Study and Analysis of Dynamic Routing Protocols in real time for Optimal Resource Utilization

[1]D.VENKATAVARA PRASAD, [2]SURESH JAGANATHAN, [3]G.P SAIKALYAN
[1&2]Department of Computer Science & Engineering
Sri Sivasubramaniya Nadar College of Engineering, Anna University
Chennai, Tamilnadu
INDIA
[3]Cisco Instructor & Engineer
Chennai, Tamilnadu
INDIA
dvvprasad@ssn.edu.in, sureshj@ssn.edu.in

Abstract: - Optimal selection of a Dynamic Routing Protocol plays a significant role in impacting the performance and resource utilization of a networking device. The suitability of the protocol is highly dependent in-terms of bandwidth usage and ability to manage increasing number of entries in the routing table. In this paper, three well-known dynamic routing protocols such as i)Routing Information Protocol (RIPv2), ii) Enhanced Interior Gateway Routing Protocol (EIGRP) and iii) Open shortest Path First protocol (OSPFv2) are studied and analyzed based on the bandwidth used during convergence, using real-time networking devices. In addition to that, testing their unequal-cost, load balancing capability in order to address the issue of choosing the optimal routing protocol based on the network specifications/requirements is also carried out.

Key-words:- Unequal-cost load balancing, Bandwidth, Convergence, Dynamic Routing Protocols, RIP, EIGRP, OSPF.

1 Introduction

Routing protocol is used by the router as a means of communication and dissemination of routing information among neighboring routers and deals with the status and reach of a set of prefixes the router is aware. These protocols make use of an algorithm that determines the best path to the destination network and enters that information into the routing table. It is always a challenge to choose the right protocol for a network that satisfies the given requirements and efficiently makes use of the available resources [1].

Each routing protocol has its own variable known as "metric" that is assigned to each route available to reach a destination [2]. Based on metric value, a router ranks them from most preferred to least preferred, which is used to select the best path when there are multiple paths to the same destination. A metric can be as simple as the number of hops between the source and destination or can have a complex dependency on several parameters like bandwidth, delay, etc. The major parameters that aid in the decision process of selecting the most suitable protocol [3, 4] are, i) Speed of Convergence, ii) Scalability, ii) Resource utilization and iv) Ease of management

2 Literature Survey

2.1 Overview of Dynamic Routing Protocols

2.1.1 Routing Information Protocol (RIP)

The Routing Information Protocol (RIP) is the oldest of the existing distance vector IP routing protocols that make use of UDP port 520 for encapsulating its messages [5]. RIP defines two types of messages, i) Request message - Informs its neighboring routers to send an update and ii) Response message – carries the actual routing table information. By configuring RIP on a router, the Request messages are sent out via all the RIP enabled interfaces as broadcast packets, following which, the requesting router listens to the response messages from its neighboring routers. Upon receiving a request message, the neighboring routers respond by sharing their routing table in the form of Response message. This process enters a loop until all the routers in the network are converged.

RIP being a Distance vector protocol; a RIP enabled router sends out an update containing the entire routing table information every 30 seconds. Whenever a new entry is included in the update, it is stored in the receiving router's routing table along with the IP address of the advertising router.

Hop count is the metric used by RIP to rate the priority of different routes obtained from its neighboring routers. The hop count is the actual number of routers traversed in a route from source to destination. Among the possible routes available in the RIP database for reaching the destination network, the route with the lowest hop count value is selected as the best path and entered into the routing table. The maximum value of the hop count is set to 15 in order to prevent routing loops, thereby indirectly limiting the size of the network. This small range of hop count makes RIP undesirable for large enterprise networks. The propagation of the inappropriate routing information throughout the network is prohibited by the usage of features like route poisoning, split horizon rule and hold down timer. In addition to that, RIP can load balance across a maximum of 6 equal-cost links [6].

RIP currently exist in 2 versions [7]: a) RIPv1 – being a classful routing protocol [8] where routing updates are sent without the subnet masks, it does not support variable-length subnet masking (VLSM). As a result, this can cause severe problems with discontinuous networks. Moreover, there is no authentication mechanism for the updates that are sent as broadcast packets. b) RIPv2 - this protocol supports classless Inter-Domain routing (CIDR) where updates are sent as multicast packets to 224.0.0.9. It also uses/provides MD5 authentications for secure exchange of routing information. Despite being a classless protocol, RIPv2 by default does automatic classful summarization when advertisements traverse through major network boundaries.

