Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 29, 2013

H. Chen
R. Li
Huawei Technologies
G. Cauchie
France Telecom
N. So
Tata Communications
L. Liu
UC Davis
February 25, 2013

Applicability of OSPF Topology-Transparent Zone
draft-chen-ospf-ttz-app-02.txt

Abstract

This document discusses the applicability of "OSPF Topology-Transparent Zone". It briefs the protocol and its operations first, and then illustrates the application scenarios of OSPF Topology-Transparent Zone. In addition, guidelines for deployment are presented and limitations of the protocol are indicated. This document is intended for accompanying "OSPF Topology-Transparent Zone" to the Internet standards track.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

The number of routers in an Autonomous System (AS) becomes larger and larger as the Internet traffic keeps growing. Thus the Open Shortest Path First (OSPF) Link State Database (LSDB) and OSPF routing table are bigger and bigger. Any link state change in an AS leads to a number of link state distributions to every router in the AS. This triggers every router in the AS to re-calculate its OSPF routes, update its Routing Information Base (RIB) and Forwarding Information Base (FIB). All these will consume network resource including network bandwidth and Central Process Unit (CPU) time. This blocks the further expansion of a network.

The current solution for this is to divide a big AS into a number of OSPF areas. However, there are a number of issues when an AS is split further into multiple areas.

At first, dividing an AS into more areas is a very challenging task since it is involved in network architecture changes.

Secondly, it is complex for a Multi-Protocol Label Switching (MPLS) Traffic Engineering (TE) Label Switching Path (LSP) crossing multiple areas to be established. In general, a TE path crossing multiple areas is computed by using collaborating Path Computation Elements (PCEs) through the PCE Communication Protocol (PCEP), which is not easy to configure by operators since the manual configuration of the sequence for a PCE chain 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.

Thirdly, some policies need to be configured on area border routers for reducing the number of link states such as summary Link-State Advertisements (LSAs) to be distributed to other routers in other areas.

Furthermore, route convergence may be slower. A router in an OSPF area can see all other routers in the same area. A link-state change anywhere in an OSPF area will be populated everywhere in the same area, and may even be distributed to other areas in the same AS indirectly.

A link state change in an area will lead to every router in the same area to re-calculate its OSPF routes, update its RIB and FIB. It may also lead to a number of link state distributions to other areas. This will trigger routers in other areas to re-calculate their OSPF routes, update their RIBs and FIBs.

This document introduces a technology called Topology-Transparent Zone (TTZ), presents a number of application scenarios of TTZ and illustrates that TTZ can resolve the issues above. In addition, guidelines for deployment are presented and limitations of the protocol are indicated.

2. Terminology

This document uses terminologies defined in RFC 2328, and RFC 2740.

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. Overview of Topology-Transparent Zone

This section briefs the concept of Topology-Transparent Zone (TTZ) and explains the TTZ in some details through an example.

4.1. Definitions of Topology-Transparent Zone

A Topology-Transparent Zone (TTZ) comprises an Identifier (ID), a group of routers and a number of links connecting the routers. A Topology-Transparent Zone is in an OSPF domain.

The ID of a Topology-Transparent Zone (TTZ) or TTZ ID is a number that is unique for identifying a node in an OSPF domain. It is not zero in general.

A TTZ may be virtualized as a different object in a different way. Two typical ways are given below.

In a first way, a TTZ may be virtualized as a group of TTZ edge routers fully connected. That is that from a router outside of the TTZ, a TTZ is seen as a group of TTZ edge routers, which are fully connected.

In a second way, a TTZ may be seen as a single router. That is that from a router outside of the TTZ, a TTZ is seen as a special single router. This router has a router ID, which is the TTZ ID or the maximum ID among all the router IDs of the routers in the TTZ. For every connection between a TTZ edge router and a router outside of TTZ, there is a connection between the special router and the router

outside of TTZ.

The virtualization of TTZ in the first way is described and used below.

