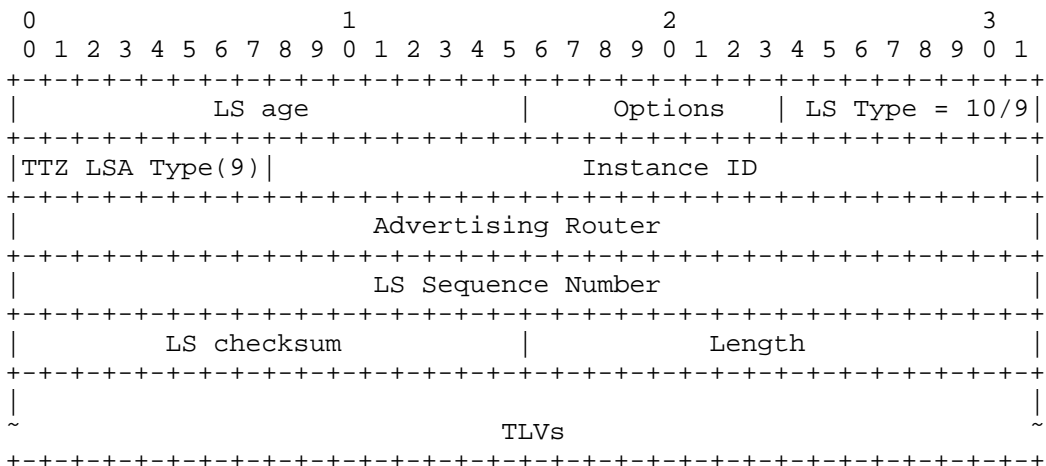
In addition, TTZ 600 comprises six TTZ-internal routers: T71, T73, T75, T77, T79, and T81. Each of them has all its adjacent routers in TTZ 600. For instance, router T71 is a TTZ-internal router since its

6. Extensions to OSPF Protocols

The link-state information about a TTZ includes router Link-State Advertisements (LSAs), which can be contained and advertised in opaque LSAs [RFC5250] within the 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 a Link-State (LS) Type = 10/9 and TTZ LSA Type, and it contains a number of TLVs.



There are three TTZ LSAs of LS Type 10 defined:

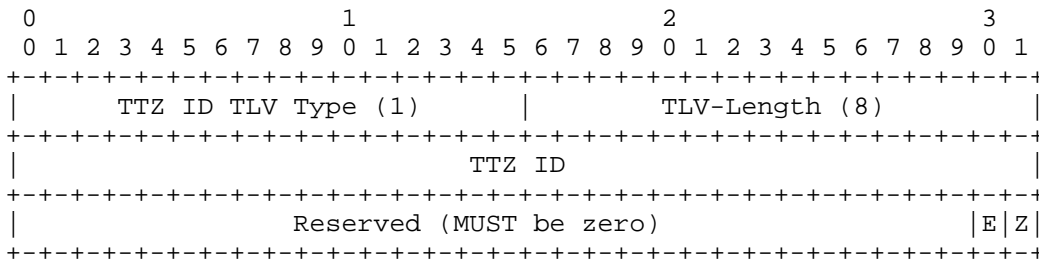
- o TTZ router LSA: a TTZ LSA containing a TTZ ID TLV and a TTZ Router TLV.
- o TTZ control LSA: a TTZ LSA containing a TTZ ID TLV and a TTZ Options TLV.
- o TTZ indication LSA: a TTZ LSA containing a TTZ ID TLV with E = 0, which indicates that the router originating this LSA is a TTZ-internal router.

There is one TTZ LSA of LS Type 9:

- o TTZ discovery LSA: a TTZ LSA containing a TTZ ID TLV and an optional TTZ Options TLV.

6.2. TTZ ID TLV

A TTZ ID TLV has the following format. It contains a TTZ ID (refer to Section 5.1) and some flags. It has the TLV-Length of 8 octets.

