

A Resolution To The Equal-Cost Multipath (ECMP) Failure In The OSPF Not-So-Stubby-Areas (NSSA)

Abhimanyu Mehta, Aparna Bhuktar, Kartik Santaluri, Sagar Karale

Abhimanyu.mehta@colorado.edu, aparna.bhuktar@colorado.edu,
kartik.santaluri@colorado.edu sagar.karale@google.com

A capstone paper submitted as partial fulfillment of the requirements for the degree of Masters in Interdisciplinary Telecommunications at the University of Colorado, Boulder, 4 May 2012.
Project directed by Prof. Jose Santos

1. Introduction

The network architecture of most datacenters today consists of the Core-Distribution-Access layer setup. Servers at the access layer, apart from having mirrored backups also have redundant links for reachability. The capacity requirement is increasing at breakneck pace considering the amount of data that flows through these links. For maximum efficiency to be reached all links need to be used. Equal Cost Multi Path (ECMP) is a necessity to achieve thorough proficiency. The use of single OSPF areas permits us to utilize all links effectively; however increased database sizes lead to stability issues. Our solution of achieving ECMP with the use of NSSA areas can result in the same effectiveness while reducing the database per area considerably. This could help a network architect to achieve a much-reduced FIB table for devices in the network thereby eliminating stability issues such as CPU and memory.

OSPF is one of the more popular interior gateway routing protocols used by service providers and enterprise networks. Being a scalable routing protocol, OSPF can be split into multiple areas with every area connecting to a single area called the “backbone area.” [1]. Also being a link-state protocol, OSPF maintains a database that has information about the links, their cost (metric), subnet information etc. This information needs to be exchanged through advertisements called Link State Advertisements between the routers running the OSPF routing protocol (LSA flooding) [1]. The excessive flooding due to the burgeoning network makes the data exchange between the routers difficult thereby causing congestion in the network [2]. This issue however is resolved by segmenting the entire OSPF domain into logical areas such as Normal Area, Stubby Area, Not-So-Stubby-Area (NSSA), and the Totally Stubby Area, which are used to limit the routing information (i.e. the number of LSA’s) entering the area. There are multiple LSAs, which are used to carry different information, that are defined in the OSPF routing protocol. The Type 3: Network Summary LSA, the Type 4: Autonomous System Border Router LSA, the Type 5: External LSA and the Type 7: NSSA External LSA are some of the LSAs that we will be using in this paper.

The NSSA area that connects networks running in different routing protocols to those in OSPF can have multiple equal cost routes to the external networks. These equal cost (ECMP) routes can be used to load balance the traffic, however that ECMP is broken at the ABR due to

the conversion of type 7 LSAs to type 5 thereby rendering the other links useless. This paper proposes a solution to retain the ECMP to maintain efficient equal path load balancing.

2. Basic Concepts

- OSPF Areas: A logical cluster of contiguous networks including the interfaces of routers and hosts participating in those networks is called an Area. Each area has its own link state database [1].
- Area Border Router (ABR): All areas must connect physically or virtually to a particular area named backbone or Area 0. A router connecting multiple OSPF areas is called the ABR [1].
- Autonomous System (AS): A set of routers exchanging routing information via a common routing protocol [1].
- Autonomous System Border Router (ASBR): A router running OSPF as well as a different routing protocol [3].
- Equal-Cost Multi path (ECMP): Multiple paths available to reach a particular destination with each path having the same metric or cost [1].
- NSSA: Not So Stubby Areas are the areas that run other routing protocols including OSPF. In the NSSA only a type 3 LSA and a default route is injected whereas the type 4 and 5 LSAs are blocked from entering the NSSA[3].
- Type 5 / External LSA: Link State Advertisements that are generated by ASBR for each destination that is external to the AS. These are flooded throughout the entire AS except the stub areas [1].
- Type 7 / NSSA External LSA: These are similar to Type-5 LSAs with two major differences: first, they originate and are advertised by the ASBR throughout the NSSA, and second, they are advertised only within a single NSSA and not the backbone area [3].
- Forwarding Address (FA): When the network between the ASBR and the adjacent AS is advertised into OSPF as an internal route, the next hop address is the forwarding address [3].
- Redistribution: The concept of advertising different routing protocols into a particular routing protocol [4].

3. Why NSSA

The primary advantage that NSSA provides over regular areas is better stability in larger networks and an improvement in the routing efficiency. This can be illustrated by considering an example large datacenter network that has 100-1000 network devices including routers, switches and servers [5].