4.2. An Example of TTZ

The figure below illustrates an example of a routing domain containing a topology-transparent zone: TTZ 600.

```
                                        TTZ 600
                                        /
                                       /

                      ^~^~^~^~^~^~^~^~^~^~^~^~^~^~^
                     (                             )
    ===[R15]========(==[R61]-----------------[R63]==)========[R29]===
        ||         (   |   \                /   |    )        ||
        ||         (   |    \              /    |    )        ||
        ||         (   |     \            /     |    )        ||
        ||         (   |      \          /      |    )        ||
        ||         (   |       \        /       |    )        ||
```

```
        ||         (   |        \      /        |    )        ||
        ||         (   |         \    /         |    )        ||
        ||         (   |    _____[R71]          |    )        ||
        ||         (   |   /     /    \         |    )        ||
        ||         (   | [R73]  /      \        |    )        ||
        ||         (   |       /        \       |    )        ||
        ||         (   |      /          \      |    )        ||
        ||         (   |     /            \     |    )        ||
        ||         (   |    /              \    |    )        ||
        ||         (   |   /                \   |    )        ||
    ===[R17]========(==[R65]-----------------[R67]==)========[R31]===
         \\          (//                          \\ )       //
          ||         //v~v~v~v~v~v~v~v~v~v~v~v~v~\\         ||
          ||        //                            \\        ||
          ||       //                              \\       ||
          ||      //                                \\      ||
          ||     //                                  \\     ||
          ||    //                                    \\    ||
           \\  //                                      \\  //
       =====[R23]======================================[R25]=====
             //                                          \\
            //                                            \\
           //                                              \\
```

Figure 1: An Example of TTZ

The routing domain comprises routers R15, R17, R23, R25, R29 and R31. It also contains a topology-transparent zone TTZ 600. The TTZ 600 comprises routers R61, R63, R65, R67, R71 and R73, and the links connecting them.

There are two types of routers in a Topology-Transparent Zone (TTZ): TTZ internal routers and TTZ edge routers. A TTZ internal router is a router inside the TTZ and every adjacent router of the TTZ internal router is a router inside the TTZ. A TTZ edge router is a router inside the TTZ and has at least one adjacent router that is outside of the TTZ and at least one adjacent router that is inside the TTZ.

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

router since it is connected to router R15, which is outside of the TTZ.

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

A TTZ may hide the information inside the TTZ from the outside. It may not distribute any internal information about the TTZ to a router outside of the TTZ.

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

In order to create a Topology-Transparent Zone (TTZ), we must configure the same TTZ ID on every link that connects routers inside the TTZ and every router in the TTZ must support TTZ feature.

For example, the same TTZ ID is configured on the nine links below:

- the link between router R61 and R65,
- the link between router R65 and R67,
- the link between router R67 and R63,
- the link between router R63 and R61,

- the link between router R71 and R61,
- the link between router R71 and R63,
- the link between router R71 and R65,
- the link between router R71 and R67 and

o the link between router R71 and R73.

Thus six routers R61, R63, R65, R67, R71 and R73, and nine links among these six routers form a topology-transparent zone TTZ 600 in the figure above.

From a router outside of the TTZ, a TTZ is seen as a group of TTZ edge routers, which are 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. These four TTZ edge routers are fully connected.

5. Applicability of Topology-Transparent Zone

Topology-Transparent Zone (TTZ) may be used in different cases. This section presents a number of application scenarios of TTZ and illustrates the benefits that TTZ brings in each scenario.

5.1. TTZ for Network Scaling

In this section, we present and analyze the application of TTZ to a network, which has 1,600 routers in one OSPF area.

Before TTZ is applied to the network, every router has the information about the network topology of 1,600 routers in its link state database. Each router in the network originates a router LSA, which is distributed to all the other routers in the network and stored in their link state databases. Thus, there are 1,600 router LSAs in every router's link state database.

Assume that TTZ is deployed in the network and divides the network into 10 Topology-Transparent Zones (TTZs). For the simplicity of analysis, we assume that every TTZ has 160 routers, six of which are edge routers of the TTZ. In this case, every router belongs to a TTZ and has two sets of link states.

