Evaluation of IS-IS Implemented in Quagga

Saša Takov and Aleksandra Smiljanić

***Abstract*—Routing protocols define rules according to which routers exchange the network topology information. Based on this topology information and a packet address, routers decide where to send incoming packets. Intermediate System to Intermediate System (IS-IS) routing protocol is the most scalable intradomain routing protocol. In this paper, we will evaluate the IS-IS functionalities implemented in the Quagga open-source software.**

***Index Terms*—Internet router, IS-IS protocol, Quagga, isisd**

I. INTRODUCTION

Internet communication is realized by exchanging corresponding packets between entities. Routers have the most important function on a packet path. In order to correctly forward each of the incoming packets a router has its lookup table. Nexthop information for all reachable networks are stored in the lookup table. Based on this information and a packet address, a router sends packets to appropriate ports.

Routers create their lookup tables according to the received network topology information. Exchanging of this information is defined strictly by the rules of the routing protocols. Although many routing protocols are developed, those that are in widespread use in Internet Protocol (IP) networks can be classified into two major types, interior (intradomain) and exterior (interdomain) gateway protocols. Interior gateway protocols are responsible for disseminating information between routers in an Autonomous System (AS), or domain, while exterior gateway protocols are used for communication between routers from different ASs. An AS is a connected group of one or more IP prefixes run by one or more network operators which has a single and clearly defined routing policy [1].

A routing protocol also can be considered as distance vector or link state protocol. A router that use some distance vector protocol advertises its routing table to all directly connected neighbors periodically. Bandwidth is not used frugally and such protocol has slow convergence. On the other hand, in link state protocols, only routing updates are propagated only when they occur, which results in more effective bandwidth utilization and faster convergence.

Pool of the IPv4 addresses are becoming insufficient resource because of constant increase of the number of Internet users. As a consequence, routing protocols need to support both IPv4 and IPv6 addresses. Routing protocols generally support either IPv4 addresses or IPv6 addresses, which hinders transition from IPv4 to IPv6 addresses.

Intermediate System to Intermediate System (IS-IS) is a link state interior gateway protocol, that can support simultaneously IPv4 and IPv6 addresses. It was initially defined as an international standard within the Open Systems Interconnection (OSI) reference design. IS-IS later was extended in order to operate in IP, which is used as the network layer protocol on the Internet. Except in OSI and TCP/IP model, IS-IS can be used in Multi Protocol Label Switching (MPLS) networks as well [4]. AS that runs IS-IS is divided into areas. Routers of one area do not have information about the routers in other areas and their interconnection. Packets bound for routers in different area are transmitted to the closest border router. Such isolation of the areas makes IS-IS to be the most scalable intradomain routing protocol.

In order to follow development of services and applications that require a lot of bandwidth, router capacity must constantly increase. Although open-source solutions in the routing area are still behind products of the large vendors regarding performance, in the most cases they are satisfactory solution. Also, as opposed to expensive vendors' routers, they allow implementation of the new ideas, which result in significant improvements. Quagga is widely used routing software in vast offer of the open-source solutions. It is applied in the networks of big companies such as Google. Quagga software suite provides implementations of the most common routing protocols, which are implemented as separate daemons [5]. Daemon implementing IS-IS protocol is denoted as isisd.

Even though the working group responsible for Quagga development have emphasized correct functionality of isisd daemon (in version 0.99.22), to best of our knowledge, there is no document in public domain which validates its functionality and describes the daemon configuration.

In this paper we will evaluate the functionalities of the isisd daemon provided in the Quagga v0.99.22.4. routing software. We created the test network environment using Linux Containers (LXC) virtualization software. In this network we were able to test all important functionalities of the isisd daemon as well as its behavior in realistic scenarios, such as initialization of new devices, or link failures. In tests we focused on the isisd functionalities in TCP/IP networks.

This paper is organized as follows. The second section gives necessary information for understanding IS-IS protocol and Quagga routing software. The third section describes used test network environment. Results of the performed tests are given in the fourth section. The paper is concluded in the last section.

II. IS-IS PROTOCOL AND QUAGGA SOFTWARE

In this section we will describe IS-IS protocol and Quagga routing software. In chapter *A.* we will explain terminology used in IS-IS and the most important features of

Saša Takov is with the Inovation Center, School of Electrical Engineering, University of Belgrade, 73 Bulevar kralja Aleksandra, 11020 Belgrade, Serbia (e-mail: sasa.takov@ ic.etf.rs).