A typical three-tier model would consist of 48-128 port access switches connecting approximately 25000 hosts [6]. As the number of servers connected to an access switch increases the subnets representing the servers also increases, causing more IGP routes to flow in the datacenter. Assuming the subnets are discontinuous, the access switches cannot summarize the routes, thus creating small prefixes and an increase in the size of the routing table on all routers up to the core. In case of OSPF, these increasing routes are nothing but the increasing number of LSAs which not only increase the RAM, Forwarding Information Base (FIB) and CPU requirements on the routers but also hamper the convergence time of the routing protocol. If you try to allocate more memory to OSPF to handle the growing LSA using “maximum-number-of-lsa <lsa type>” command, other memory-intensive protocols, for instance Border Gateway Protocol (BGP) running on the same router, will have to pay the price [7].

To advertise external networks on an ASBR, we use a concept called redistribution. This redistribution generates Type 5 LSAs in a regular area [1]. These Type 5 LSAs cannot be filtered on any Area Border Router (ABR) and are therefore injected in all the areas. Since the majority of the routes in a datacenter are server routes, these routes increase the size of the routing table drastically.

NSSA helps to address the above issues. By its nature, NSSA blocks the Type 4 and 5 LSAs from being injected into itself thereby blocking the majority of the external routes. This relaxes the routers inside the NSSA greatly as now the memory can be shared between the different routing protocols or could be used in implementing faster switching and traffic engineering solutions such as Multi-Protocol Label Switching (MPLS) and Quality of Service (QoS) [8].

4. ECMP

The original aim behind the formation of the IP data layer was to ensure that there was absolute sturdiness. The basic idea was to ensure that there is a fault tolerant network, which would recover immediately in case of any failures within the network, as long as there existed enough physical connectivity. However the one shortcoming of that was that the subsequent rerouting process was extremely slow for today's packet networks. One suggested way to improve the re-routing process was to adjust the routing protocol timers, however even this was evidently too slow for bandwidth consuming applications, which were mission critical such as multimedia [9]. This caused a necessary requirement in concepts to speed up the recovery performances of data networks. This is what brought the ECMP feature into existence. ECMP load balances flows across multiple IP paths [10]. This feature, available to various routing protocols allowed a distribution of packets on various outgoing links between the source and the destination.

Thus with ECMP, OSPF can route packets in a round robin fashion to the destination via multiple equal cost paths. Further, this can be observed on a router by issuing the "show ip route" command, which would list multiple next hop addresses [10].

5. Research Problem

In reality, an enterprise or a service provider's network can have different routing protocols running and as discussed earlier, in order to achieve end-to-end connectivity to all destinations the concept of redistribution is used. This redistribution of routes from different routing protocols into OSPF gives rise to our research problem. To illustrate the research problem we have used an example network as shown below.