One set of link states describes the topology of the TTZ to which the router belongs. There are 160 router LSAs in this set, which may be stored in a link state database in a router.

The other set of link states describes the topology outside of the

TTZ. This topology contains the information about the link states for all the edge routers of all the TTZs in the network. There are ten TTZs in the network, each of which has six edge routers. Thus there are 60 router LSAs in this set.

After TTZ is applied to the network, every router in a TTZ has only 220 router LSAs in its link state databases. 160 of them describe the topology of the TTZ and the other 60 the topology outside of the TTZ.

The size of the link state database in every router is reduced significantly after TTZ is applied to the network. In the sense of the number of router LSAs, the size is reduced to 220 from 1,600. The saving is (1,600 - 220)/1,600 = 86.25 %.

5.2. One Area Network

Many networks start with one area. A network with only one area is easy to operate and maintain. As a network with one area becomes bigger and bigger because the increasing traffic in the network drives the expansion of the network, it needs to be split into multiple areas in general.

5.2.1. Issues on Splitting Network into Areas

Splitting a network with only one area into multiple areas is a very challenging task and may raise a number of issues.

1. Significant Changes on Network Architecture

There are significant changes on network architecture when splitting a network with one area into multiple areas. Originally the network has only one area, which is backbone area. This original backbone area will be split into a new backbone area 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 area 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.

Before splitting the network into areas, every router in the network has the information about the network topology. However, after splitting the network into areas, each router in an area has the information of the topology of the area, and it does not have the

information of the topology of any other area. It has only the summary information about the other areas.

2. Complex for MPLS TE Tunnel Setup

Each of the MPLS TE LSP tunnels originally in one area, which has its ingress and egress in different areas after the network splitting, needs to be re-configured and re-established. It is very complex for a MPLS TE LSP tunnel crossing areas to be set up.

In order to reduce the manual configurations for a MPLS TE LSP tunnel crossing multiple areas, we use PCEs to compute the path for the tunnel. Thus we must configure PCEs for the network split into multiple areas.

In addition, we need to provide a sequence of areas for the tunnel through manual configurations. The tunnel will go through the sequence of areas provided.

More critically, there are some issues on using PCEs. One of them is that the path computed by PCEs for the tunnel may not be optimal. If the optimal path for the tunnel is not in the sequence of areas configured by uesrs, the path found by PCEs for the tunnel will not be optimal.

3. Policy Configurations

Some policies need to be applied on area border routers for reducing the number of link states such as summary LSAs to be distributed to other routers in other areas.

On an ABR between a non backbone area (say area A) and the backbone area (say area B), there are four sets of policies, which may be configured. One set of them is the policies that block some summary LSA distributions from area B to area A. The second set of them is the policies that aggregate some routes for reducing the number of summary LSAs to be distributed from area B to area A. The third set is the policies that block some summary LSA distributions from area A to area B. The forth set of them is the policies that aggregate some routes for reducing the number of summary LSAs to be distributed to area A from area B.

In addition, OSPF need to be disabled before any of the policy commands is issued. And then OSPF must be enabled after the policies are configured.

4. Slow Convergence

A router in an OSPF area can see all other routers in the same area. A link-state change anywhere in an OSPF area will be populated everywhere in the same area, and may even be distributed to other areas in the same AS indirectly. For example, all the routers and links in a Point-Of-Presence (POP) in an OSPF area will be seen by all the other routers in the same area. Any link state change in the POP will be distributed to all the other routers in the same area and may be distributed to routers in other areas indirectly.

A link state change in an area will lead to every router in the same area to re-calculate its OSPF routes, update its RIB and FIB. It may also lead to a number of link state distributions to other areas. This will trigger routers in other areas to re-calculate their OSPF routes, update their RIBs and FIBs. All these will consume network resource including network bandwidth and Central Process Unit (CPU) time. Thus route convergence is slow.