Aleksandra Smiljanić is with the School of Electrical Engineering, University of Belgrade, 73 Bulevar kralja Aleksandra, 11020 Belgrade, Serbia (e-mail: aleksandra@ etf.rs).

this protocol, such as routers' types and Network Service Access Point (NSAP) addresses, which are used in OSI environment. Also, IS-IS will be compared with Open Shortest Path First (OSPF), which is the widely used intradomain protocol [6]. Chapter *B.* describes architecture and basic principles of the Quagga routing suite.

A. IS-IS protocol

In IS-IS terminology, a domain is equal to an Autonomous System in Border Gateway Protocol (BGP), Intermediate System (IS) is router, and End System (ES) is a term for node/host. IS-IS domain can be partitioned into areas. Borders between areas in IS-IS are on links, which connect two routers. As a consequence, a router belongs to a single area. This is opposite in OSPF, where a router can be part of two or more areas, because borders between areas in OSPF are on the Area Border Routers (ABRs). An example of the IS-IS network is shown in Fig.1.

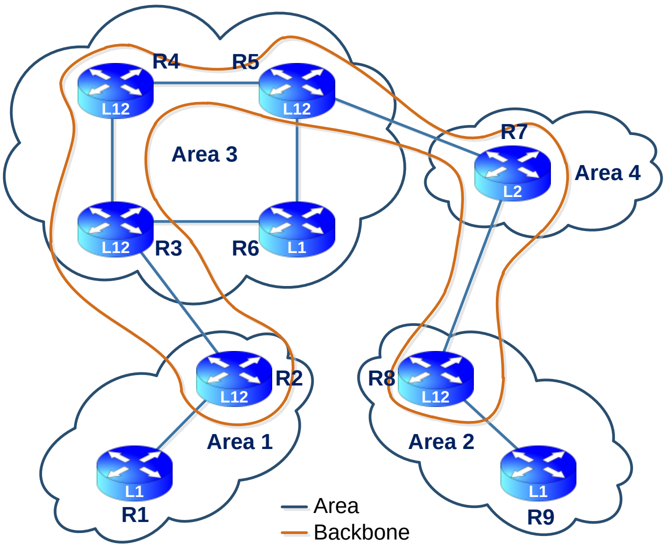

Fig. 1. An example of the IS-IS network

Packets exchanged in the IS-IS protocol are known as Protocol Data Units (PDUs). There are four types of PDUs. Intermediate System-to-Intermediate System Hello (IIH) are used for establishing of adjacencies between routers. A Link-state PDU (LSP) collects the information about each router and its interfaces in IS-IS network and based on these PDUs, routers form their Link State Databases (LSDBs). Each router creates its own routing table using Dijkstra algorithm based on its LSDB. LSPs are like Link-State Advertisements (LSAs) in OSPF. Complete Sequence Number PDU (CSNP) contains LSP ID, LSP lifetime, sequence number and checksum field of each entry in the LSDB. Partial Sequence Number PDU (PSNP) is used when routers request LSP information from their neighbors and for confirmation that requested LSP is received.

IS-IS provides two-level hierarchical routing, intra-area and inter-area. Intra-area routing is routing across single IS-IS area and routers responsible for it are denoted as Level-1 (L1). Level-2 (L2) routers are responsible for inter-area routing between different areas. Also, there are Level-1/Level-2 (L1L2) routers, which forward packets to the destinations in its own area and can communicate with routers in other areas as well. A L1 router only knows the topology of its area and as neighbors has other L1 or L1L2 routers from this area. It possesses the Level-1 Link-State Database (L1 LSDB) with the complete routing information for the area. In order to communicate with entities in other areas, L1 router sends packets to the closest L1L2 router in its area, which then forwards the packets to the L2 (or L1L2) router in the destination area. When the destination L2 (or L1L2) router receives packets, it routes them toward the destination according to its routing table. A L2 router has L2 LSDB for routing between areas and may establish adjacencies with L2 (or L1L2) routers from all areas. It does not know topology within any L1 area, but it should be informed about all reachable networks by the other routers from level-2. Unlike L1 and L2 routers, a L1L2 router has two separate LSDBs (L1 LSDB and L2 LSDB) and forms adjacencies with L1 (or L1L2) routers in its area and L2 (or L1L2) routers in other areas.