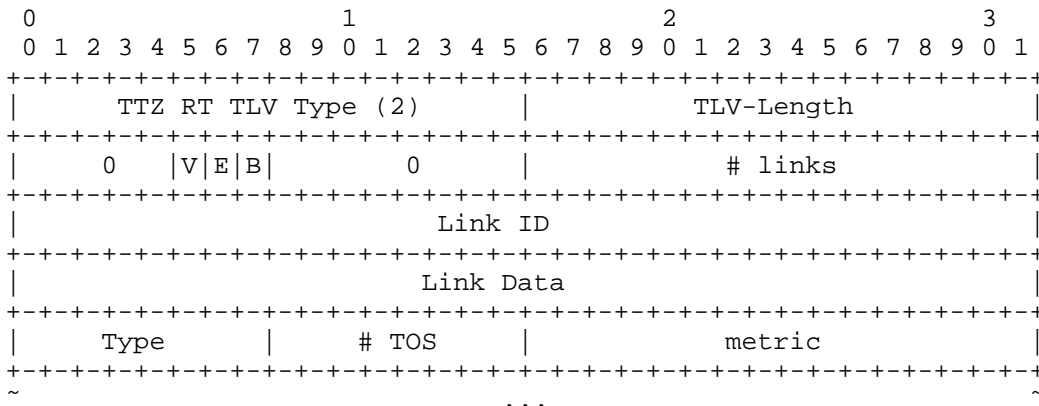


E = 1: Indicating a router is a TTZ edge router
Z = 1: Indicating a router has migrated to TTZ

When a TTZ router originates a TTZ LSA containing a TTZ ID TLV, it MUST set 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 MUST set flag Z to 1 after it has migrated to TTZ and to 0 before it migrates to TTZ or after it rolls back from TTZ (refer to Section 6.4).

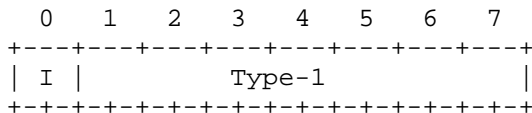
6.3. TTZ Router TLV

The format of a TTZ Router TLV is as follows. It has the same content as a standard OSPF router LSA (RFC 2328) with the following modifications.



Note: TOS = Type of Service

For a router link, the existing 8-bit Link Type field for a router link is split into two fields as follows:



I bit flag:

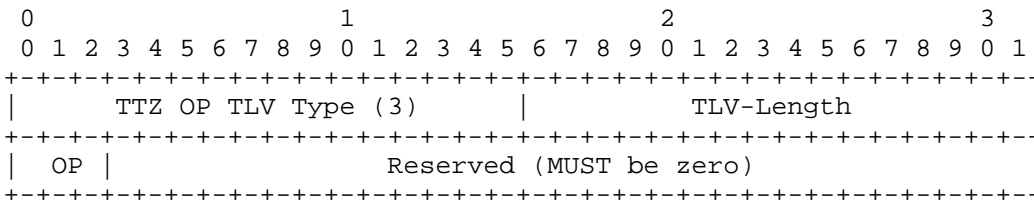
- 1: Router link is a TTZ-internal link.
- 0: Router link is a TTZ-external link.

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.

The Link Type field is 8 bits and the values 128-255 of the field are reserved (refer to [RFC4940]), which allows the reuse of the bottom 7 bits to indicate the type of TTZ-internal or external link.

6.4. TTZ Options TLV

The format of a TTZ Options TLV is as follows.



OP Value	Meaning (Operation)
0x001 (T):	Advertising TTZ Topology Information for Migration
0x010 (M):	Migrating to TTZ
0x011 (N):	Advertising Normal Topology Information for Rollback
0x100 (R):	Rolling Back from TTZ

An OP field of 3 bits is defined. It may have a value of 0x001 for T, 0x010 for M, 0x011 for N, or 0x100 for R, which indicates one of the four operations above. When any of the other values is received, it is ignored.

Advertising TTZ Topology Information for Migration (T):
After a user configures a TTZ router to advertise TTZ topology information, advertising TTZ topology information for migration is triggered. The TTZ router originates a TTZ control LSA having a TTZ Options TLV with OP for T. It also originates its other

TTZ LSA such as a TTZ router LSA or TTZ indication LSA. When another TTZ router receives the LSA with OP for T, it originates its TTZ LSA as described in Section 7.

