OSPF Topology-Transparent Zone

draft-chen-ospf-ttz-09

Huaimo Chen (huaimochen@huawei.com)
Renwei Li (renweili@huawei.com)
Anil Kumar S N (anil.sn@huawei.com)
Gregory Cauchie (greg.cauchie@gmail.com)
Alvaro Retana (aretana@cisco.com)
Ning So (ningso01@gmail.com)
Vic Liu (liuzhiheng@chinamobile.com)
Mehmet Toy (mehmet_toy@cable.comcast.com)
Lei Liu (liulei.kddi@gmail.com)

Prototype Implementation

- Introduction
- What are Implemented and Tested
  1. CLIs for TTZ
  2. Extensions to OSPF Protocols for TTZ
  3. Smooth Migration to TTZ
  4. Add a Router to TTZ
  5. Leak TTZ Loopbacks Outside
- Implementation Experience
- Conclusions

Introduction

What are the issues as a network grows bigger & bigger?

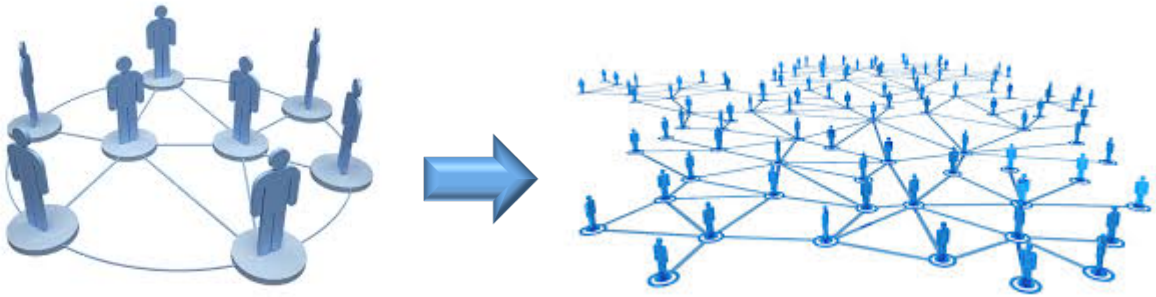

Need for Split

- Very challenging (network architecture changes)
- Time consuming (months/even year's planning and executing)
- Service interruptions (service interruptions while splitting)
- Complex for E2E services (need PCEs' help for E2E path)

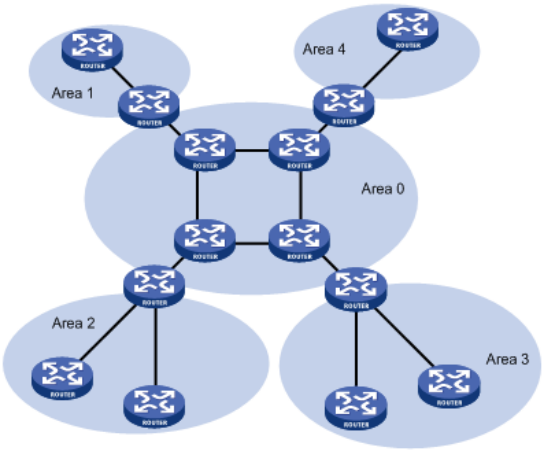

TTZ resolves these issues

CLIs for TTZ

5 New Commands

- **ttz** *ttz-id* **internal**

  at area level, configure a TTZ internal router

```
ospf 1 router-id 4.4.4.4
 opaque-capability enable
 area 0.0.0.0
  ttz 100 internal
```

- **ttz** *ttz-id* **edge**

  at area level, configure for a TTZ edge router

```
ospf 1 router-id 2.2.2.2
 opaque-capability enable
 area 0.0.0.0
  ttz 100 edge
```

- **ospf ttz-enable** *ttz-id*

  at interface level, for a TTZ link on a TTZ edge router

```
interface Ethernet3/0/1
 undo shutdown
 ip address 2.2.2.1 255.255.255.0
 ospf ttz-enable 100
 ospf network-type p2p
 ospf enable 1 area 0.0.0.0
```

- **ttz state** *ttz-id* **prepare-migrate**

  at area level, triggers every TTZ router to prepare for migration to TTZ

```
ospf 1 router-id 2.2.2.2
 area 0.0.0.0
  ttz state 100 prepare-migrate
```

- **ttz state** *ttz-id* **migrate**

  at area level, triggers every TTZ router to migrate to TTZ

```
ospf 1 router-id 2.2.2.2
 area 0.0.0.0
  ttz state 100 migrate
```

Extensions to OSPF Protocols for TTZ

All “Extensions to OSPF Protocols” except for rolling back from TTZ:

- A TTZ is virtualized to outside as its edges fully connected.
- Link state on TTZ is contained within the TTZ.
- 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.
- TTZ is backward compatible. Any router outside of a TTZ does not need to support or know TTZ.