In OSPF all areas are connected to the backbone, which is a particular area through which inter-area traffic must pass. On the other hand, the IS-IS backbone is the interconnected collection of the L2 and L1L2 routers, as shown in Fig. 1. This feature allows much easier network extension in IS-IS than in OSPF. On the other hand, should the level-2 backbone become partitioned, there is no provision for use of level-1 links to repair a level-2 partition [7]. As a consequence, when a link between routers R4 and R3 in Fig.1 fails, router R5 cannot communicate with Area 1, even though it has path to router R2 via router R6. IS-IS is mostly used in networks of the largest Internet Service Providers (ISPs), while OSPF is still the most popular protocol in enterprise networks.

Because it was originally developed as an OSI routing protocol, IS-IS must use NSAP addresses even in case of pure IP environment. Each router in IS-IS is uniquely specified by special address, that is called Network Entity Title (NET). NET is a variant of NSAP address. It is hexadecimal value and has length from 8 to 20 bytes. Generally, each NET address consists of four parts, as shown in Fig. 2.

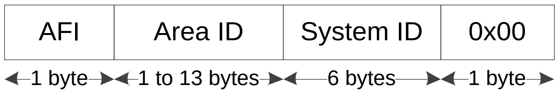

AFI	Area ID	System ID	0x00
1 byte	1 to 13 bytes	6 bytes	1 byte

Fig. 2. Format of a NET address

An Authority and Format Identifier (AFI) is 1-byte long and it is equivalent to the IP address space assigned to an AS. Area ID has variable length between 1 and 13 bytes and represents the area address of a router. System ID is 6-byte long value and it is unique label of a router or a node in the area. The last byte in a NET address is always 0x00. All routers in one area must have the same Area ID and different System IDs. Also, each L2 (or L1L2) router has unique System ID in a domain.

Although method for NET determination is not strictly defined, among ISPs there is a common rule. A typical NET address is 49.0001.1921.6800.1001.00. This address is 10 bytes long. The first byte is 0x49 and it is AFI for private addresses, such as the range 192.168.0.0/16 in IP networks. The next two bytes represent the Area ID and in the case of above example it is value 0x0001, which means that the router belongs to the Area 1. Similarly, Area IDs of Area 2, Area 3, etc. have values 0x0002, 0x0003, etc., respectively. System ID is determined according to router's loopback address by filling in all leading zeros, and then repositioning the decimal points to form three two-byte numbers. For example, if the loopback address of a router is 192.168.1.1, the result filled with zeros is 192.168.001.001. By rearranging the decimal points, we get 1921.6800.1002 as a System ID.

B. *Quagga routing software*

Quagga consists of several daemons corresponding to different routing protocols. Curently, Quagga has daemons implementing RIPv1, RIPv2, RIPng, OSPFv2, OSPFv3, BGPv4+ and IS-IS protocols. RIPv1, RIPv2 and OSPFv2 are IPv4 protocols, RIPng and OSPFv3 are IPv6 protocols, while BGPv4+ and IS-IS support both IPv4 and IPv6 addressing. Besides, there are daemons for wireless routing and MPLS, which have not validated yet, i.e. they have some known bugs. All these daemons are clients to the core daemon zebra, which acts as a routing manager.

Zebra is responsible for selecting the best route to each available destination from Routing Information Base (RIB), which represents collection of routes received from clients. Also, zebra communicates with kernel in order to update Forwarding Information Base (FIB) which are used by a router for transmitting packets. Each of these client daemons is independent from the other and can work standalone, but it must communicate with the zebra daemon. Consequently, before running a client daemon, zebra need to be initiated.

Each daemon (including zebra) has separate configuration file (denoted as *name_of_the_daemon*.conf). All these files must be placed in the same directory, which is set by the configuration option *sysconfdir* during the Quagga installation process. By changing daemons' configuration files, it is possible to create different network topologies, which implement active routing protocols.