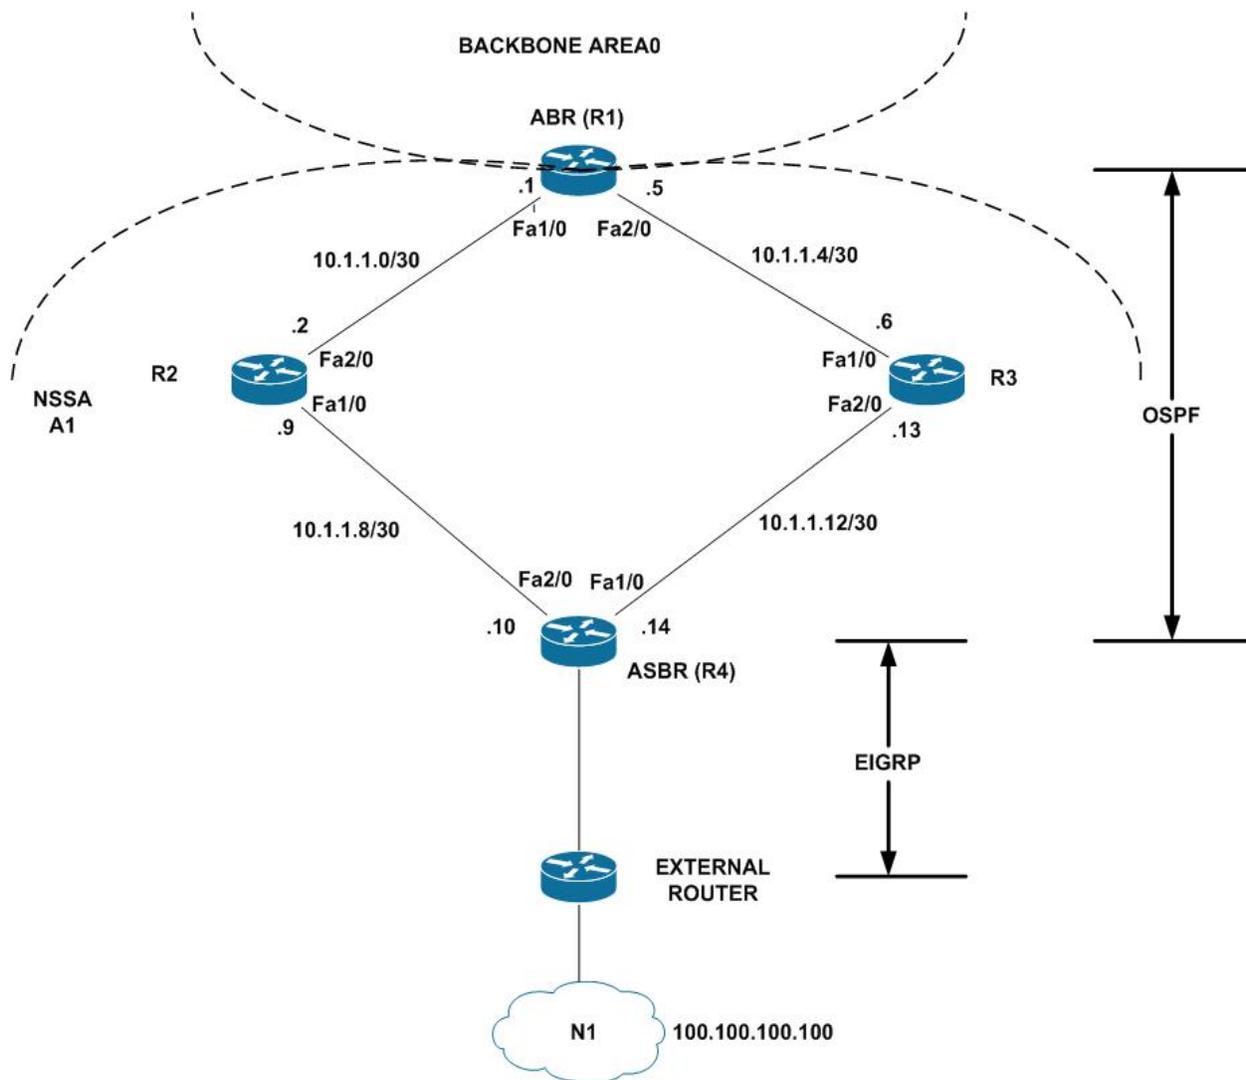


Figure 1: Sample OSPF topology using NSSA area

In Figure 1, R1 is the ABR that connects two OSPF areas, backbone area (A0) and NSSA area (A1). R2 and R3 routers are internal routers in the NSSA. R4 is the ASBR that runs OSPF with the internal routers, R2 and R3, and EIGRP with the External router. Network N1 (100.100.100.100) represents the external network that needs to be redistributed into OSPF on the ASBR so that all the OSPF speaking routers know how to reach this network. This redistribution on the ASBR causes a Type 7 LSA to be generated within the NSSA area. The ABR converts the Type 7 LSA to a Type 5 LSA and floods it across to all the other areas (in this case area 0) [11]. Inside the Type 7 LSA, there exists a field called Forwarding Address (FA), which is a non-zero value that leads the ABR to the ASBR and eventually to the external network [12]. This FA is the IP address of one of the physical interfaces on the ASBR. In this case, it will be either 10.1.1.10 or 10.1.1.14. Let us assume the FA in this case is 10.1.1.14. The ABR will search for all possible shortest paths to reach the FA. Ideally in this case, we would want the ABR to use both the paths i.e. through R2 and R3 to reach the ASBR. However, the FA being 10.1.1.14, the ABR will have only a single shortest path to reach the ASBR i.e. through

R3. OSPF determines the best path to a destination by using the cost associated with all the links along the path [13]. In our example, the cost to reach the FA (10.1.1.14) from the ABR is explained as follows (assuming the cost of each link is configured to be 1):

$$\text{Cost via R3} = \text{R1's fa2/0} + \text{R3's fa2/0} = 1 + 1 = 2$$

$$\text{Cost via R2} = \text{R1's fa1/0} + \text{R2's fa1/0} + \text{R4's Fa1/0} = 1 + 1 + 1 = 3 \quad (1)$$

Thus, there is only a single shortest path from the ABR to the FA and therefore to the external network, leaving the other path through R2 idle defeating the very purpose of ECMP. This inefficient routing overburdens the only link being utilized to reach the external network and leaves the rest of the links underutilized. This not only impacts the efficiency but also the economics involved in the network, as the cost for running the underutilized link is not recovered. Assuming an 8 port-10G access switch worth \$25,000 with an approximate cost/port equal to \$3,000 we might have 4 uplinks from this 10G switch out of which only one is being utilized. This causes the enterprise to lose about 30Gigs worth of bandwidth as well as $\$3,000 \times 3 = \$9,000$ [14]. Imagine such a loss in a datacenter on hundreds of switches.

6. Research Question and Hypothesis

The main purpose of this paper is to address the problem of the ECMP failure on the ABRs in NSSAs. Therefore our research question is as follows: Is there a solution to retain the ECMP on the ABRs in NSSAs?