2.1.2 EIGRP

Enhanced Interior Gateway Routing Protocol is a Cisco proprietary protocol that uniquely blends in the best features of both distance vector and link-state protocols. This hybrid routing protocol employs an extremely powerful and efficient algorithm known as Diffused Update Algorithm (DUAL) for the selection of best paths to a destination. This algorithm has certainly made EIGRP one of the best routing protocols that offers a rapid network convergence coupled with reliability by guaranteeing a loop-free operation [9]. EIGRP also has a reputation of being very easy on the router's memory and CPU utilization during normal operation.

Scalability is one of the major advantages of EIGRP as it has the potential to accommodate very large enterprise networks. Apart from that, EIGRP features a fairly simple configuration when compared to other existing routing protocols. Even

though EIGRP is classless in nature like RIPv2, it does automatic summarization by default where prefixes between major network boundaries are summarized to classful boundaries. However, this feature can be disabled based on the requirement.

In addition to that, EIGRP extends its ability to support routing information exchange for multiple network protocols like IPv4, IPv6, IPX and AppleTalk. This is made possible with the assistance of Protocol Dependent Module (PDM) concept, where separate routing table is used for each protocol in the network layer.

EIGRP makes use of 5 possible parameters for metric calculation, i) Bandwidth - inverse of the lowest bandwidth (in kbps) along the path is selected, which is then scaled by a factor of [10^7 *256], ii) Delay - cumulative delay (in tens of microseconds) along each route, which is then scaled by a factor of 256, iii) Reliability - Lowest reliability along the path is selected, whose values are ranging from 0 to 255 with 255 being the preferred value, iv) Load - highest load value along the path is selected for computation, which ranges from 0% to 100% with 0% load being the preferred value and v) MTU - smallest Maximum Transmission Unit along each path is tracked by EIGRP, although this parameter is not used for composite metric calculation. It is just considered to be potential tie-breaker.

These parameters can be adjusted accordingly to fine tune the link characteristics for optimal route selection [10]. By default, EIGRP uses only the bandwidth and delay parameters for selecting best paths. In simple terms, Bandwidth is nothing but a measure of the data carrying capacity of a link and delay is the propagation time taken by the data packets to reach its destination.

It is highly recommended that parameters like reliability and load are not modified as they can easily cause routing loops in the network, resulting in high CPU utilization. Interestingly, these parameters have corresponding K values using which we can manipulate the composite metric value. These K values must match between any two EIGRP routers for them to become neighbours. Composite metric is computed as:

$$metric = [k1 * bandwidth + \frac{k2*bandwidth}{256-load} + k3 * delay]$$

$$(1)$$

$$If(k5! = 0)$$
$$metric = metric * [k5/(reliability + k4)]$$
$$(2)$$

EIGRP Terminologies – DUAL are, i) Advertised Distance (AD) - Cost from the next hop

to the destination network, which is sometimes referred to as Reported Distance (RD), ii) Local Distance (LD) - It is a composite metric to reach the local neighbor that advertised the path, iii) Feasible Distance (FD) - Cost between the local router and the final destination prefix. In other words, it is the sum of Advertised Distance (AD) and Local Distance (LD), making it the final composite metric of the best path, vi) Successor - Best path used by EIGRP to reach the destination and v) Feasible Successor – It is a backup path to the destination.

One of the major advantages in using DUAL is that it offers sub-second network convergence by immediately replacing with the feasible successor in the routing table when the successor fails. However, a feasible successor is entered into the topology table only when it satisfies the Feasibility Condition- Next hop's AD should be less than FD of the current successor.

Five types of EIGRP packets are used for communication, i) Hello, ii) Update, iii) Acknowledgment, iv) Query and v) Reply. EIGRP Process creates and maintains three tables, i) Neighbour table - Lists all the directly connected routers that are running EIGRP with which the local router has established a neighborship, ii) Topology table - Lists all possible routes learned from each of the local router's EIGRP neighbors and this includes both Successor and Feasible Successor for a route that satisfied the Feasibility Condition and iii) Routing table- Lists all the best routes with the lowest composite metric selected from the EIGRP topology table.

2.1.3 OSPF

Open Shortest Path First (OSPF) is an open standard protocol that was developed by the Internet Engineering Task Force (IETF) as a replacement for RIP. This is one of the most commonly used link-state protocols that makes use of a highly efficient but processor intensive Dijkstra's Shortest Path First (SPF) algorithm.