Migrating to TTZ (M):

After a user configures a TTZ router to migrate to TTZ, migrating to TTZ is triggered. The TTZ router originates a TTZ control LSA having a TTZ Options TLV with OP for M and migrates to TTZ. When another TTZ router receives the LSA with OP for M, it also migrates to TTZ. When a router migrates to TTZ, it computes routes using the TTZ topology and the topology outside of the TTZ. For a TTZ-internal router, it also updates its TTZ indication LSA with Z = 1. For a TTZ edge router, it updates its TTZ router LSA with Z = 1 and its router LSA for virtualizing the TTZ. A TTZ router determines whether it is internal or edge based on configurations (refer to Section 11.1).

Advertising Normal Topology Information for Rollback (N):

After a user configures a TTZ router to advertise normal topology information, advertising Normal topology information for rollback is triggered. The TTZ router originates a TTZ control LSA having a TTZ Options TLV with OP for N. 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 OP for N, it forwards the LSA, advertises its normal LSAs, and stops advertising its TTZ LSAs.

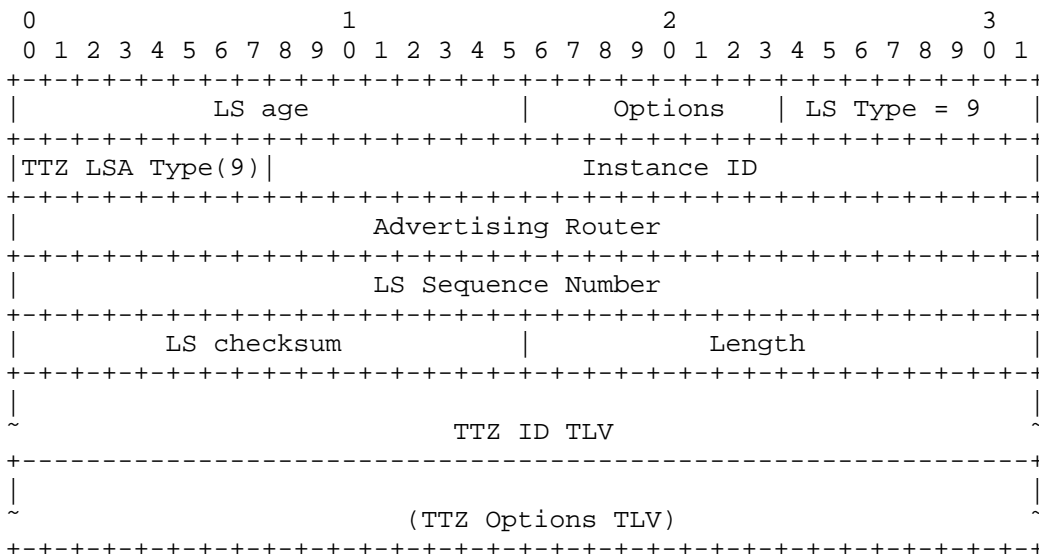
Rolling back from TTZ (R):

After a user configures a TTZ router to roll back from TTZ, rolling back from TTZ is triggered. The TTZ router originates a TTZ control LSA having a TTZ Options TLV with OP for R and rolls back from TTZ. When another TTZ router receives the LSA with OP for R, 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 flushes the TTZ control LSA if OP in the LSA is set for the configuration and the configuration is removed.

6.5. Link Scope TTZ LSA

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



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

7. Constructing LSAs for TTZ

For a TTZ, its topology is represented by the LSAs generated by its TTZ routers for the link states in the TTZ, which include TTZ router LSAs by TTZ edge routers, TTZ indication LSAs by TTZ-internal routers, normal router LSAs, and network LSAs. The TTZ router LSAs and TTZ indication LSAs MUST be generated after advertising TTZ topology information for migration is triggered.

A TTZ edge router generates a TTZ router LSA that 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 is set to 1, which indicates the originator of the LSA is a TTZ edge router. The TTZ Router TLV contains the TTZ-external links to the routers outside of the TTZ and the TTZ-internal links to the routers inside of the TTZ as described in Section 6. The TTZ router LSA containing this TLV is constructed and advertised within the TTZ.