Now that we understand the cause behind the failure of ECMP to be the FA, we wanted to understand if changing the value of that FA to a particular value would help in retaining the ECMP. Based on this, our research is divided into the following sub questions: Can the “forwarding address” field present in the Type 7 LSA be used to route packets from the ABR to reach the external network while maintaining ECMP? What are the possible addresses this field can take?

This research was conducted after taking into consideration certain assumptions. The test-bed was laid out in the Access-Distribution-Core format (Figure 2). The hosts, representing the external networks were connected to the access layer and routed through the Distribution and Core layers. The routers used in the network were Cisco 3640s and running OSPF as the interior routing protocol. The ASBR was running OSPF as well as EIGRP that connected to the external networks. All the links were Fast Ethernet and no serial interfaces were used. The cost of all the links in the test network is configured to be 1.

7. Audience & Scope

The intended audience for this paper is enterprise and service providers who have deployed or are in the phase of implementation of OSPF as an interior routing protocol for their networks. From an academic perspective this paper could also benefit doctorates or scholars who intend to conduct research at a packet level and develop or propose an algorithm in order to make this solution acceptable. The experiments, results and observations are only limited to the NSSAs in OSPF. Other fields inside the type 7 LSA such as the N bit, P bit, E bit, and the different arguments involved in the redistribution process are out of scope of this paper.

8. Previous Research & their Relevance

This research is important for efficient and effective routing. For the most part, previous studies on OSPF have been either model or simulation based [8], [15] or have concentrated on measuring OSPF implementation behavior in a single area [16], [18]. The research topic by A. Shaikh et al. highlighted key points such as the scale of deployment of OSPF in an enterprise network, the effects of LSAs and the need to control them when the network grows, the importance of segregation of the networks, and different types of LSAs [15]. However, the paper fails to describe NSSA areas and equal cost routing in an enterprise network. Another paper by Dzida et al. [17] deals with ECMP routing to derive the metrics used by the routing protocol to optimize the load balancing, but the paper does not consider the ECMP routing inside the NSSA. Our research aims to find a solution to load balance the traffic inside the NSSA thereby eliminating the limitations from the papers mentioned above.

9. Our Approach

We reproduced the problem in the lab and worked towards answering our research questions. As explained in the research problem section, the FA was the address of one of the physical interfaces on the ASBR. With reference to Figure 2, the ASBR has five physical interfaces, out of which four are running OSPF (OSPF-enabled interfaces) and one is running the EIGRP routing protocol. Upon redistributing the OSPF networks into EIGRP and vice-versa on the ASBR, the FA will be selected from any one of these OSPF-enabled physical interfaces thus breaking the ECMP (as explained in the Research Problem section) [3]. In order to avoid this, the first step was to ensure that the FA is selected as the configured loopback address on the ASBR instead of the physical address of one of the OSPF enabled ASBR interfaces. Now the cost to reach the loopback interfaces on R8 from the ABR will be equal via R3, R4, R5, and R6. The math involved in the cost calculation is as follows:

$$\begin{aligned} \text{Cost via R3} &= \text{R1/R2's fa1/0} + \text{R3's fa2/0} + \text{R7's loopback} = 1 + 1 + 1 = 3 \\ \text{Cost via R4} &= \text{R1/R2's fa2/0} + \text{R4's fa3/0} + \text{R7's loopback} = 1 + 1 + 1 = 3 \\ \text{Cost via R5} &= \text{R1's Fa2/1 /R2's fa3/0} + \text{R5's fa1/1} + \text{R7's loopback} = 1 + 1 + 1 = 3 \\ \text{Cost via R6} &= \text{R1's Fa3/0 /R2's fa3/1} + \text{R6's fa1/1} + \text{R7's loopback} = 1 + 1 + 1 = 3 \end{aligned} \quad (2)$$

By achieving the FA to be the loopback of the ASBR, we now have four equal cost paths to reach the external networks thus helping us retain the ECMP inside the NSSA.

The OSPF neighborship was formed between, R3 & R7, R4 & R7, R5 & R7, and R6 & R7 over the loopback interfaces. In order to achieve this, we had to encapsulate OSPF packets inside the IP protocol. This encapsulation is possible through a concept called 'tunnels'. There are several types of tunnels such as Generic Routing Encapsulation (GRE), Manually Configured Tunnel (MCT), 6 to 4 tunnel, Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) etc. [19]. However, the links between the routers in the test network are point-to-point links therefore it eliminates the option of 6 to 4 and ISATAP tunnels since they are multipoint tunnels [20]. Also, we required the tunnels to carry OSPF packets and exchange the OSPF routes with other routers. GRE tunnels meet our requirement as they help to encapsulate a protocol (OSPF) inside a GRE packet and send it across the network [20]. Therefore, the choice for the type of tunnel used was narrowed down to GRE.