One of the major advantages of using OSPF is that it provides scalability to much larger production networks [11]. It is achieved with the concept of areas, where the router and its links are associated with a particular area. This enhances the possibility of constructing hierarchical network topologies for easier administration and drastically reduces the routing protocol's impact on memory and CPU.

While breaking down the physical topology into different areas, a lot of factors need to be considered:

- Number of routers that can exist in a particular area [12].

- Number of OSPF neighbours that can be supported by a router, which is dependent on the local router's software and hardware platform.
- Number of areas that can be handled by the router.

However, stability and redundancy are the most significant factors that should be taken into account while assigning areas to a router's interface. An optimal selection on the size of the area needs to be done in such a way that it enhances the stability of the network without affecting the router's CPU resources. Because, even for a small change in the state of the link used for a particular route, all the routers in that area should run SPF algorithm for recalculation of its routes.

In addition to that, OSPF being a classless dynamic routing protocol supports super netting and VLSM for efficient management of IP addresses. OSPF exhibits less susceptibility to unwanted routing information by making use of secure MD5 authentication mechanisms that hashes the data packets exchanged between two OSPF enabled routers. In order to maximize the utilization of multiple paths available for routing packets, this protocol supports equal-cost load balancing too.

In OSPF, routers are classified as Area Border Routers (ABR), Autonomous System Boundary Routers (ASBR) and Internal routers. ABR is a router that has at least one of its interfaces belonging to the Backbone area (Area 0), thereby connecting multiple areas to the main backbone area [13]. ASBR is the type of router that runs more than one routing protocol including OSPF and performs a mutual exchange of OSPF routing information with routers that are running other protocols. In OSPF, all the routers belonging to a particular area should have identical OSPF Database [14]. Even though this protocol supports manual summarization, it can be performed only at ABRs and ASBRs to satisfy the above condition.

Each OSPF enabled router exchanges the routing information with its adjacencies in the form of Link State Advertisements (LSA). There are different types of LSAs based on the routing information and the type of router that originated it. The most commonly known LSAs are Router LSA (LSA 1), Network LSA (LSA 2), Network Summary LSA (LSA3), ASBR Summary LSA (LSA 4), AS External LSA (LSA 5) and NSSA External LSA (LSA 7) [15].

OSPF has classified the areas into five types for better control over data flow inside the network as, i) Backbone (area 0) - Permits Router LSA, Network LSA, Network Summary LSA, ASBR Summary LSA and AS External LSA, ii) Non-

backbone, non-stub - Permit Allows Router LSA, Network LSA, Network Summary LSA, ASBR Summary LSA and AS External LSA, iii) Stub - Permits Router LSA, Network LSA, Network Summary LSA and blocks others types of LSA, iv) Totally Stub - Permits Router LSA and Network LSA and blocks other types of LSA and v) Not-so-stubby - Permits Router LSA, Network LSA, Network Summary LSA, ASBR Summary LSA and NSSA External LSA.

Once OSPF databases are exchanged and synchronized between two adjacent routers, best path selection process begins with the help of SPF algorithm [16]. OSPF uses Cost as a metric for path selection, which is computed based on the bandwidth of the link being used [17]. Cost is inversely proportional to the bandwidth value:

$$Default\ Cost = 100\ Mbps/Link\ Bandwidth$$

(3)

where, Bandwidth is in Mbps.

In other words, higher the bandwidth value, lower is the cost assigned to that link. For instance, an Ethernet link with a bandwidth of 10Mbps would be assigned a cost of 10. The cost of a path/route is the final sum of the individual costs of all the outgoing interfaces involved in reaching the destination. Finally, SPF selects the link with the lowest end-to-end cost as the best path and enters that into the routing table [18].

3   Real-Time Comparison

Fig. 1 shows the graphical representation of the experimental setup. Platforms used are i) 5 Cisco Routers (comprising of Cisco 2801 Integrated Services Router, IOS 12.5T, RAM 256 MB, Flash 64 MB, Cisco 2811 Integrated Services Router, IOS 12.5T, RAM 256 MB, Flash 128 MB), ii) Dell XPS Laptop and iii) Ethernet Switch.

a. Experimental setup
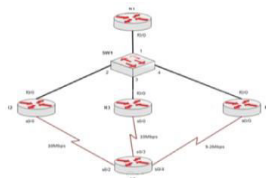
b. Graphical Representation