Configure and Smoothly Migrate to TTZ

1: Configure TTZ (on all TTZ routers, TTZ links on TTZ edge routers)

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
#
interface Ethernet3/0/0
 undo shutdown
 ip address 1.1.1.2 255.255.255.0
 ospf network-type p2p
 ospf enable 1 area 0.0.0.0
#
interface Ethernet3/0/1
 undo shutdown
 ip address 2.2.2.1 255.255.255.0
 ospf ttz-enable 100
 ospf network-type p2p
 ospf enable 1 area 0.0.0.0
#
interface LoopBack0
 ip address 2.2.2.2 255.255.255.255
 ospf enable 1 area 0.0.0.0
#
ospf 1 router-id 2.2.2.2
 opaque-capability enable
 area 0.0.0.0
  ttz 100 edge
```

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
#
interface Ethernet3/0/0
 undo shutdown
 ip address 3.3.3.1 255.255.255.0
 ospf network-type p2p
 ospf enable 1 area 0.0.0.0
#
interface Ethernet3/0/2
 undo shutdown
 ip address 2.2.2.2 255.255.255.0
 ospf network-type p2p
 ospf enable 1 area 0.0.0.0
#
interface LoopBack0
 ip address 4.4.4.4 255.255.255.255
 ospf enable 1 area 0.0.0.0
#
ospf 1 router-id 4.4.4.4
 opaque-capability enable
 area 0.0.0.0
  ttz 100 internal
#
```

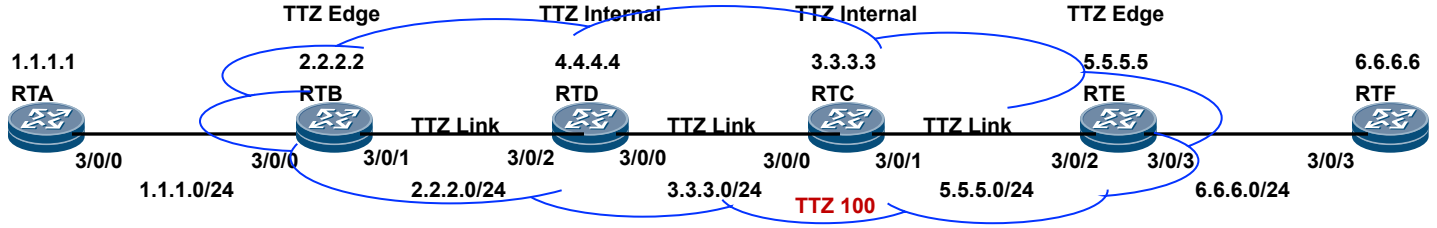

2: Prepare for migration (1 command)

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
#
ospf 1 router-id 2.2.2.2
 area 0.0.0.0
+  ttz state 100 prepare-migrate
#
[~RTB-ospf-1-area-0.0.0.0]
```

3: Migrate (1 command) after checking LSDB

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
#
ospf 1 router-id 2.2.2.2
 area 0.0.0.0
+  ttz state 100 migrate
#
[~RTB-ospf-1-area-0.0.0.0]
```