10.Methodology

The research was a set of experiments conducted in the lab that involved configuring OSPF on eight Cisco 3640 routers. The two ABRs, ABR1 and ABR2, were configured to be a part of the backbone area (Area 0) and the NSSA area (Area 10). Our scenario is depicted in the network diagram below:

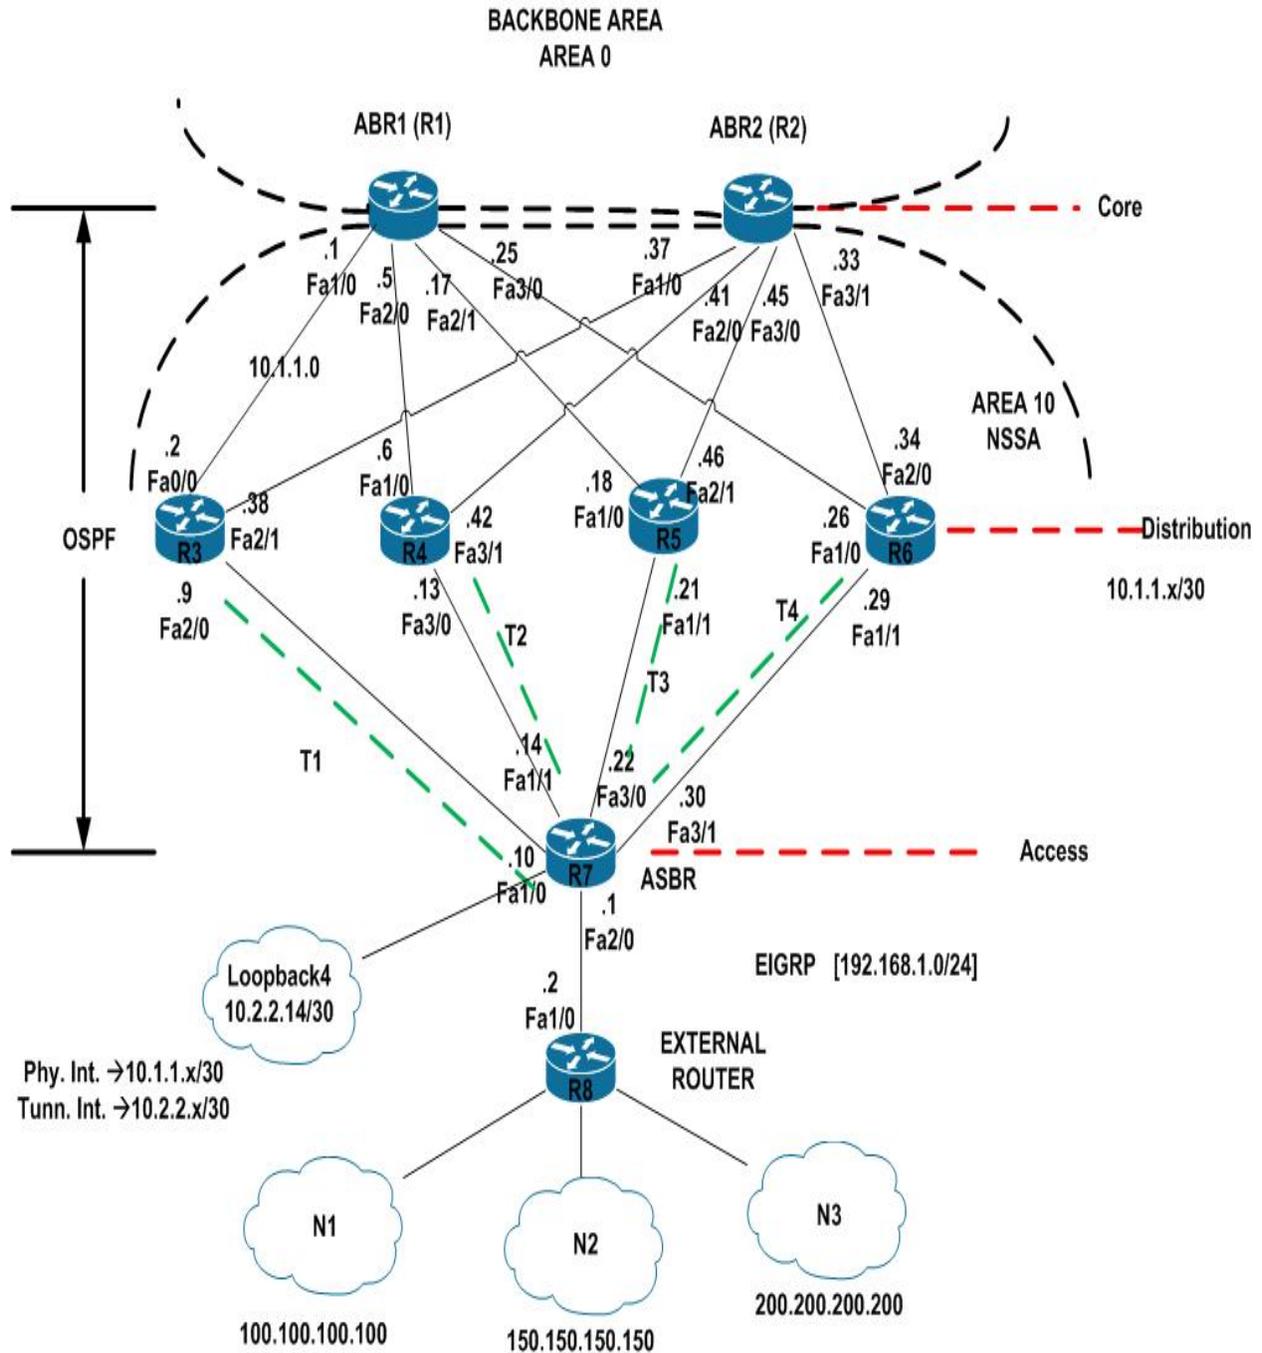


Figure 2: Test network for the research

R3, R4, R5, and R6 are the internal routers (IR) inside the NSSA area (Area 10). The ABRs, the IRs and the ASBR were running the OSPF routing protocol. The ASBR was also running the EIGRP routing protocol with the External Router (R8). Net1, Net2, and Net3 are the external networks that are advertised in EIGRP. GRE tunnels were configured between R3-R7, R4-R7, R5-R7, and R6-R7 and OSPF was enabled on the tunnel interfaces. The IP address of the tunnel was configured using “*ip unnumbered loopback<number>*” command. The source and destination address on the tunnels was the physical address of the respective routers. The ip unnumbered interface command takes the IP address from the interface specified in the configuration to the interface on which it is configured. This results in both interfaces having the same the IP address [21]. We used this concept to ensure that the tunnels have the same ip address as the loopback on the ASBR which would in turn result in our FA using the highest loopback (tunnel) ip address, thereby change the value of the FA.

11. Configurations

Having spoken about our approach and methodology it is only pertinent to include some of our major configurations.

As mentioned earlier we had created four loopbacks on the ASBR (R7) and one each on the internal routers (i.e. R3, R4, R5 and R6). We used these loopbacks on the ASBR so that we could use them as the tunnel ips with the help of the *ip unnumbered loopback <number>* command. This would help us in achieving our objective of changing the Forwarding Address to a loopback ip, which would in turn help us achieve the ECMP. Another observation that was made was that the Forwarding Address obtained the highest loopback ip on the ASBR as its address. Therefore the configurations included only show that of Loopback 4.

