Experimental results about Multiprotocol Routing and Route Redistribution

Eugen PETAC, *Member IEEE*, Bogdan MUSAT

***Abstract* — This paper considers an experimental platform and some results about the problem of multi protocol routing and route redistribution. Our particular lab platform, with Zebra/ Quagga routing software packages, is analyzed. The particularity of this case study is that we can link normal computers running Zebra / Quagga together with Cisco dedicated equipments. Interrupting various links between routers for studying convergence time as well as shutting down various routers to study network flexibility led to results according to routing protocols specifications.**

***Keywords* — IGRP, OSPF, Quagga, RIP, route redistribution, routing, VLAN, Zebra**

I. Introduction

THIS paper studies a lab platform which is used to experiment multiprotocol routing and route redistribution. The particularity of this case study platform is that, besides existent dedicated network equipment, a number of general computer systems is added. On these systems runs the Zebra / Quagga software package. Zebra [6] and Quagga [5] are two routing software packages which support traditional IPv4 routing protocols but also support IPv6 routing protocols. Quagga is an alternative to GNU Zebra (last version is zebra-0.94), the latter being developed by Kunihiro Ishiguro. These two software packages provide TCP/IP services and routing protocols, like: RIP version 1 (RIPv1), RIP version 2 (RIPv2), OSPF version 2, OSPF version 3, BGP version 4 and BGP version 4+.

II. Platform Description

The case study is developed on a scenario in which it is supposed that a company owns a building in Bucharest as headquarters and two branches in Constanta and Iasi. The Constanta branch has four departments: IT, Design, Human Resource and Public Relations. Each department has its own computer network. To ease organizing the network by departments, VLANs (Virtual LAN) were used, because they allow for logical grouping, no matter of physical location. All computers within a department were connected to the same VLAN, like this: IT – VLAN ID 10, Design – VLAN ID 20, Human Resources – VLAN ID 30, Public Relations – VLAN ID 40.

The connection between branches and headquarters is based on a Frame Relay network, with the two ends of each connection being Ethernet based. Because of didactic reasons, the Frame Relay network does not appear on network schematic, the link between Zebra2 and Cisco2513 being a direct one. The Iasi branch has a single LAN and the connection with headquarters is done through a serial link. Headquarters in Bucharest also has a LAN, but it also has an Internet connection. This Internet connection is used also by branches.

Each branch can choose its own routing protocol. Each IT department from each city chose a routing protocol, like this: Constanta branch has chosen OSPF, Iasi branch has chosen EIGRP and headquarters from Bucharest has chosen RIP.

Between each pair of protocols there will be route redistribution, because, by default, routes learned from a routing protocol does not automatically redistribute into another routing protocol.

Network topology is a laboratory platform (figure 1) on which all scenarios were tested. The diagram in figure 1 contains the following network equipments: 2 Cisco 2501 routers (named Cisco2501A and Cisco2501B), 1 Cisco 2511 router, 1 Cisco 2513 router, 4 computers with Linux Redhat 9 and Fedora Core 2, running Zebra / Quagga software; 2 WAN modems, RAD brand, model ASM-31; 2 Allied Telesyn AT8024 switches; 6 host computers, 1 server running Zebra / Quagga connected to Internet. The two RAD modems can be used to setup a WAN connection between Cisco2501A and Cisco2511 routers thru a single pair of copper wires on a distance of maximum 5.5 km.

- Cisco2501A router has two connected interfaces: eternet0 and serial0. Ethernet0 interface is linked with Ethernet interface of the Internet server via a crossover cable. Serial0 interface is connected to one of the RAD modems.
- Cisco2511 router has three connected interfaces:

E. Petac, Ph.D. is with the Faculty of Mathematics and Informatics, “Ovidius” University of Constantza, Romania and with the ICT Foundation, Constantza, Romania, Timisana 15, 900524 Constantza, Romania (phone: 40/241/617298; fax: 40/241/617298; e-mail: epetac@gmail.com).

B. Musat, M. Sc. was with the Faculty of Mathematics and Informatics, “Ovidius” University of Constantza, Romania. He is now with ICT Foundation, Constantza, Timisana 15, 900524 Constantza, Romania (phone: 40/241/617298; fax: 40/241/617298; e-mail: musatb2000@yahoo.com).