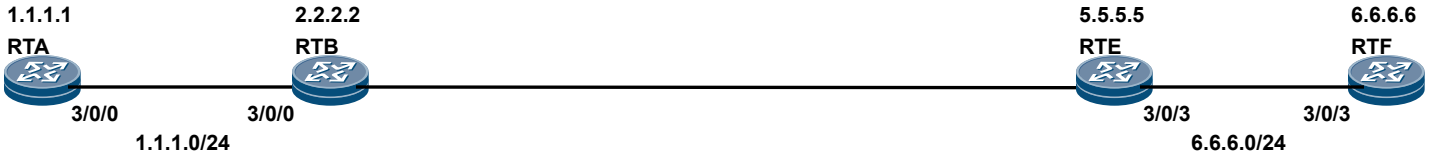

LSDBs during Preparation for Migration

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
  #
 ospf 1 router-id 2.2.2.2
  area 0.0.0.0
+  ttz state 100 prepare-migrate
  #
[~RTB-ospf-1-area-0.0.0.0]
```

Preparation for Migration:

✓TTZ peers are discovered

✓ TTZ link states are distributed within TTZ in Opaque LSAs

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
<RTC>display ospf lsdb

         OSPF Process 1 with Router ID 3.3.3.3
                 Link State Database

                         Area: 0.0.0.0
 Type      LinkState ID    AdvRouter       Age  Len   Sequence       Metric
 Router    1.1.1.1         1.1.1.1         288  60    80000087            1
 Router    2.2.2.2         2.2.2.2         470  84    80000085            1
 Router    3.3.3.3         3.3.3.3        1454  84    80000054            1
 Router    4.4.4.4         4.4.4.4         213  84    80000055            1
 Router    5.5.5.5         5.5.5.5         905  84    80000083            1
 Router    6.6.6.6         6.6.6.6         839  60    8000004f            1

             Type 10 Opaque (Area-Local Scope) Database
 Type      LinkState ID    AdvRouter       Age  Len   Sequence    Area
 Opq-Area  10.0.0.0        2.2.2.2          22  96    80000001    0.0.0.0
 Opq-Area  10.0.0.0        3.3.3.3          17  96    80000002    0.0.0.0
 Opq-Area  10.0.0.0        4.4.4.4          18  96    80000002    0.0.0.0
 Opq-Area  10.0.0.0        5.5.5.5          21  96    80000001    0.0.0.0
```

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
<RTE>display ospf lsdb

         OSPF Process 1 with Router ID 5.5.5.5
                 Link State Database

                         Area: 0.0.0.0
 Type      LinkState ID    AdvRouter       Age  Len   Sequence       Metric
 Router    1.1.1.1         1.1.1.1         532  60    80000087            1
 Router    2.2.2.2         2.2.2.2         837  84    80000085            1
 Router    3.3.3.3         3.3.3.3        2478  84    80000054            1
 Router    4.4.4.4         4.4.4.4         409  84    80000055            1
 Router    5.5.5.5         5.5.5.5        1560  84    80000083            1
 Router    6.6.6.6         6.6.6.6        1449  60    8000004f            1

             Type 10 Opaque (Area-Local Scope) Database
 Type      LinkState ID    AdvRouter       Age  Len   Sequence    Area
 Opq-Area  10.0.0.0        2.2.2.2          88  96    80000001    0.0.0.0
 Opq-Area  10.0.0.0        3.3.3.3          83  96    80000002    0.0.0.0
 Opq-Area  10.0.0.0        4.4.4.4          84  96    80000002    0.0.0.0
 Opq-Area  10.0.0.0        5.5.5.5          87  96    80000001    0.0.0.0

          Type 9 Opaque (Link-Local Scope) Database.  Area: 0.0.0.0
 Type      LinkState ID    AdvRouter         Age  Len   Sequence   Interfaces
 Opq-Link  4.0.0.0         3.3.3.3            88  40    80000002   0.0.0.0
 Opq-Link  4.0.0.0         5.5.5.5            87  40    80000001   0.0.0.0
<RTE>
```

LSDBs after Migration to TTZ

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
#
ospf 1 router-id 2.2.2.2
 area 0.0.0.0
+  ttz state 100 migrate
#
[~RTB-ospf-1-area-0.0.0.0]
```

RTA | RTB | RTD | RTC | RTE | RTF

```
ay ospf lsdb router 5.5.5.5

)SPF Process 1 with Router ID 1.1.1.1
                Area: 0.0.0.0
        Link State Database

Type      : Router
Ls id     : 5.5.5.5
Adv rtr   : 5.5.5.5
Ls age    : 150
Len       : 120
Options   :  E
seq#      : 80000084
chksum    : 0x5141
Link count: 8
   Link ID: 2.2.2.2
   Data   : 0.0.0.8
   Link Type: P-2-P
   Metric : 3
   Link ID: 1.2.2.1
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 4
   Link ID: 3.3.3.3
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 1
   Link ID: 4.4.4.4
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 2
   Link ID: 6.6.6.6
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 1
   Link ID: 6.6.6.6
   Data   : 6.6.6.1
   Link Type: P-2-P
   Metric : 1
   Link ID: 6.6.6.0
   Data   : 255.255.255.0
```