To begin with, the important configurations are as follows: -

ASBR (R7)

```
interface Loopback4
ip address 10.2.2.14 255.255.255.252
!
interface Tunnel4
ip unnumbered Loopback4.....This sets the ip of tunnel 4 as 10.2.2.14
ip ospf cost 1.....This sets the ospf cost of the tunnel to 1
tunnel source 10.1.1.30.....This creates the tunnel on the physical interface to R6
tunnel destination 10.1.1.29.....This sets the tunnel destination to the physical interface on R6
!
interface FastEthernet3/1
ip address 10.1.1.30 255.255.255.252
ip ospf cost 1.....This sets the ospf cost of the physical interface to 1
duplex auto
speed auto
!
router eigrp 1
network 192.168.1.0.....Here we are advertising the interface between R7 and R8 in !EIGRP
auto-summary
```

```

!
router ospf 1
router-id 7.7.7.7
log-adjacency-changes
area 10 nssa.....Defining area 10 as an NSSA (This is done on all Area 10 routers)
redistribute eigrp 1 metric-type 1 subnets....redistributing eigrp AS -1 into OSPF
network 10.2.2.0 0.0.0.3 area 10
network 10.2.2.4 0.0.0.3 area 10
network 10.2.2.8 0.0.0.3 area 10
network 10.2.2.12 0.0.0.3 area 10....These 4 statements advertise only the tunnels in OSPF area
! 10 (Note we are not advertising the physical interfaces)

```

R6

```

interface Loopback1
ip address 10.2.2.13 255.255.255.252
!
interface Tunnel1
ip unnumbered Loopback1.....This sets the ip of tunnel 1 as 10.2.2.13
ip ospf cost 1.....This sets the ospf cost of the tunnel to 1
tunnel source 10.1.1.29.....This creates the tunnel on the physical interface to R7
tunnel destination 10.1.1.30.....This sets the tunnel destination to the physical interface on R7
!
interface FastEthernet1/1
ip address 10.1.1.29 255.255.255.252
ip ospf cost 1.....This sets the ospf cost of the physical interface to 1
duplex auto
speed auto
!
router ospf 1
router-id 6.6.6.6
log-adjacency-changes
area 10 nssa.....Defining area 10 as an NSSA
network 10.1.1.24 0.0.0.3 area 10
network 10.1.1.32 0.0.0.3 area 10....These two statements put the uplinks to the ABR's (R1 and
! R2 in OSPF
network 10.2.2.12 0.0.0.3 area 10....This puts the tunnel 1 in OSPF