A TTZ-internal router generates a TTZ indication LSA that has a TTZ ID TLV containing the ID of the TTZ to which the router belongs and flag E is set to 0, which indicates the originator of the LSA is a

TTZ-internal router. For a TTZ-internal router, its regular router LSA is still generated. If a TTZ router is a Designated Router (DR), it originates its regular network LSA.

After receiving a trigger to migrate to TTZ such as a TTZ control LSA with OP for M, a TTZ edge router MUST originate its normal router LSA for virtualizing a TTZ, which comprises three groups of links in general.

The first group is 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 is 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 (P2P) 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.

7.1. TTZ Migration Process

After migration to TTZ is triggered, a TTZ router computes routes using its TTZ topology (refer to Section 10), and a TTZ edge router originates its normal router LSA for virtualizing the TTZ in two steps:

Step 1: The router updates its router LSA by adding a P2P link to each of the other known edge routers in the TTZ and also by adding the stub links for the loopback addresses in the TTZ to be accessed outside of the TTZ according to configuration policies of operators.

Step 2: After `MaxLSAGenAdvTime` (0.3 s) or `sr-time + MaxLSAAadvTime` (0.1 s), it removes the TTZ links from its router LSA, where `sr-time` is the time from updating its router LSA to receiving the ack for its router LSA and receiving the updated router LSAs originated by the other TTZ edge routers. In other words, it removes the TTZ links from its router LSA after sending its updated router LSA and receiving the updated router LSAs originated by the other TTZ edge routers for `MaxLSAAadvTime` or after sending its updated router LSA for `MaxLSAGenAdvTime`. `MaxLSAAadvTime` and `MaxLSAGenAdvTime` SHOULD be set to 100 ms and 300

ms, respectively, but MAY be configurable. The former is the maximum time for an LSA to be advertised to all the routers in an area. The latter is 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.

This is to avoid a possible route down or change in a TTZ-external router while the TTZ is being virtualized. If each TTZ edge router originates its router LSA by adding its P2P links to the other TTZ edge routers and removing its TTZ links in one step, a route taking a path through the TTZ in the TTZ-external router may be down or changed before all the router LSAs generated by the TTZ edge routers reach the TTZ-external router. When the TTZ-external router computes routes with some router LSAs originated by the TTZ edge routers, bidirectional checks for some of the P2P links will fail. Thus, the route taking the path through the shortest path for the P2P link failing the bidirectional check will be down or changed.

To roll back from a TTZ smoothly after receiving a trigger to roll back from TTZ, a TTZ edge router MUST originate its normal router LSA in the above two steps in a reverse way.

Step 1: Initially, it updates its normal router LSA by adding the normal links for the links configured as TTZ links into the LSA.

Step 2: It then removes the P2P 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 MaxLSAAadvTime or after sending its updated router LSA for MaxLSAGenAdvTime.

8. Establishing Adjacencies

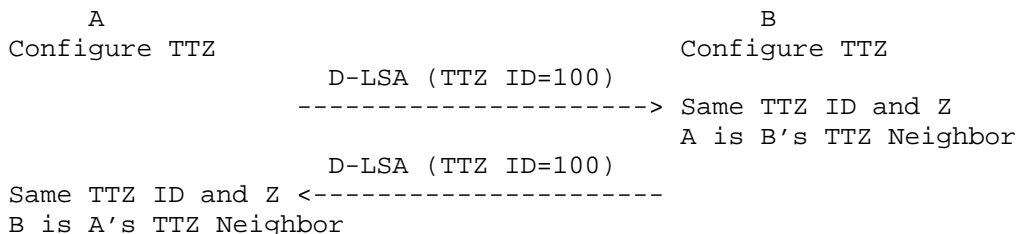
This section describes the TTZ adjacencies.

8.1. Discovery of TTZ Neighbors

When two routers A and B are connected by a P2P link and have a normal adjacency, they TTZ discover 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 or only one end is configured with a TTZ ID, TTZ adjacency over the link MUST NOT 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.