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
<RTA>display ospf lsdb router 2.2.2.2

        OSPF Process 1 with Router ID 1.1.1.1
                Area: 0.0.0.0
        Link State Database

Type      : Router
Ls id     : 2.2.2.2
Adv rtr   : 2.2.2.2
Ls age    : 131
Len       : 120
Options   :  E
seq#      : 80000086
chksum    : 0x8a50
Link count: 8
   Link ID: 5.5.5.5
   Data   : 0.0.0.7
   Link Type: P-2-P
   Metric : 3
   Link ID: 1.2.2.1
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 1
   Link ID: 3.3.3.3
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 2
   Link ID: 4.4.4.4
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 1
   Link ID: 6.6.6.6
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 4
   Link ID: 1.1.1.1
   Data   : 1.1.1.2
   Link Type: P-2-P
   Metric : 1
   Link ID: 1.1.1.0
   Data   : 255.255.255.0
   Link Type: StubNet
```

- ✓ Router LSA for virtualizing TTZ originated from TTZ edge routers
- ✓P2P link to each of the other edges in LSA
- ✓ Loopback addresses in TTZ as stubs in LSA

Routing Tables after Migration to TTZ

- ✓ No route change on TTZ routers
- ✓ No change in any route on non TTZ router
- ✓ Routes for nets in TTZ disappear on non TTZ router

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
<RTC>display ospf routing

        OSPF Process 1 with Router ID 3.3.3.3
                Routing Tables

 Routing for Network
 Destination        Cost      Type       NextHop         AdvRouter
 1.1.1.0/24         3         Stub       3.3.3.1         2.2.2.2
 1.2.2.1/32         3         Stub       3.3.3.1         1.1.1.1
 2.2.2.0/24         2         Stub       3.3.3.1         4.4.4.4
 2.2.2.2/32         2         Stub       3.3.3.1         2.2.2.2
 3.3.3.0/24         1         Direct     3.3.3.2         3.3.3.3
 3.3.3.3/32         0         Direct     3.3.3.3         3.3.3.3
 4.4.4.4/32         1         Stub       3.3.3.1         4.4.4.4
 5.5.5.0/24         1         Direct     5.5.5.1         3.3.3.3
 5.5.5.5/32         1         Stub       5.5.5.2         5.5.5.5
 6.6.6.0/24         2         Stub       5.5.5.2         5.5.5.5
 6.6.6.6/32         2         Stub       5.5.5.2         6.6.6.6
```

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
<RTF>display ospf routing

        OSPF Process 1 with Router ID 6.6.6.6
                Routing Tables

 Routing for Network
 Destination        Cost      Type       NextHop         AdvRouter
 1.1.1.0/24         5         Stub       6.6.6.1         2.2.2.2
 1.2.2.1/32         5         Stub       6.6.6.1         1.1.1.1
 2.2.2.0/24         4         Stub       6.6.6.1         4.4.4.4
 2.2.2.2/32         4         Stub       6.6.6.1         2.2.2.2
 3.3.3.0/24         3         Stub       6.6.6.1         3.3.3.3
 3.3.3.3/32         2         Stub       6.6.6.1         3.3.3.3
 4.4.4.4/32         3         Stub       6.6.6.1         4.4.4.4
 5.5.5.0/24         2         Stub       6.6.6.1         5.5.5.5
 5.5.5.5/32         1         Stub       6.6.6.1         5.5.5.5
 6.6.6.0/24         1         Direct     6.6.6.2         6.6.6.6
 6.6.6.6/32         0         Direct     6.6.6.6         6.6.6.6