```

Note that the configs on R3, R4 and R5 are similar to R6, but I've only shown those for R6 because the Forwarding Address gets selected as the highest Loopback ip on the ASBR (R7), which also happens to be the tunnel ip between R7 and R6. That being said, the tunnels created between R3-R7, R4-R7, R5-R7 and R6-R7, provide us ECMP thereby helping us use all 4 links between the internal routers and the ASBR. This can be seen in th outputs obtained in the next section.

12.Results & Findings

After having set up our network as explained above, we observed the following results. The Type 7 LSA generated on the ASBR (R7) before and after creating the tunnel is as follows:

```
R7#sh ip ospf database nssa-external

      OSPF Router with ID (7.7.7.7) (Process ID 1)

      Type-7 AS External Link States (Area 10)

LS age: 104
Options: (No TOS-capability, Type 7/5 translation, DC)
LS Type: AS External Link
Link State ID: 100.0.0.0 (External Network Number )
Advertising Router: 7.7.7.7
LS Seq Number: 80000002
Checksum: 0xF5C
Length: 36
Network Mask: /8
  Metric Type: 1 (Comparable directly to link state metric)
  TOS: 0
  Metric: 20
  Forward Address: 10.1.1.30
  External Route Tag: 0

R7#sh ip ospf database nssa-external

      OSPF Router with ID (7.7.7.7) (Process ID 1)

      Type-7 AS External Link States (Area 10)

LS age: 990
Options: (No TOS-capability, Type 7/5 translation, DC)
LS Type: AS External Link
Link State ID: 100.0.0.0 (External Network Number )
Advertising Router: 7.7.7.7
LS Seq Number: 80000005
Checksum: 0x4135
Length: 36
Network Mask: /8
  Metric Type: 1 (Comparable directly to link state metric)
  TOS: 0
  Metric: 20
  Forward Address: 10.2.2.14
  External Route Tag: 0
```

Figure 3: Forwarding address in Type 7 LSA before and after the tunnel creation respectively.

We can see that the Forwarding Address for the external network 100.0.0.0 in the first case is 10.1.1.30, which is a physical address on the ASBR. In the second output, i.e. after using the GRE tunnels, the FA is the highest loopback address, 10.2.2.14 of the ASBR. This change in the FA changes the number of paths available to reach the external networks from the ABRs. This can be observed in the following output.

```

ABR-R1#sh ip route 100.100.100.100
Routing entry for 100.0.0.0/8
  Known via "ospf 1", distance 110, metric 22, type extern 1
  Last update from 10.1.1.26 on FastEthernet3/0, 00:03:36 ago
  Routing Descriptor Blocks:
  * 10.1.1.26, from 2.2.2.2, 00:03:36 ago, via FastEthernet3/0
    Route metric is 22, traffic share count is 1

ABR-R1#sh ip route 150.150.150.150
Routing entry for 150.150.0.0/16
  Known via "ospf 1", distance 110, metric 22, type extern 1
  Last update from 10.1.1.26 on FastEthernet3/0, 00:03:41 ago
  Routing Descriptor Blocks:
  * 10.1.1.26, from 2.2.2.2, 00:03:41 ago, via FastEthernet3/0
    Route metric is 22, traffic share count is 1

ABR-R1#sh ip route 200.200.200.200
Routing entry for 200.200.200.0/24
  Known via "ospf 1", distance 110, metric 22, type extern 1
  Last update from 10.1.1.26 on FastEthernet3/0, 00:03:47 ago
  Routing Descriptor Blocks:
  * 10.1.1.26, from 2.2.2.2, 00:03:47 ago, via FastEthernet3/0
    Route metric is 22, traffic share count is 1

ABR-R1#

```

Figure 4: Single path to reach external networks

As seen above in Figure 4., the ABR had a single route with a metric of 22 to reach the external networks, before creating the tunnels. However, after configuring the GRE tunnels we observed four equal cost paths to reach the same external networks as seen in the output below (Figure 5.): -

```

ABR-R1#sh ip rout 100.100.100.100
Routing entry for 100.0.0.0/8
  Known via "ospf 1", distance 110, metric 23, type extern 1
  Last update from 10.1.1.6 on FastEthernet2/0, 00:08:14 ago
  Routing Descriptor Blocks:
  10.1.1.26, from 2.2.2.2, 00:08:14 ago, via FastEthernet3/0
    Route metric is 23, traffic share count is 1
  10.1.1.18, from 2.2.2.2, 00:08:14 ago, via FastEthernet2/1
    Route metric is 23, traffic share count is 1
  * 10.1.1.6, from 2.2.2.2, 00:08:14 ago, via FastEthernet2/0
    Route metric is 23, traffic share count is 1
  10.1.1.2, from 2.2.2.2, 00:08:14 ago, via FastEthernet1/0
    Route metric is 23, traffic share count is 1

```

Figure 5: Equal cost paths to reach external network, 100.100.100.100

Thus we could achieve ECMP in an NSSA by changing the FA to a loopback ip on the ASBR.