Fig. 1 Graphical representation of Experimental setup

We have emulated a network topology as shown in Fig.1 with routers R1 and R5 being the source and destination of the data traffic respectively. There are three transit routers on the path namely R2, R3 and R4 and hence the data packets have the option to make use of all or some of the available paths based on the routing protocol employed. Among those three links to router R5, two of them carry the same bandwidth of 10Mbps, whereas the other link between R4 and R5 has a bandwidth of 9.2Mbps. By deploying various dynamic routing protocols like RIP, EIGRP and OSPF in the given network, and generating data traffic from loopback interface of R1 destined to loopback interface of R5, we have made a detailed analysis on the utilisation of available links by these protocols by their support to unequal-cost load balancing feature and efficient usage of available resources.

By generating a continuous ping of 1000 Bytes Echo packets with a repeat count of 500, we tested the unequal cost load balancing capability of the protocols deployed.

3.1 Results for EIGRP

EIGRP is deployed among the routers; automatic summary is disabled and all the prefixes are advertised into the network. After convergence of routing and topology tables, if the network is stable, we can see the following topology table output on router R1 as shown in Fig. 2.

Fig. 2 EIGRP topology table output on the router R1

It can be seen that EIGRP was able to recognize all the three available paths to the prefix 5.5.5.5/32 as all of them satisfied the Feasibility Condition. But if we notice the routing table of R1, only two paths with equal cost are chosen as best paths and inserted as routing table entries as shown in Fig. 3.

Fig. 3 EIGRP best paths in the routing table

By performing a trace route operation on R1 for the destination 5.5.5.5 with its loopback interface as a source, we can observe the following output as shown in Fig. 4.

```
R1#traceroute 5.5.5.5 source 1.1.1.1

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 192.168.1.3 16 msec
    192.168.1.2 56 msec
    192.168.1.3 68 msec

  2 192.168.25.5 20 msec
    192.168.35.5 24 msec *
```

Fig. 4 EIGRP Trace route operation on R1

As we can see, the traffic is load balanced across R2 and R3, but not R4. Now by using the feature "variance" available with EIGRP, unequal-cost load balancing can be achieved. So by setting variance 2 under the EIGRP process, it is possible to obtain the following output as Fig. 5.

```
R1#show ip route eigrp

D    192.168.45.0/24 [90/815616] via 192.168.1.4, FastEthernet0/0
D    192.168.25.0/24 [90/793600] via 192.168.1.2, FastEthernet0/0
     5.0.0.0/32 is subnetted, 1 subnets
D       5.5.5.5 [90/943616] via 192.168.1.4, FastEthernet0/0
                [90/921600] via 192.168.1.3, FastEthernet0/0
                [90/921600] via 192.168.1.2, FastEthernet0/0
```

Fig. 5 output of variance 2 under the EIGRP process

It can be seen that EIGRP was able to recognize all the three available paths to the prefix 5.5.5.5/32 from its topology table and put those entries in the routing table as all of them satisfied the Feasibility Condition. So if we notice the routing table of R1, two paths with equal cost and the one with unequal cost are chosen as best paths and inserted as routing table entries.

By performing a trace route operation on R1 for the destination 5.5.5.5, we can observe the following output as shown in Fig. 6.

```
R1#traceroute 5.5.5.5 source 1.1.1.1

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 192.168.1.4 16 msec
    192.168.1.3 72 msec
    192.168.1.2 32 msec
  2 192.168.45.5 64 msec
    192.168.35.5 28 msec
    192.168.25.5 28 msec
```

Fig. 6 EIGRP trace route operation on R1 for the destination 5.5.5.5

As we can see, the traffic is load balanced across R2, R3 and R4. Generating ICMP data traffic from R1 with size 1000 bytes each and with a packet count of 500 to R5, we can see that all the traffic is being sent successfully as shown in Fig. 7.

```
R1#ping 5.5.5.5 source 1.1.1.1 size 1000 repeat 500

Type escape sequence to abort.
Sending 500, 1000-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!
Success rate is 100 percent (500/500), round-trip min/avg/max = 4/19/56 ms
```

Fig. 7 ICMP data traffic from R1

Now, the incoming interface (Fig. 8) statistics of router R5 are checked to see in order to verify that unequal-cost load balancing has been supported by EIGRP. It can be seen that ICMP data traffic generated from R1 to R5 was received from all the available incoming interfaces on R5.

```
R5#show interfaces serial 0/2 | include packets input
    193 packets input, 183580 bytes, 0 no buffer
R5#show interfaces serial 0/3 | include packets input
    171 packets input, 161492 bytes, 0 no buffer
R5#show interfaces serial 0/4 | include packets input
    169 packets input, 159484 bytes, 0 no buffer
```