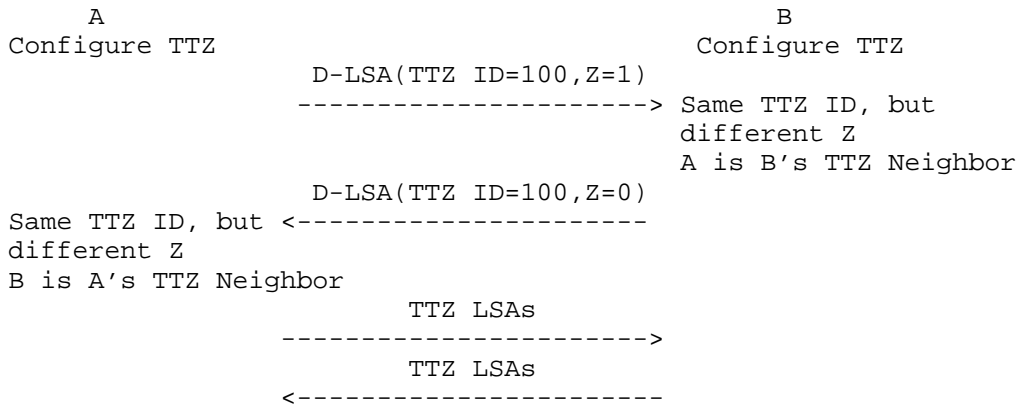
A sends B a D-LSA with TTZ ID after the TTZ is configured on it. B sends A a D-LSA with TTZ ID after the TTZ is configured on it.

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 LSA when one of the following conditions is met.

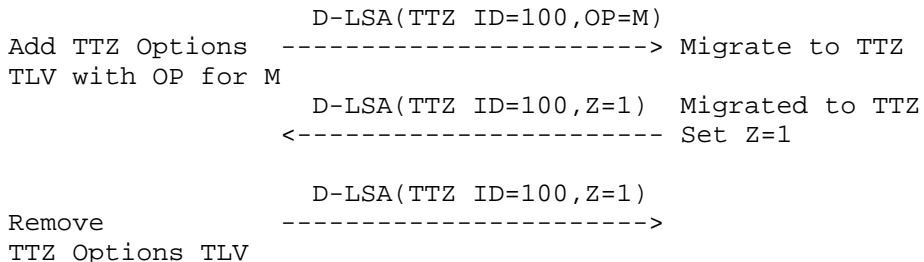
- o Z = 0 and there is a TTZ LSA with OP for T.
- o Z = 1.

B is symmetric to A and acts similarly to A.

If two ends of the link have the same TTZ ID but the Z flags are different, a TTZ adjacency over the link MUST be "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 a TTZ Options TLV with OP for M 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.

When B receives the D-LSA from A with OP for M, it starts to migrate 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 TTZ Options TLV. It also updates and advertises its LSAs as needed.

When a number of routers connected through a broadcast link have normal adjacencies among them, they also TTZ discover each other through D-LSAs. The Designated Router (DR) for the link MUST "form" 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 MUST NOT "form" any TTZ adjacency with any router attached to the link.

When a number of routers connected through a broadcast link have TTZ adjacencies among them, if a misconfigured router is introduced on the broadcast link, the DR for the link MUST NOT "form" any TTZ adjacency with this misconfigured router.

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

A TTZ adjacency over a link MUST be removed when one of the following events happens.

- o TTZ ID on one end of the link is changed to a different one.

- o TTZ ID on one end of the link is removed.
- o The D-LSA is not received after the D-LSA-MAX-RETRANSMIT-TIME or is explicitly flushed. The D-LSA-MAX-RETRANSMIT-TIME SHOULD be set to 60 minutes, but MAY be configurable.
- o Normal adjacency over the link is down.

When the TTZ ID on one end of the link is removed, the corresponding D-LSA is flushed.

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 only within the TTZ. It is not advertised to any router outside of the TTZ. For example, a router LSA generated for a TTZ-internal router is advertised only within the TTZ.

Any network LSA generated for a broadcast or Non-Broadcast Multi-Access (NBMA) network in a TTZ is advertised only within the TTZ. It is not advertised outside of the TTZ.