13. Conclusion

This paper finds a solution to retain the ECMP in the NSSA environment thereby achieving efficient routing. The Forwarding Address of the LSA was modified by using GRE tunnels so that the ABR could reach the external networks via multiple equal cost paths. Thus, this solution could further motivate enterprises and service providers to use NSSA areas that would help them in recovering the investments made on the existing infrastructure.

Although our solution is achieving the desired ECMP, it would still require tunnels to be configured between the ASBR and each internal router and on even bigger networks it could result in extremely large configurations. Therefore further research could be done on reducing these large tunnel configurations.

References

- [1] J. Moy, "OSPF Version 2," IETF RFC 2328," Apr, 1998.
- [2] T. Thomas, *OSPF network design solution*. Indianapolis, IN: Cisco Press, 1998.
- [3] P. Murphy, "The OSPF not-so-stubby (NSSA) area option," IETF RFC 3101, Jan, 2003.
- [4] Cisco Systems Inc. *Redistributing routing protocols* [Online]. Available: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009487e.shtml
- [5] A. Vahdat, "Scale and efficiency in data center networks," Dept. Elec. Eng., Stanford Univ., Stanford, CA, Nov. 2009.
- [6] M. Al-Fares, et al., "A scalable, commodity data center network architectures," in *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, Seattle, WA, Aug. 17-22, 2008.
- [7] OSPF Commands. Available FTP: ftp.hp.com Directory: pub/networking/software/ File: Bk3_Ch10_OSPF_cmds.pdf
- [8] A. Basu & J. G. Riecke, "Stability issues in OSPF routing," Dept. Comp. Sci. & Software Eng., Auburn Univ., Auburn, AL, Mar. 2003.
- [9] A. Iselt, et al., "Resilient routing using MPLS and ECMP," in *High performance switching and routing, 2004. HPSR. 2004 Workshop*, Phoenix, AZ, 2004, pp. 345-349.
- [10] Cisco Systems Inc. *Server cluster designs with Ethernet* [Online]. Available: http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_Infra2_5/DCInfra_3.html#wp1088785
- [11] C. S. I. "Understanding redistribution of OSPF routes into BGP, Document ID: 5242," 10 August 2005. [Online]. Available: <http://www.cisco.com/image/gif/paws/5242/bgp-ospf-redis.pdf>.
- [12] A. Zinin, "Link-State Routing Protocols," in *Cisco IP Routing: Packet forwarding and intra-domain routing protocols*. Boston, MA: Addison – Wesley, 2001.
- [13] W. Odom, "Routing protocols configuration and troubleshooting," in *CCNA ICND2 official exam certification guide*. Indianapolis, IN: Cisco Press, 2007, pp. 343-376.
- [14] A. Greenberg and D. A. Maltz, "What goes into a data center," in *SIGMETRICS 2009*, Jun. 19, 2009.
- [15] A. Shaikh, et al., "An OSPF topology server: design and evaluation," *IEEE J. Sel. Areas in Comm.*, vol. 20, no. 4, pp. 746-755, May, 2002.
- [16] A. Shaikh and A. Greenberg, "OSPF Monitoring: Architecture, design and deployment experience," *USENIX Symp. on Networked systems design and implementation*, San Francisco, CA, 2004.
- [17] M. Dzida, et al., "Optimization of the shortest-path routing with equal-cost multi path load balancing," in *Proceedings of Conference on Transparent Optical Networks International Conference, 2006*, Nottingham, UK, Jun. 18-22, 2006.
- [18] A. Shaikh, et al., "A case study of OSPF behavior in a large enterprise network," in *Proceedings of Internet Measurement Conference (IMC)*, Marseille, France, Nov. 6-8, 2002
- [19] D. Farinacci, et al., "Generic routing encapsulation (GRE)," IETF RFC 2784, Mar, 2000.
- [20] Cisco Systems Inc. (2001, June 7). *Implementing tunneling of IPv6* [Online]. Available: <http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-tunnel.html>
- [21] Cisco Systems Inc. *Understanding and configuring the ip unnumbered command* [Online]. Available: http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094e8d.shtml