 Total Nets: 11
 Intra Area: 11  Inter Area: 0  ASE: 0  NSSA: 0
<RTF>
```

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
        OSPF Process 1 with Router ID 3.3.3.3
                Routing Tables

 Routing for Network
 Destination        Cost      Type       NextHop         AdvRouter       Area
 1.1.1.0/24         3         Stub       3.3.3.1         2.2.2.2         0.0.0.0
 1.2.2.1/32         3         Stub       3.3.3.1         1.1.1.1         0.0.0.0
 2.2.2.0/24         2         Stub       3.3.3.1         4.4.4.4         0.0.0.0
 2.2.2.2/32         2         Stub       3.3.3.1         2.2.2.2         0.0.0.0
 3.3.3.0/24         1         Direct     3.3.3.2         3.3.3.3         0.0.0.0
 3.3.3.3/32         0         Direct     3.3.3.3         3.3.3.3         0.0.0.0
 4.4.4.4/32         1         Stub       3.3.3.1         4.4.4.4         0.0.0.0
 5.5.5.0/24         1         Direct     5.5.5.1         3.3.3.3         0.0.0.0
 5.5.5.5/32         1         Stub       5.5.5.2         5.5.5.5         0.0.0.0
 6.6.6.0/24         2         Stub       5.5.5.2         5.5.5.5         0.0.0.0
 6.6.6.6/32         2         Stub       5.5.5.2         6.6.6.6         0.0.0.0

 Total Nets: 11
```

10.193.46.23 | RTA | RTB | RTD | RTC | RTE | RTF

```
        OSPF Process 1 with Router ID 6.6.6.6
                Routing Tables

 Routing for Network
 Destination        Cost      Type       NextHop         AdvRouter       Area
 1.1.1.0/24         5         Stub       6.6.6.1         2.2.2.2         0.0.0.0
 1.2.2.1/32         5         Stub       6.6.6.1         5.5.5.5         0.0.0.0
 2.2.2.2/32         4         Stub       6.6.6.1         2.2.2.2         0.0.0.0
 3.3.3.3/32         2         Stub       6.6.6.1         5.5.5.5         0.0.0.0
 4.4.4.4/32         3         Stub       6.6.6.1         5.5.5.5         0.0.0.0
 5.5.5.5/32         1         Stub       6.6.6.1         5.5.5.5         0.0.0.0
 6.6.6.0/24         1         Direct     6.6.6.2         6.6.6.6         0.0.0.0
 6.6.6.6/32         0         Direct     6.6.6.6         6.6.6.6         0.0.0.0

 Total Nets: 8
 Intra Area: 8  Inter Area: 0  ASE: 0  NSSA: 0
<RTF>
```

Migrate to TTZ vs Split Area

	Migrate to TTZ	Split Area
Operations (configuration, migration)	**Easy** (N nodes to a TTZ, E edges with T TTZ links) # CLIs = N + E*T + 2 = 122 (for N=100, E=4, T=5)	**Complex** (big network architecture change) (N nodes to an area, each node has L Links) # CLIs = N + N*(1 or L) = 200 or 1100 (for N=100, L=10)
Interruptions on services?	No	Yes (change area for links and nodes. Links down and up)