Any opaque LSA generated for a TTZ-internal TE link is advertised only 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. The TTZ edge router determines whether an LSA is about a TTZ-internal link state by checking if the advertising router of the LSA is a TTZ-internal router (i.e., there is a TTZ indication LSA generated by the TTZ-internal router that has the same advertising router).

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 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 inside and outside of 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. It computes routes using the TTZ router LSAs and the regular LSAs, excluding the router LSAs for virtualizing the TTZ. That is, it computes routes using the TTZ topology and the topology outside of the TTZ, excluding the links for virtualizing 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 the same TTZ ID as a TTZ link, the TTZ is determined. A router with some links in a TTZ and some links not in this TTZ is a TTZ edge router. A router with all its links in a TTZ is a TTZ-internal router.

2. Configuring TTZ on Routers in TTZ

A same TTZ ID is configured on every TTZ-internal 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 option 1 above.

For a TTZ edge router X with different TTZ IDs on its different links, router X connects two or more different TTZs. In this case, router X originates its router LSA for virtualizing the TTZs. This LSA includes the normal links connecting to routers outside of these TTZs and the virtual links to the other edge routers of each of these TTZs. Router X also originates its TTZ router LSA for each of the TTZs. The TTZ router LSA for TTZ N includes the links to the routers outside of these TTZs, the virtual links to the other edge routers of the other TTZs, and the TTZ links to the routers in TTZ N.

11.2. 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 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 or advertise its TTZ topology information. It will discover its TTZ neighbors.

Second, after configuring the TTZ, a user issues a configuration 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 configuration, it originates a TTZ control LSA with OP for T (indicating TTZ information generation and advertisement for migration). It also originates its TTZ LSA, such as TTZ router LSA or TTZ indication LSA, and advertises the LSA to its TTZ neighbors. When another router in the TTZ receives the LSA with OP for T, it originates its TTZ 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.

Third, 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 TTZ LSAs, including TTZ router LSAs from TTZ edge routers and TTZ indication LSAs from TTZ-internal routers. This information may be displayed on a router through a configuration.

Then, a user activates the TTZ through using a configuration such as migrate to TTZ on one router in the TTZ. The router migrates to TTZ, generates and advertises a TTZ control LSA with OP for M (indicating Migrating to TTZ) after it receives the configuration. After another router in the TTZ receives the TTZ control LSA with OP for M, 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 configurations according to certain procedures. In an abnormal case, for example, two conflicting configurations are issued on two TTZ routers in a TTZ at the same time, and a TTZ router issues an error and logs the error when it detects a conflict.

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

- 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.

After receiving a configuration on a router to migrate to TTZ, which is for 3), the router determines whether 2) is performed by 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 configuration at this time.

After a router receives a TTZ LSA with OP for M for 3) from another router, it determines whether 2) is performed by checking if it has received/originated TTZ LSAs. If it has not, it issues an error and logs the error, and it does not migrate to TTZ. In this case, it

does not originate its router LSA for virtualizing the TTZ if it is a TTZ edge router.

After receiving a configuration on a router to generate and advertise TTZ information, which is for 2), the router determines whether 1) is performed by 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 configuration at this time.

For rolling back from TTZ, the correct sequence of configurations is below.

- 1) configure on one router in the TTZ to trigger every router in the TTZ to advertise normal LSAs and stop advertising TTZ LSAs; and
- 2) configure on one router in the TTZ to trigger every router in the TTZ to roll back from TTZ.

After receiving a configuration on a router to roll back from TTZ, which is for 2), the router determines whether 1) is performed by checking if it has received TTZ LSA with OP for N. If it has not, it issues an error to an operator (advertise normal LSAs and stop advertising TTZ LSAs as rolling back from TTZ is not done yet) and rejects the configuration at this time.

After a router receives a TTZ LSA with OP for R for 2) from another router, it determines whether 1) is performed by checking if it has received TTZ LSA with OP for N. If it has not, it issues an error and logs the error, and it does not roll back from TTZ.

After receiving a configuration on a router to advertise normal LSAs and stop advertising TTZ LSAs for rolling back from TTZ, which is for 1), the router checks whether it has any TTZ LSAs. If it does not, it issues an error to an operator (no TTZ to be rolled back) and rejects the configuration 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 migrated TTZ router (say T1), and there is a normal adjacency between them over the link, a user can configure TTZ on both ends of the link to add R1 into the TTZ to which T1 belongs. They TTZ discover each other as described in Section 8.