III. TEST ENVIRONMENT

In order to evaluate the IS-IS implementation in Quagga routing suite, we created the appropriate test network environment. Network was created on PC under Ubuntu 12.10. operating system using the LXC tool. LXC is a lightweight virtualization technology. It allows creating multiple containers on one physical host, where each of them behaves as a standalone operating system. This means that a container has its own processes, system and network files, etc. LXC does not modify kernel of the host, which is its most important advantage compared to other virtualization softwares, such as VMware, Qemu and OpenVZ.

As a first step in test network generation, we created 11 containers using LXC software. By configuring virtual ethernet (veth) interfaces in containers and then by joining appropriate pairs of these interfaces to the same LXC bridge, we established direct connections between containers. Quagga software was installed on all containers. In this way, a container becomes a virtual router. We used the most recent version 0.99.22.4 of the Quagga routing suite. There are three methods of installation: using tar.gz file, yum repositorium, or the source code from GitHub. Our choice was installation by tar.gz file.

By generating configuration files of zebra and isisd daemon in each container, we created IS-IS network, shown in Fig. 3. Routers are denoted as R1-R11 and each of them belongs to one of three areas, according to their NET addresses, which are shown in Table I. The NET address is assigned to the router by *net* command in its isisd.conf file.

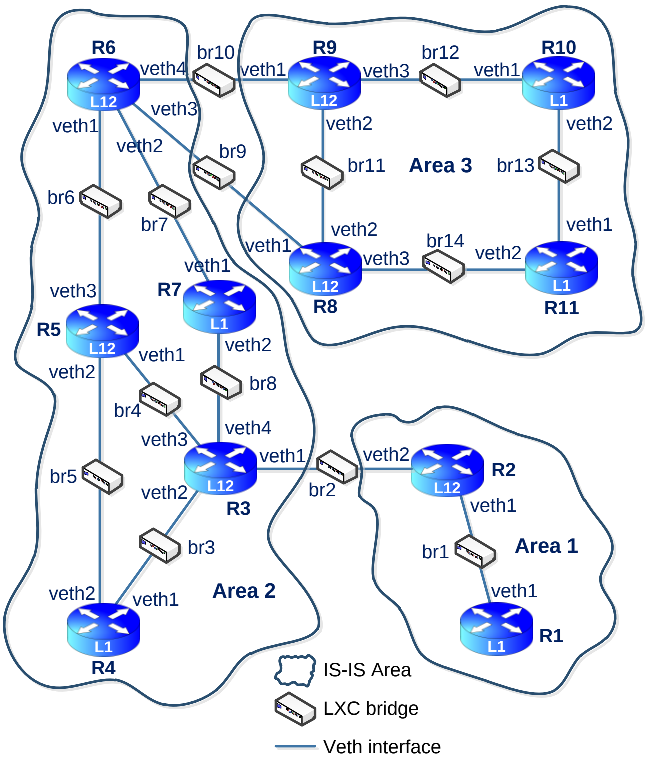

Fig. 3. Created test network environment for IS-IS evaluation

TABLE I
NET ADDRESSES OF THE ALL ROUTERS

Router	NET address
R1	49.0001.1921.6810.0001.00
R2	49.0001.1921.6810.0002.00
R3	49.0002.1921.6810.0003.00
R4	49.0002.1921.6810.0004.00
R5	49.0002.1921.6810.0005.00
R6	49.0002.1921.6810.0006.00
R7	49.0002.1921.6810.0007.00
R8	49.0003.1921.6810.0008.00
R9	49.0003.1921.6810.0009.00
R10	49.0003.1921.6810.0010.00
R11	49.0003.1921.6810.0011.00

Also, in Fig. 3., types of routers are L1 or L1/L2. A router gets one of these types by the *is-type* command in its isisd configuration file. All interfaces in Fig. 3. are denoted as vethX, where X is the ordinal number of the specific interface on a router. Each of these interfaces is attached to one of fourteen LXC bridges. IP addresses of the interfaces are shown in Table II and they are set by the *ip address* command in zebra configuration files.