5.2.2. Use of TTZ in One Area Network

The issues mentioned above on splitting network into areas disappear if we do not split network into areas and use OSPF Topology Transparent Zone (TTZ) instead.

TTZ may be applied to a group of routers and links in the network directly. For a group of routers and links connecting the routers in the group in the network, no matter where it localtes in the network, we may configure it as an OSPF TTZ as long as each router in the group can reach the other routers in the group through those links.

1. No Significant Changes on Network Architecture

There is not any significant changes on network architecture when an OSPF TTZ is applied to a group of routers and links in the network directly.

At first, we do not add any new connection to the network, or remove any existing connection from the network.

Secondly, every router outside of the TTZ is not aware of the TZZ.

Even the router directly connecting to the TTZ is not aware of the TTZ.

Furthermore, every router in the network still has a topology view of the network. Except for those internal TTZ routers and links, which are hidden, every router outside of the TTZ has the link state information about all the routers and links in the network.

2. Easy for MPLS TE Tunnel Setup

After a group of routers and links in the network is configured as an OSPF TTZ, a MPLS TE LSP tunnel with an ingress router and an egress router, which are anywhere in the network, can be configured in a way, which is the same as or similar to the way in which a MPLS TE LSP tunnel in one area network is configured.

For example, in the network in Figure 1 above, a MPLS TE LSP tunnel from ingress router R15 to egress router R29 can be configured in the same way as a MPLS TE LSP tunnel in one area network through provisioning the ingress router R15's IP address, the egress router R29's IP address and some constraints for the tunnel on the ingress router R15.

We do not need any PCEs for computing the constrained path for a MPLS TE LSP tunnel in the network with OSPF Topology Transparent Zones (TTZs). After a MPLS TE LSP tunnel with an ingress and egress anywhere in the network with OSPF TTZs is configured, the ingress computes the constrained path for the tunnel from the ingress to the egress in the same way as it computes the constrained path for the tunnel in one area network. The constrained path computed may go through some OSPF TTZs.

For example, in the network in Figure 1 above, the constrained path computed for the tunnel from ingress router R15 to egress router R29 may be from ingress router R15, to edge router R61 of TTZ 600, to edge router R63 of TTZ 600 and then to egress router R29.

As soon as the constrained path for a MPLS TE LSP tunnel is computed or given through configuration, the LSP can be established along the path by the signalling protocol RSVP-TE.

3. No Policy Configurations and Fast Convergence

There is not any policy configuration for the use of TTZ in one area network. When we apply a TTZ to a group of routers and links connecting the routers, we just configure a TTZ link attribute TTZ ID on each of the links.

Link state changes such as a link down or a router down inside a TTZ is only advertised internally inside the TTZ. The routers inside the TTZ will do internal routing table re-calculation. The routers outside of the TTZ will not see any topology change, and thus the routing table is not re-calculated and route convergence is fast.

5.3. Multi-Area Network

For a network with mutiple areas, it typically needs to be split into more areas when the size of the network becomes larger and larger as

the traffic in the network keeps growing.

5.3.1. Issues on Splitting Network into More Areas

What would happen when we split a network with multiple areas into even more areas?

1. Significant Changes on Network Architecture

The changes on network architecture are significant when a network with multiple areas is split into even more areas. In the network before splitting, there is one backbone area, which is surrouned by a number of non backbone areas. Each of the non backbone areas is connected to the backbone area. There is not any direct connection between any two non backbone areas in general.

Splitting the network into more areas is involved in re-arranging a number of routers and links to create new areas and make some of the existing areas to become smaller. Some of the routers and links in a new area may be from the backbone area, and the other from some of the non backbone areas.

In the network after splitting, there is still one backbone area, which must be changed and be surrounded by the new non backbone areas and the existing non backbone areas some of which have been changed.

Each of the non backbone areas is connected to the backbone area.

2. More Configurations for Tunnel Setup

For reducing the manual configurations for a MPLS TE LSP tunnel crossing multple areas, we use PCEs to compute the path for the tunnel. Thus more configurations for tunnel setup is needed. We must configure PCEs for each of the new areas and peer relations among the PCEs for the new areas and the PCEs for the existing areas.