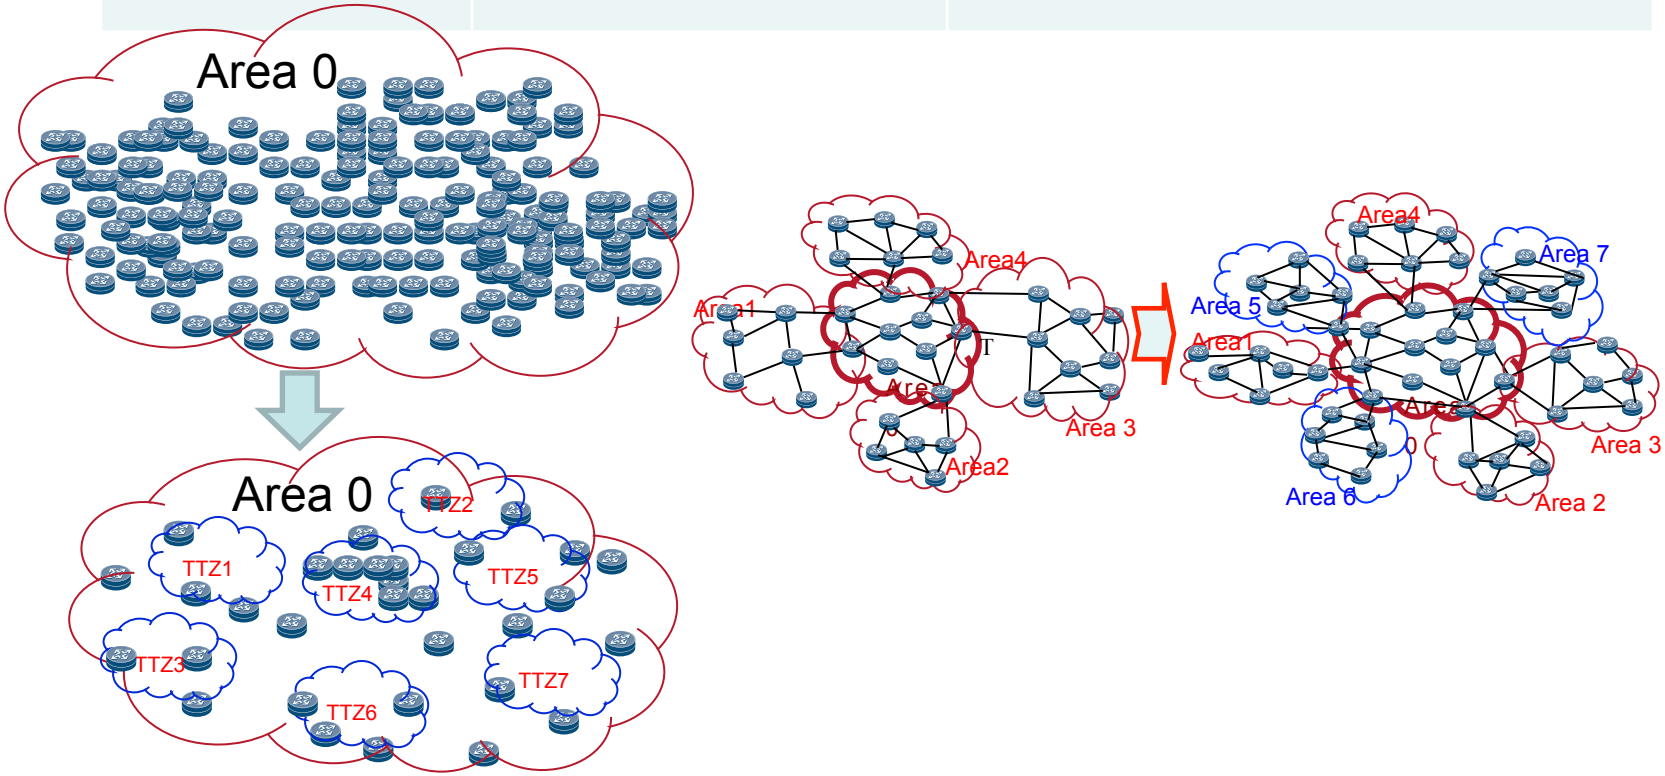

Implementation Experience

1. Relatively easy to implement TTZ
2. One topology for both

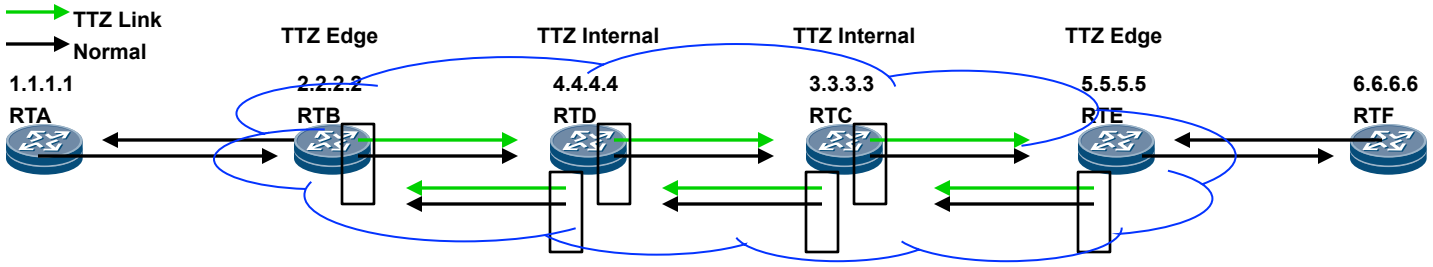

One Topology for Both TTZ links and normal links (during migration)

For each TTZ link in a TTZ opaque LSA, there is an additional flag to differentiate between a TTZ link and a Normal link.

3. Normal links are used before migration to TTZ
4. TTZ/normal links are used after migration
5. Reverse link check with additional flag match

Conclusions

- **Relatively easy to implement**
- **Smooth migration [No service interruption]**
- **Simple steps for migration [Simple operation and easy to maintain]**
- **Virtualizing a big number of nodes to a few [Improves Scalability]**

Next Step

Welcome comments