TABLE II
IP ADDRESSES OF THE INTERFACES ON ALL ROUTERS

Router	Interface name	Interface IP address	LXC bridge
R1	veth1	192.168.1.1/24	br1
R2	veth1	192.168.1.2/24	br1
	veth2	192.168.2.1/24	br2
R3	veth1	192.168.2.2/24	br2
	veth2	192.168.3.1/24	br3
	veth3	192.168.4.1/24	br4
	veth4	192.168.8.2/24	br8
R4	veth1	192.168.3.2/24	br3
	veth2	192.168.5.2/24	br5
R5	veth1	192.168.4.2/24	br4
	veth2	192.168.5.1/24	br5
	veth3	192.168.6.1/24	br6
R6	veth1	192.168.6.2/24	br6
	veth2	192.168.7.1/24	br7
	veth3	192.168.9.1/24	br9
	veth4	192.168.10.2/24	br10
R7	veth1	192.168.7.2/24	br7
	veth2	192.168.8.1/24	br8
R8	veth1	192.168.9.2/24	br9
	veth2	192.168.11.2/24	br11
	veth3	192.168.14.1/24	br14
R9	veth1	192.168.10.1/24	br10
	veth2	192.168.11.1/24	br11
	veth3	192.168.12.1/24	br12
R10	veth1	192.168.12.2/24	br12
	veth2	192.168.13.1/24	br13
R11	veth1	192.168.13.2/24	br13
	veth2	192.168.14.2/24	br14

IV. TEST RESULTS

In this section we will describe tests performed in the network environment described in the previous section. Based on the test results, we will evaluate functionality of the IS-IS implementation in Quagga routing suite. We will describe behavior of the routers of both levels in the case of link failure and when all links are up. Routers R11 and R6 are taken as representatives of the L1 and L1L2 routers, respectively.

Routing table of the router R11 when all of the links in Fig. 3. are up is shown in Fig. 4. It can be observed that router R11 has route to each destination in its area, Area 3. In this area, router R11 forwards packets with respect to the shortest path rule. Thus, packets with destination in network 192.168.9.0/24 and 192.168.11.0/24 will be sent through router R8 (interface veth2), while network 192.168.12.0/24 is reached through router R10 (interface veth1).

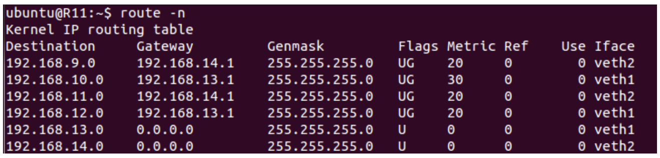

```
ubuntu@R11:~$ route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
192.168.9.0     192.168.14.1    255.255.255.0   UG    20     0        0 veth2
192.168.10.0    192.168.13.1    255.255.255.0   UG    30     0        0 veth1
192.168.11.0    192.168.14.1    255.255.255.0   UG    20     0        0 veth2
192.168.12.0    192.168.13.1    255.255.255.0   UG    20     0        0 veth1
192.168.13.0    0.0.0.0         255.255.255.0   U     0      0        0 veth1
192.168.14.0    0.0.0.0         255.255.255.0   U     0      0        0 veth2
```

Fig. 4. Routing table of the router R11 when all links are up

For destinations in network 192.168.10.0/24, router R11 has two IS-IS paths with equal costs, as shown in Fig. 5. In that case, the router choose the route, which is first received in zebra. For router R11, it is the route through interface veth1.

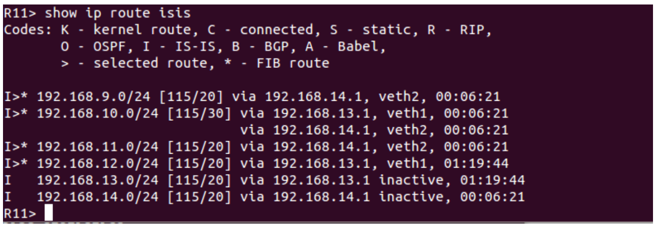

```
R11> show ip route isis
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

I>* 192.168.9.0/24 [115/20] via 192.168.14.1, veth2, 00:06:21
I>* 192.168.10.0/24 [115/30] via 192.168.13.1, veth1, 00:06:21
                             via 192.168.14.1, veth2, 00:06:21
I>* 192.168.11.0/24 [115/20] via 192.168.14.1, veth2, 00:06:21
I>* 192.168.12.0/24 [115/20] via 192.168.13.1, veth1, 01:19:44
I   192.168.13.0/24 [115/20] via 192.168.13.1 inactive, 01:19:44
I   192.168.14.0/24 [115/20] via 192.168.14.1 inactive, 00:06:21
R11> 
```