3. More Policy Configurations and Slow Convergence

More policy configurations may be needed. For each of the new non backbone areas and each of the existing non backbone areas that have been changed, some policies may need to be configured on the area border router connecting it to the backbone area to aggregate the routes in the non backbone area for transmisstion to the backbone area.

A link state change in every area including any of the new areas and the existing areas will lead to every router in the same area to re-calculate its OSPF routes, update its RIB and FIB. It may also lead to a number of link state distributions to other areas. This will

trigger routers in other areas to re-calculate their OSPF routes, update their RIBs and FIBs. All these will consume network resource including network bandwidth and CPU time. Thus route re-convergence is slow.

5.3.2. Use of TTZ in Multi-Area Network

The issues described above on splitting network into even more areas disappear if we do not split network into more areas and use OSPF TTZ instead.

A TTZ may be applied to a group of routers and links in any area in the network directly. For a group of routers and links connecting the routers in the group in an area, no matter where it localtes in the area, we may configure it as an OSPF TTZ as long as each router in the group can reach the other routers in the group through those links.

1. No Significant Changes on Network Architecture

We can see that there is not any significant changes on network architecture when an OSPF TTZ is applied to a group of routers and links in an area in the network directly.

At first, we do not add any new connection to the network, or remove any existing connection from the network.

Secondly, every router outside of the TTZ is not aware of the TZZ. Even the router directly connecting to the TTZ is not aware of the TTZ.

Furthermore, every router in the area still has a topology view of the area. Except for those internal TTZ routers and links, which are hidden, every router outside of the TTZ has the link state information about all the routers and links in the area.

2. No Extra Configurations for Tunnel Setup

After a group of routers and links in an area in the network is configured as an OSPF TTZ, there is not any extra configuration for supporting setup of a MPLS TE tunnel. We do not need to configure any new PCE since there is not any new area generated after applying a TTZ to a group of routers and links in an area.

3. No More Policy Configurations and Fast Convergence

The architecture of areas is not changed after applying a TTZ to a group of routers and links in an area. Thus there is no more policy

configuration on the existing area border routers connecting existing non backbone areas to the existing backbone area to aggregate the routes in the non backbone area for transmission to the backbone area in general.

Link state changes such as a link down or a router down inside a TTZ is only advertised internally inside the TTZ. And thus, routers inside the TTZ will do internal routing table re-calculation. The routers outside of the TTZ will not see any topology change, and thus the routing table is not re-calculated.

5.4. Use of TTZ on Routers in IPRAN

In an IP RAN network, a connection between a cell site Base Station (BS) and a Radio Network Controller (RNC) may cross over a number of OSPF areas or ASes. The upper part of the figure below illustrates that a connection between a BS and a RNC crosses over two OSPF areas: area 1 and area 0.

```
                                          RNC
                                         //
                               ^~^~^~^~^~^~^
                ^~^~^~^~^     (             )
               (         )   (               )
              (           [R2]               )
              (  Area 1   )  (     Area 0    )
              (           )  (               )
               (         )   (               )
                v~v~v~v~v     (             )
              //               v~v~v~v~v~v~v
             BS
                            ||
                            ||
                            \/             RNC
                                         //
                 ^~^~^~^~^~^~^~^~^~^~^~^~^~^
                (  ^~^~^~^~^                )
               (  (         )                )
               ( (          [R2]             )
               ( (    TTZ    )      Area 0   )
               ( (           )               )
               (  (         )                )
                (  v~v~v~v~v                )
              // v~v~v~v~v~v~v~v~v~v~v~v~v~v
             BS
```

Figure 2: TTZ for IP RAN Case

The lower part of the figure shows that applying an OSPF TTZ for area 1 converts the network into only one backbone area, which is area 0. Operating the network with one area is easier than operating one with multiple areas. Moreover, it is easy to provide a MPLS LSP

connection within one area between a BS and a RNC. In addition, a MPLS TE LSP within one area can be protected using FRR. Thus the service over the connection is protected.