Fig. 8 Incoming interface

3.2 Results for RIPv2

RIPv2 is deployed among the routers; automatic summary is disabled and all the prefixes are advertised into the network. After convergence of routing tables and if the network is stable, we can see the following output on router R1 Fig. 9.

```
R1#show ip rip database
1.0.0.0/8    auto-summary
1.1.1.1/32    directly connected, Loopback0
5.0.0.0/8    auto-summary
5.5.5.5/32
     [2] via 192.168.1.2, 00:00:06, FastEthernet0/0
     [2] via 192.168.1.3, 00:00:24, FastEthernet0/0
<output omitted>
```

Fig. 9 RIPv2 output on the router R1

It can be seen that RIP was able to recognize only two paths to the prefix 5.5.5.5/32 and entered those paths alone into its database. Because of the absence of third path in RIP database of R1, only two paths with equal cost are chosen as best paths and inserted as routing table entries and that can be seen from the following routing table as shown in Fig. 10.

```
R1#show ip route rip
R    192.168.45.0/24 [120/1] via 192.168.1.4, FastEthernet0/0
R    192.168.25.0/24 [120/1] via 192.168.1.2, FastEthernet0/0
     5.0.0.0/32 is subnetted, 1 subnets
R       5.5.5.5 [120/2] via 192.168.1.3, FastEthernet0/0
                [120/2] via 192.168.1.2, FastEthernet0/0
R    192.168.35.0/24 [120/1] via 192.168.1.3, FastEthernet0/0
```

Fig. 10 RIPv2 best paths in the routing table for R1

By performing a trace route operation on R1 for the destination 5.5.5.5 with its loopback interface as a source, we can observe the following output as in Fig. 11.

```
R1#traceroute 5.5.5.5 source 1.1.1.1

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 192.168.1.3 36 msec
    192.168.1.2 24 msec
    192.168.1.3 28 msec
  2 192.168.25.5 28 msec
    192.168.35.5 48 msec *
```

Fig. 11 Trace route operation on R1 for the
destination 5.5.5.5 with the
loopback interface as a source

As we can see, the traffic is load balanced across R2 and R3, but not on R4, because that path does not even show up in the RIP database of R1. Generating ICMP data traffic from R1 with size 1000 bytes each and with a packet count of 500 to R5, we can see that all the traffic is being sent successfully as in Fig. 12.

```
R1#ping 5.5.5.5 source 1.1.1.1 size 1000 repeat 500

Type escape sequence to abort.
Sending 500, 1000-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!
Success rate is 100 percent (500/500), round-trip min/avg/max = 4/19/40 ms
```

Fig. 12 ICMP data traffic from R1

Now, the incoming interfaces' statistics of router R5 are checked to see in order to verify that unequal-cost load balancing has not been supported by RIP.

```
R5#show interfaces serial 0/2 | include packets input
    254 packets input, 251456 bytes, 0 no buffer
R5#show interfaces serial 0/3 | include packets input
    253 packets input, 251380 bytes, 0 no buffer
R5#show interfaces serial 0/4 | include packets input
    4 packets input, 456 bytes, 0 no buffer
```

Fig. 13 Incoming Interface

It can be seen that ICMP data traffic generated from R1 to R5 was received on only equal cost interfaces on R5.

3.3. Results for OSPF

OSPF is deployed among the routers; automatic summary is disabled and all the prefixes are advertised into the network. After convergence of routing tables and if the network is stable, we can see the following output on router R1 as shown in Fig. 14.

```
R1#show ip ospf database router 5.5.5.5

          OSPF Router with ID (1.1.1.1) (Process ID 1)

             Router Link States (Area 0)

  LS Type: Router Links
  Link State ID: 5.5.5.5
  Advertising Router: 5.5.5.5
  Number of Links: 7

   Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 4.4.4.4
    (Link Data) Router Interface address: 192.168.45.5
    Number of TOS metrics: 0
      TOS 0 Metrics: 64

   Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 3.3.3.3
    (Link Data) Router Interface address: 192.168.35.5
    Number of TOS metrics: 0
      TOS 0 Metrics: 64

   Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 2.2.2.2
    (Link Data) Router Interface address: 192.168.25.5
    Number of TOS metrics: 0
      TOS 0 Metrics: 64

   Link connected to: a Stub Network
<output omitted>
```