Fig. 5. IS-IS routes of the router R11 when all links are up

On the other hand, router R11 does not have any route to destinations in the other areas, which is according to the IS-IS protocol. Namely, according to the IS-IS protocol, a L1L2 router should set Attached Bit to value 1 in LSP packets, which it sends to all L1 routers in its own area. Based on this bit and the shortest path rule, each L1 router should choose the closest L1L2 router as a default gateway for destinations in the other area. As a result, the default route should be installed in the routing tables of all L1 routers via chosen L1L2 router. In the case of the router R11, its routing table should contain the default route via router R8. In the routing table of the router R11, however, there is no such route, as shown in Fig. 4. Consequently, router R11 cannot send packets to destinations outside Area 3. This problem occurs because routers R8 and R9 (which are L1L2 routers in Area 3) do not set Attached Bit in their LSPs to router R11, as denoted with the yellow rectangle in Fig. 6. Other L1 routers in network from Fig. 3. have the same problem.

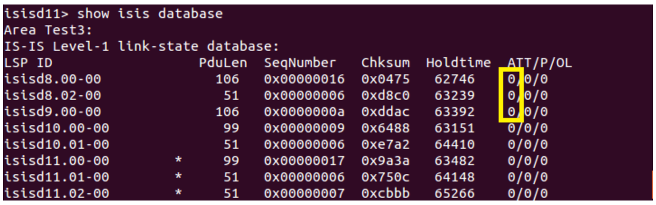

```
isisd11> show isis database
Area Test3:
IS-IS Level-1 link-state database:
LSP ID                      PduLen  SeqNumber   Chksum  Holdtime  ATT/P/OL
isisd8.00-00                   106  0x00000016  0x0475    62746    0/0/0
isisd8.02-00                    51  0x00000006  0xd8c0    63239    0/0/0
isisd9.00-00                   106  0x0000000a  0xddac    63392    0/0/0
isisd10.00-00                   99  0x00000009  0x6488    63151    0/0/0
isisd10.01-00                   51  0x00000006  0xe7a2    64410    0/0/0
isisd11.00-00            *      99  0x00000017  0x9a3a    63482    0/0/0
isisd11.01-00            *      51  0x00000006  0x750c    64148    0/0/0
isisd11.02-00            *      51  0x00000007  0xcbbb    65266    0/0/0
```

Fig. 6. Level-1 link-state database of the router R11

Fig. 7. shows the routing table of the router R6 when all links are up. This L1L2 router has routes to all destinations in its area and to the networks connected to the L1L2 routers from other areas. It can be observed that router R6 sends packets to router R2 (network 192.168.1.0/24) via router R5 (interface veth1), although it has path with the same cost via router R7 (veth2). However, according to the IS-IS protocol routing table of the router R6 is correct, since router R7 is L1 router and router R5 belongs to level-2 as router R2. On the other hand, router R6 has two paths to the Area 3, as a consequence of the direct connections with routers R8 and R9. It sends packets to the destinations in Area3 using the shortest path rule. Thus, it forwards traffic to the networks 192.168.12.0/24 and 192.168.14.0/24 via interfaces veth4 and veth3, respectively. In Area 2, router R6 behaves as an L1 router. Consequently, R6 stores routes to the destinations in this area in its routing table similarly as the R11.

Problem with routers at level-2 is that they do not have routes to the networks connected to L1 routers from other areas. As a result, L1L2 routers cannot send packets to these networks. According to the IS-IS protocol, each L1L2 router should advertise routes from level-1 to level-2. In such a way, it allows a L1L2 (or L2) router from another area to communicate with destinations, which are connected to L1 routers in its area. This functionality of the isisd daemon does not work properly, as we can see in Fig. 7, which does not show route from router R6 to the network 192.168.13.0/24.