5.5. Use of TTZ on Routers in POP

A Point of Presence (POP) comprises the routers in a room, a floor, a building or a group of buildings. These routers are normally in an AS or OSPF area.

We may increase the network scalability significantly through configuring a POP as an OSPF TTZ. When a POP becomes a TTZ, the link state information about every link and every router inside the POP is hidden from a router outside of the POP. Only very small amount of link state information virtualizing the TTZ for the POP is distributed to the router outside of the POP. Thus, the size of the LSDB on every router in the network is reduced significantly, and the speed for the router to compute the shortest path to every destination is increased dramatically.

We may also improve the network availability when we use a TTZ for a POP. In the case that a link or a router in the POP is down, the traffic may be interrupted without using any TTZ for the POP. The link state information about the link or router down needs to be distributed to every router in the network, and every router needs to compute the shortest path to every destination and download the path to its FIB. The traffic is forwarded according to the latest FIB on every router. During this process, there may be inconsist views on the network topology on different routers.

The traffic interruption is minimized if we use a TTZ for the POP. The link state information about the link or router down is hidden from every router outside of the POP, which is not aware of the link or router down in the POP. Thus every router outside of the POP has the same topology view when the link or router is down. It does not compute the shortest path or download the path to its FIB.

5.6. Use of TTZ on Routers from Same Vendor

In a network, we may separate the routers from different vendors through using TTZ in order to alleviate the possible multi-vendor inter-operability issue. For example, the routers from a same vendor can be configured as a TTZ, and the routers outside of this TTZ are developed by different vendors.

5.7. Use of TTZ on Routers in a Power Saving Group

A power saving group is a set of routers and links, wherein the routers are connected through the links and there is a redundant route or path from a router in the group to another router in the group. The redundant path is within the group. That is that every hop in the redundant path is within the group.

In a power saving group, when the usage of a link within the group between two routers crosses a given threshold value for shutting down the link to save energy, the link will be shut down. This link down in the power saving group will not be distributed to any router outside of the group. The traffic outside of the group will not be affected by the link down inside the group.

From the characteristics of a power saving group, we can see that a power saving group is very suitable to be configured as a TTZ.

6. Deployment Considerations

When deploying a topology transparent zone, there should be well defined administrative policies governing the selection of routers belonging to the zone. Furthermore, special considerations should be given to the number of edge routers of the zone. In general, the number of edge routers of a topology transparent zone should be small.

7. Limitations

A topology transparent zone is within an OSPF area. It can not be in more than one OSPF areas.

A router may belong to a topology transparent zone. It can not belong to more than one topology transparent zones.

8. Security Considerations

This document does not introduce any new security issues.

9. Contributors

Yuanbin Yin
Huawei Technologies
Beijing,
China

Email: yinyuanbin@huawei.com

Figure 3:

10. Acknowledgement

The authors would like to thank Alvaro Retana, Acee Lindem, and Dean Cheng for their valuable comments on this draft.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.

11.2. Informative References

[RFC2370] Coltun, R., "The OSPF Opaque LSA Option", RFC 2370, July 1998.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC5786] Aggarwal, R. and K. Kompella, "Advertising a Router's

Local Addresses in OSPF Traffic Engineering (TE) Extensions", RFC 5786, March 2010.

[OSPF-TTZ]
Chen, H., Li, R., Cauchie, G., So, N., and L. Liu, "OSPF Topology-Transparent Zone", draft-chen-ospf-ttz, Work in Progress, July 2012.

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

Gregory Cauchie
France Telecom
38-40 avenue du General LECLERC
Issy-les-Moulineaux 92130,
FRANCE

Email: greg.cauchie@gmail.com

Ning So
Tata Communications
2613 Fairbourne Cir.
Plano, TX 75082
USA

Email: ning.so@tatacommunications.com

Lei Liu
UC Davis
USA

Email: liulei.kddi@gmail.com