Fig. 14 OSPF output on the router R1

It can be seen that OSPF was able to recognize all the three available paths to the prefix 5.5.5.5/32 and all of them are entered into the OSPF link state database of R1. But if we notice the routing table of R1, only two paths with equal cost are chosen as best paths and inserted as routing table entries as shown in Fig. 15.

```
R1#show ip route

<output omitted>

Gateway of last resort is not set

     1.0.0.0/32 is subnetted, 1 subnets
C       1.1.1.1 is directly connected, Loopback0
O    192.168.45.0/24 [110/30] via 192.168.1.4, FastEthernet0/0
O    192.168.25.0/24 [110/20] via 192.168.1.2, FastEthernet0/0

     5.0.0.0/32 is subnetted, 1 subnets
O       5.5.5.5 [110/21] via 192.168.1.3, FastEthernet0/0
                [110/21] via 192.168.1.2, FastEthernet0/0

O    192.168.35.0/24 [110/20] via 192.168.1.3, FastEthernet0/0
C    192.168.1.0/24 is directly connected, FastEthernet0/0
```

Fig. 15 OSPF best paths in the routing table for R1

By performing a trace route operation on R1 for the destination 5.5.5.5 with its loopback interface as a source, we can observe the following output as in Fig. 16.

```
R1#traceroute 5.5.5.5 souce 1.1.1.1

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 192.168.1.3  56 msec
    192.168.1.2  60 msec
    192.168.1.3  24 msec

  2 192.168.25.5  44 msec
    192.168.35.5  76 msec *
```

Fig. 16 OSPF trace route operation on R1 for the
destination 5.5.5.5 with its
loopback interface as a source

As we can see, the traffic is load balanced across R2 and R3, but not on R4, because that path does not even show up in the RIP database of R1. Generating ICMP data traffic from R1 with size 1000 bytes each and with a packet count of 500 to R5, we can see that all the traffic is being sent successfully as in Fig 17.

```
R1#ping 5.5.5.5 source 1.1.1.1 size 1000 repeat 500

Type escape sequence to abort.
Sending 500, 1000-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
||||||||||||||||||||||||||||||||||||||||||||||||||
||||||||||||||||||||||||||||||||||||||||||||||||||
||||||||||||||||||||||||||||||||||||||||||||||||||
||||||||||||||||||||||||||||||||||||||||||||||||||
||||||||||||||||||||||||||||||||||||||||||||||||||
||||||||||||||||||||||||||||||||||||||||||||||||||
||||||||||||||||||||||||||||||||||||||||||||||||||
||||||||||
Success rate is 100 percent (500/500), round-trip min/avg/max = 8/20/60 ms
```

Fig. 17 OSPF ICMP data traffic from R1

Now, the incoming interfaces (Fig 18) statistics of router R5 are checked to see in order to verify that unequal-cost load balancing has not been supported by OSPF.

```
R5#show interfaces serial 0/2 | include packets input
    253 packets input, 251440 bytes, 0 no buffer
R5#show interfaces serial 0/3 | include packets input
    253 packets input, 251440 bytes, 0 no buffer
R5#show interfaces serial 0/4 | include packets input
    5 packets input, 548 bytes, 0 no buffer
```

Fig. 18 Incoming Interface

It can be seen that ICMP data traffic generated from R1 to R5 was received on only equal cost interfaces on R5.

We have configured the networks on each router, for each routing protocol and have determined the bandwidth usage of various routing protocols through graphs at fa0/1 interface (100Mbps) of router during network convergence. We have also estimated the bandwidth usages of the selected protocols for Ethernet (100Mbps), T1 (2.544Mbps) and Serial (512Kbps) links and compared them. The following graphs (Fig 19) of 'Time (X-axis) vs.