When a number of non-TTZ routers are connected via a broadcast or NBMA link to a migrated TTZ router (say T1), 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 TTZ discovery process described in Section 8.

12. Manageability Considerations

Section 11 ("Operations") outlines the configuration process and deployment scenarios for a TTZ. The configurable item is enabling a TTZ on a router and/or an interface on a router. The TTZ function may be controlled by a policy module and assigned a suitable user privilege level to enable. A suitable model may be required to verify the TTZ status on routers participating in the TTZ, including their role as an internal or edge TTZ router. The mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those indicated in [RFC2328].

13. Security Considerations

A notable beneficial security aspect of TTZ is that the TTZ is enclosed in a single area, and TTZ could be used to mask the internal topology. External routers that are not participating in the TTZ will not be aware of the internal TTZ topology. It should be noted that a malicious node could inject TTZ LSAs with the OP field set to M or R, which could trigger the migration into/from a TTZ and may result in the isolation of some routers in the network. Good security practice might reuse the OSPF authentication and other security mechanisms described in [RFC2328] and [RFC7474] to mitigate this type of risk.

14. IANA Considerations

Under the registry name "Opaque Link-State Advertisements (LSA) Option Types" [RFC5250], IANA has assigned a new Opaque Type registry value for TTZ LSA as follows:

Registry Value	Opaque Type	reference
9	TTZ LSA	This document

IANA has created and will maintain a new registry:

- o OSPFv2 TTZ LSA TLVs

Initial values for the registry are given below. The future assignments are to be made through IETF Review [RFC5226].

Value	OSPFv2 TTZ LSA TLV Name	Definition
-----	-----	-----
0	Reserved	
1	TTZ ID TLV	see Section 6.2
2	TTZ Router TLV	see Section 6.3
3	TTZ Options TLV	see Section 6.4
4-32767	Unassigned	
32768-65535	Reserved	

15. References

15.1. 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>>.
- [RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <<http://www.rfc-editor.org/info/rfc7474>>.
- [RFC4940] Kompella, K. and B. Fenner, "IANA Considerations for OSPF", BCP 130, RFC 4940, DOI 10.17487/RFC4940, July 2007, <<http://www.rfc-editor.org/info/rfc4940>>.

15.2. Informative References

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

Appendix A. Prototype Implementation

A.1. What Is Implemented and Tested

1. Command-Line Interface (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 "Extensions to OSPF Protocols" (Section 6) are implemented and tested except for rolling back from TTZ. The testing results illustrate:

- o As seen from the outside, a TTZ is virtualized as its edge routers connected to each other. Any router outside of the TTZ sees the edge routers (as normal routers) connecting to 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 to become 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.

A.2. Implementation Experience

The implementation of TTZ reuses the existing OSPF code along with additional simple logic. A couple of engineers started to work on implementing the TTZ from the middle of June 2014 and finished coding it just before the end of July 2014. 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.

Acknowledgements

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

Contributors

The following people contributed significantly to the content of this document and should be considered co-authors:

Mehmet Toy
United States of America
Email: mehmet.toy@verizon.com

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

Anil Kumar SN
India
Email: anil.sn@huawei.com

Ning So
United States of America
Email: ningso01@gmail.com

Lei Liu
United States of America
Email: lliu@us.fujitsu.com

We also acknowledge the contribution of the following individuals:

Veerendranatha Reddy Vallem
India
Email: veerendranatharv@huawei.com

William McCall
United States of America
will.mccall@rightside.co

Authors' Addresses

Huaimo Chen
Huawei Technologies
Boston, MA
United States of America

Email: huaimo.chen@huawei.com

Renwei Li
Huawei Technologies
2330 Central expressway
Santa Clara, CA
United States of America

Email: renwei.li@huawei.com

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Rd.
Raleigh, NC 27709
United States of America

Email: aretana@cisco.com

Yi Yang
Sockrate
Cary, NC
United States of America

Email: yyang1998@gmail.com

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

Email: liu.cmri@gmail.com