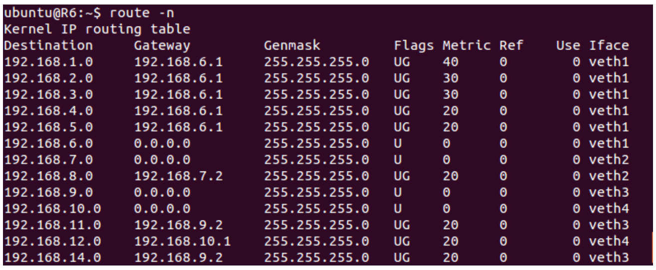

```
ubuntu@R6:~$ route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
192.168.1.0     192.168.6.1     255.255.255.0   UG    40     0        0 veth1
192.168.2.0     192.168.6.1     255.255.255.0   UG    30     0        0 veth1
192.168.3.0     192.168.6.1     255.255.255.0   UG    30     0        0 veth1
192.168.4.0     192.168.6.1     255.255.255.0   UG    20     0        0 veth1
192.168.5.0     192.168.6.1     255.255.255.0   UG    20     0        0 veth1
192.168.6.0     0.0.0.0         255.255.255.0   U     0      0        0 veth1
192.168.7.0     0.0.0.0         255.255.255.0   U     0      0        0 veth2
192.168.8.0     192.168.7.2     255.255.255.0   UG    20     0        0 veth2
192.168.9.0     0.0.0.0         255.255.255.0   U     0      0        0 veth3
192.168.10.0    0.0.0.0         255.255.255.0   U     0      0        0 veth4
192.168.11.0    192.168.9.2     255.255.255.0   UG    20     0        0 veth3
192.168.12.0    192.168.10.1    255.255.255.0   UG    20     0        0 veth4
192.168.14.0    192.168.9.2     255.255.255.0   UG    20     0        0 veth3
```

Fig. 7. Routing table of the router R6 when all links are up

Fig. 8. shows a routing table of the router R11 when the link fails between this router and router R8. We simulated the link failure by shutting down the veth2 interface on router R11. In Fig. 8., it can be observed that the router R11 sends traffic to the networks 192.168.9.0/24 and 192.168.11.0/24 via router R10 (interface veth1) after the link failure. These routes differ from the routes to the same networks in the case when the link between routers R11 and R8 is up, which are via router R8 (veth2), as shown in Fig.4. Also, after the link failure, router R11 gets a route to the network 192.168.14.0/24 via interface veth1, because it has lost a direct connection to this network. Based on the above, it can be concluded that a L1 router successfully updates its routing table in the case of a link failure.

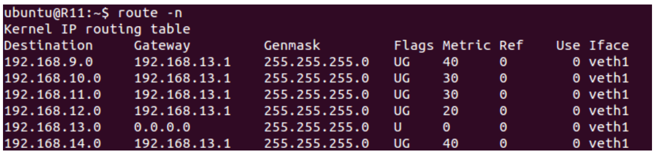

```
ubuntu@R11:~$ route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
192.168.9.0     192.168.13.1    255.255.255.0   UG    40     0        0 veth1
192.168.10.0    192.168.13.1    255.255.255.0   UG    30     0        0 veth1
192.168.11.0    192.168.13.1    255.255.255.0   UG    30     0        0 veth1
192.168.12.0    192.168.13.1    255.255.255.0   UG    20     0        0 veth1
192.168.13.0    0.0.0.0         255.255.255.0   U     0      0        0 veth1
192.168.14.0    192.168.13.1    255.255.255.0   UG    40     0        0 veth1
```

Fig. 8. Routing table of the router R11 after the link failure

At the end of this section, we will describe scenario when fails link between routers R5 and R3. This scenario will show changes in the router's R6 routing table after the link failure. The routing table is shown in Fig. 9. In section II, we mentioned that the IS-IS backbone must be a collection of the interconnected L2 (or L1L2) routers. When a link between routers R5 and R3 fails, level-2 backbone becomes partitioned. As a consequence, router R6 lose route to the router R2 (network 192.168.1.0/24), as shown in Fig. 9. This occurs even though router R6 still has a path to the router R2 via router R7, which is the L1 router and does not participate in level-2 routing. After the link failure, router R6 also changes level-1 routes. It forwards packets to the networks 192.168.2.0/24 and 192.168.4.0/24 (attached to the router R3) via router R7 (interface veth2) instead through router R5, as it was before the link failure. This occurs because paths via router R7 are shorter than the route via router R5 after the link failure. Therefore, it can be concluded that the routing table of a L1L2 router is correctly updated in case of the link failure.