Bandwidth (Y-axis)' are plotted for the selected protocols.
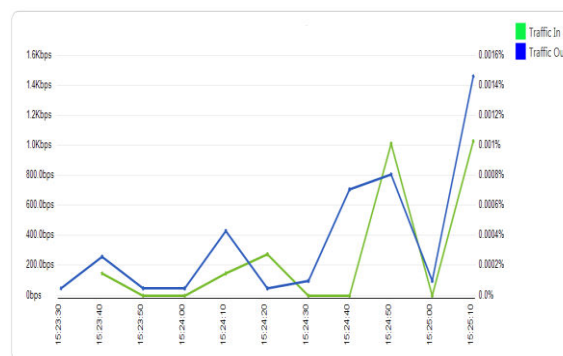
a. EIGRP protocol results
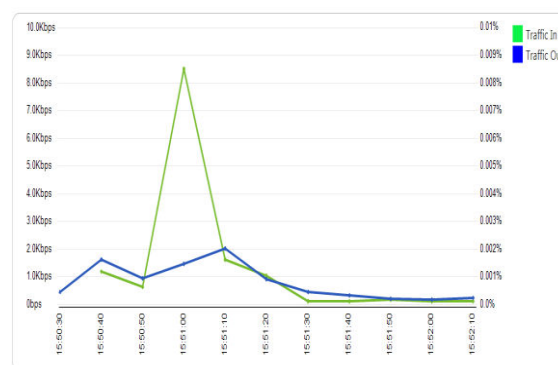
b. RIPv2 protocol results
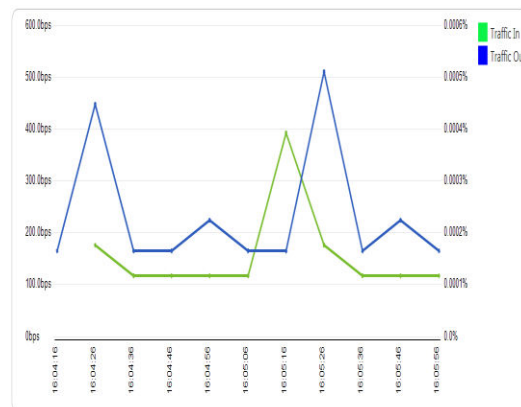
c. OSPF protocol results

Fig. 19 Bandwidth Usage

Table 1 shows the comparison of bandwidths used during convergence.

Table: 1 Comparison of Bandwidths Used During
Convergence

Protocol	Fast Ethernet	Ethernet	T1	Serial
EIGRP	0.0016	0.016	0.1056	2.475
RIP	0.002	0.02	0.129	3.022
OSPF	0.0004	0.004	0.0264	0.619

From the above table, we can infer the following points:

1. If the network involves more Serial or T1 links, it better to avoid RIP and choose OSPF over EIGRP and RIP.
2. If the network is not a stable one, that is, it has got frequent convergences happening, OSPF has to be chosen over the other two protocols as it consumes the least amount of bandwidth.

4   Analysis of Routing Protocols

A lot of factors play an important role in deducing the correct protocol for a network. Bandwidth utilization is one of the most important factors to be considered. Through unequal-cost load balancing, all the available links can be utilized, thereby causing better utilization of the available resources. But as we can see from Case I, this feature is supported by EIGRP alone, but not by OSPF or RIP though all the three protocols do support equal-cost load balancing. Hence, if there are many unequal cost links in the network, it is suggested to go for EIGRP as it alone uses the given resources in an optimal manner. Also if the network cannot allocate much bandwidth to control plane traffic, it better to avoid RIP and choose OSPF over EIGRP and RIP and specifically if the network involves low bandwidth links, this solution would prove optimal as we can see from case II. And for a network involving frequent convergences, then OSPF would be optimal taking bandwidth utilised during convergence as the parameter.

When it comes to the order of complexity of configuring protocols, RIP has a lesser complexity order compared to EIGRP and OSPF [19]. Since EIGRP has the flexibility to configure multiple AS on a single router, it can support very large networks [20]. OSPF, which is dependent on the concept of areas, supports large networks too. Another important factor is the ease with which the network can be scaled. Scaling limits the entries in the routing table. When the three protocols are compared, EIGRP performs best by using Autonomous System (AS), followed by OSPF, which uses areas. RIP does not scale easily.

Considering CPU resources, OSPF performs poorly as it uses a complex SPF algorithm where as EIGRP and RIP consume less. When networks run IPX or AppleTalk, the go-to protocol is EIGRP, as it is the only protocol that can route these packets. RIP easily beats the other two, as all routing software and devices support it. For example, when networks contain old Unix routers, RIP is the best choice.

5   Conclusion

The same protocol may not be the best for all kinds of networks. It all depends on the network environment and business needs. Considering both unequal-cost load balancing and bandwidth utilisation during convergence in real time as key parameters, comparative analysis of routing protocols based on the optimal resource utilization has been done and based up on the results, choosing the correct protocol for a network with bandwidth as a bottle-neck factor or with unequal-cost links, becomes easier.