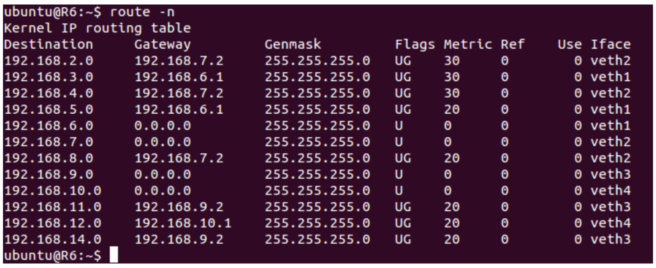

```
ubuntu@R6:~$ route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
192.168.2.0     192.168.7.2     255.255.255.0   UG    30     0        0 veth2
192.168.3.0     192.168.6.1     255.255.255.0   UG    30     0        0 veth1
192.168.4.0     192.168.7.2     255.255.255.0   UG    30     0        0 veth2
192.168.5.0     192.168.6.1     255.255.255.0   UG    20     0        0 veth1
192.168.6.0     0.0.0.0         255.255.255.0   U     0      0        0 veth1
192.168.7.0     0.0.0.0         255.255.255.0   U     0      0        0 veth2
192.168.8.0     192.168.7.2     255.255.255.0   UG    20     0        0 veth2
192.168.9.0     0.0.0.0         255.255.255.0   U     0      0        0 veth3
192.168.10.0    0.0.0.0         255.255.255.0   U     0      0        0 veth4
192.168.11.0    192.168.9.2     255.255.255.0   UG    20     0        0 veth3
192.168.12.0    192.168.10.1    255.255.255.0   UG    20     0        0 veth4
192.168.14.0    192.168.9.2     255.255.255.0   UG    20     0        0 veth3
ubuntu@R6:~$ 
```

Fig. 9. Routing table of the router R6 after the link failure

V. CONCLUSION

Results obtained from realized tests in the created network environment show that some functionalities of the IS-IS implementation in Quagga do not work properly. As we described in the previous section, there are routing problems in the IS-IS domains with multiple areas. It has been observed that L2 (or L1L2) routers in an area do not set Attached Bit in LSPs to L1 routers, as they should. As a result, L1 routers do not have default routes through the closest L2 (or L1L2) routers in their routing tables and cannot forwarding packets outside their area. Furthermore, L1L2 routers do not set networks reachable via L1 routing in LSPs to other L1L2 (or L2) routers. Consequently, L2 (or L1L2) routers cannot reach subnets on L1 routers from other areas.

In the case of the domain with a single area or when all routers are configured as L2 (or L1L2), the isisd daemon creates routing tables correctly. Also, routing tables of all routers are updated successfully in the case of the link failure. Since IS-IS users (including the largest ISPs) typically use one area per domain (usually just a backbone), the Quagga's IS-IS implementation can serve their needs.

On the other hand, Quagga is open-source routing software, which allows users to adapt it to their needs. Through testing of the most relevant networking scenarios, this paper presented the space for further improvements of the Quagga's IS-IS deamon.

ACKNOWLEDGMENT

We express gratitude to the Serbian Ministry of education, Science and Technologies, Telekom Srbija and Informatika for their financial support.

REFERENCES

[1] J. Hawkinson, T. Bates, “Guidelines for creation, selection, and registration of an Autonomous System (AS)”, RFC 1930, Network Working Group, 1996

[2] Ashwin Sridharan, Roch Guérin, Christophe Diot, “Achieving near-optimal traffic engineering solutions for current OSPF/IS-IS networks”, TON, April 2005

[3] Athina Markopoulou, Gianluca Iannaccone, Supratik Bhattacharyya, Chen-Nee Chuah, Yashar Ganjali, Christophe Diot, “Characterization of failures in an operational IP backbone network”, TON, August 2008

[4] Eljack, S.M., Abdelkarim, S.B., “Effect of the interior gateway routing protocols in the multiprotocol label switching networks a simulated comprison between IS-IS & OSPF routing protocols”, International Conference on Computing, Electrical and Electronics Engineering (ICCEEE), Khartoum, Sudan, August 2013

[5] Quagga Routing Suite [Online].
Available: http://www.nongnu.org/quagga/

[6] O. Sharon, “Dissemination of routing information in broadcast networks: OSPF versus IS-IS”, Network, IEEE, 2001

[7] R. Callon, “Use of OSI IS-IS for Routing in TCP/IP and Dual Environments”, RFC 1195, Network Working Group, December 1990