References

[1] GoyalM., Soperi. M, Baccelli, E.,Choudhury, G.,Shaikh, Hosseini, Trivedi K., "Improving Convergence Speed and Scalability in OSPF: A Survey", 14(2), pp:443 – 463, *Communications Surveys & Tutorials*, IEEE, 2012.

[2] Jeff Doyle, Jennifer Carroll, "CCIE Professional Development Routing TCP/IP", Volume1, Second Edition, *Cisco Press*.

[3] D Venkatavara Prasad, Kalyan G.P.S, "Optimal selection of Dynamic Routing Protocol with real time case studies", Proceedings in *International Conference on Recent Advances in Computing and Software Systems (RACSS)*, 2012, vol., no., pp.219,223, 25-27 April 2012, doi: 10.1109/RACSS.2012.6212727

[4] B. R. Bellur and R. G. Ogier, "A reliable, efficient topology broadcast protocol for dynamic networks" , In Proceedings of *IEEE INFOCOM*, 1999, pp.178–186.

[5] K. Mishra, "Traffic Engineering Solutions for OSPF Based Best Effort Networks", *Master's Thesis*, K. R. School of Information Technology, IIT Bombay, 2007.

[6] Jun Wang, Yaling Yang, Li Xiao, KlaraNahrstedt, "Edge-based traffic engineering for OSPF networks" *Computer Networks*, 48(4), 15 July 2005, Pages 605-625.

[7] Cisco Systems Inc, "The Internetworking Technologies Handbook - Conflicts, Empire and National Identity", Fourth Edition, *Cisco Press*.

[8] Russ White, Don Slice, Alvaro Retana, "Optimal Routing Design", *Cisco Press*, June 7, 2005.

[9] Alaettinoglu, Cengiz, Van Jacobson, and Haobo Yu. "Towards milli-second IGP convergence" *IETF Draft* (2000)

[10] N. Wang, K. Ho, G. Pavlou, and M. Howarth, "An overview of routing optimization for internet traffic engineering", *IEEE*

Communication Surveys & Tutorials, vol. 10, no. 1, 2008.

[11] Poprzen, N.; Gospic, N, "Scaling and convergence speed of EIGRPv4 and OSPFv2 dynamic routing protocols in hub and spoke network", *IEEE International Conference on Telecommunication in Modern Satellite, Cable, and Broadcasting Services*, 2009.

[12] M. Goyal, K. Ramakrishnan, and W. Feng, "Achieving faster failure detection in OSPF networks", In Proceedings of *IEEE International Conference on Communications (ICC 2003)*, pp. 296–300

[13] Zinin, Lindem, D. Yeung, "Alternative Implementations of OSPF Area Border Routers", *Request for Comments: 3509*, April 2003.

[14] J. McQuillan, "The birth of link-state routing," *IEEE Annals of the History of Computing*, 31(1), January/March 2009.

[15] M. Takashi, K. Takashi, and A. Michihiro, "Enhancing the network scalability of link-state routing protocols by reducing their flooding overhead", In Proc. *IEEE Workshop on High Performance Switching and Routing (HPSR)*, June 2003, pp. 263–268.

[16] NemanjaPoprzen, NatasaGospic, "Performance Comparison of EIGRPv4 and OSPFv2 Dynamic Routing Protocols in Network With One Redundant Path", *In Proceedings of INFOTECH 2009*, Pale, Bosnia and Herzegovina, 2009.

[17] M. Goyal, W. Xie, S. H. Hosseini, K. Vairavan, and D. Rohm, "Improving OSPF dynamics on a broadcast LAN", *Simulation,* 82(2), pp. 107–129, 2006

[18] S. Venkatesh, "Smart adjacency establishment in OSPF routing protocol," *Master's thesis, University of Wisconsin - Milwaukee*, 2006

[19] Thorenoor, S.G., "Dynamic Routing Protocol Implementation Decision between EIGRP, OSPF and RIP Based on Technical Background Using OPNET Modeler", *In Proceedings of Second International Conference on Computer and Network Technology (ICCNT)*, pp.191-195, 2010, doi: 10.1109/ICCNT.2010.66

[20] Wijaya, C, "Performance Analysis of Dynamic Routing Protocol EIGRP and OSPF in IPv4 and IPv6 Network", International Conference on Informatics and Computational Intelligence, 2011, pp:355-360, doi:10.1109/ICI.2011.64.