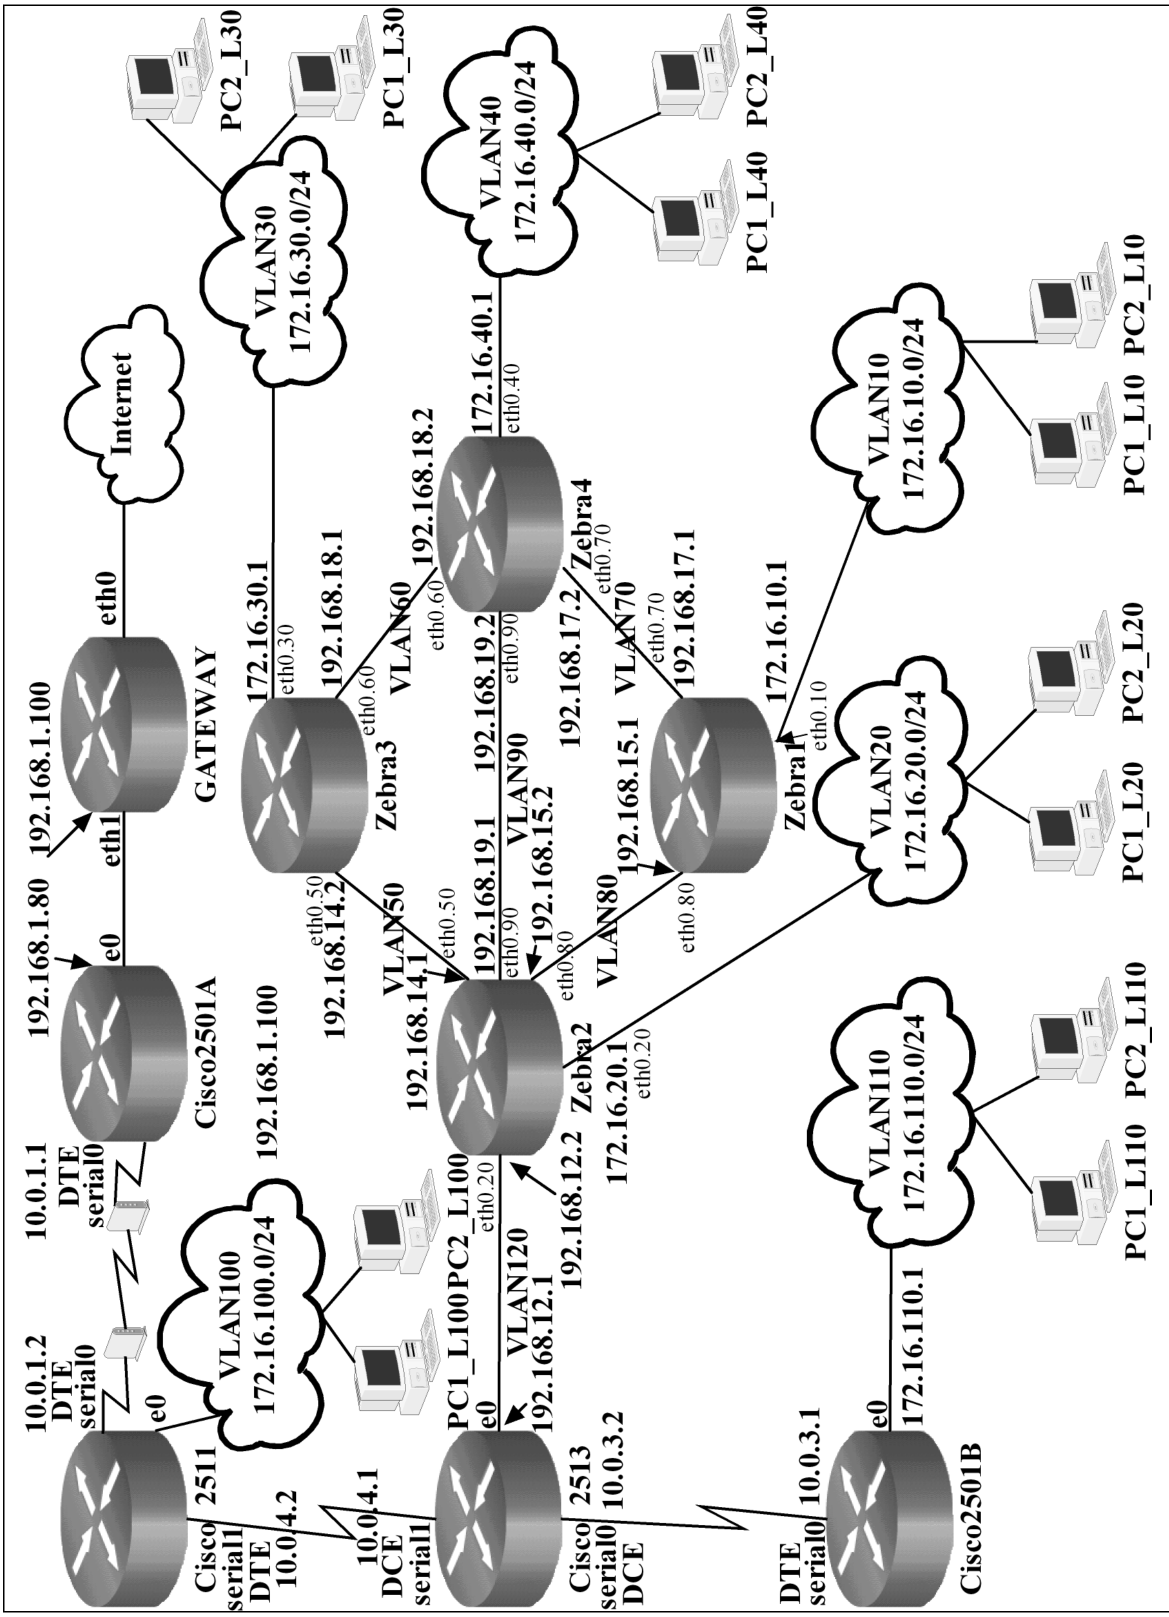

Figure 1: Network schematic

Ethernet0, Serial0 and Serial1. Ethernet0 interface is connected to Switch2 on port 1 (VLAN ID 100). In this VLAN there is also PC1_L100 host. Serial0 interface connects to the other RAD modem. Serial1 interface is connected with serial1 interface of Cisco2513 router.

• Cisco2513 router has three connected interfaces: ethernet0, serial0 and serial1. Ethernet0 interface is connected to Switch2 on port 2 (VLAN ID 120). In this VLAN there is also eth0.120 sub-interface of Zebra2 router. Serial0 interface is connected to serial0 interface of Cisco2501B router. Serial1 interface is connected with serial1 interface of Cisco2511 router.

• Cisco2501B router has two connected interfaces: ethernet0 and serial0. Ethernet0 interface is connected to Switch2 on port 3 (VLAN ID 110). In this VLAN there is also PC1_L110 host. Serial0 interface is connected to serial0 interface of Cisco2513 router.

• Zebra1 router is a computer running Linux Redhat 9. It has a physical interface (eth0) under which there are defined three sub-interfaces (eth0.10, eth0.70, eth0.80) for 10, 70, 80 VLAN IDs. Physical interface eth0 is connected to Switch1 on port 1.

• Zebra2 router is a computer running Linux Redhat 9. It has a physical interface (eth0) under which there are defined five sub-interfaces (eth0.20, eth0.50, eth0.80, eth0.90, eth0.120) for 20, 50, 80, 90, 120 VLAN IDs. Physical interface eth0 is connected to Switch1 on port 2.

• Zebra3 router is a computer running Fedora Core 2. It has a physical interface (eth0) under which there are defined three sub-interfaces (eth0.30, eth0.50, eth0.60) for 30, 50, 60 VLAN IDs. Physical interface eth0 is connected to Switch1 on port 3.

• Zebra4 router is a computer running Fedora Core 2. It has a physical interface (eth0) under which there are defined four sub-interfaces (eth0.40, eth0.60, eth0.70, eth0.90) for 40, 60, 70, 90 VLAN IDs. Physical interface eth0 is connected to Switch1 on port 4.

• In Switch1 are connected the eth0 physical interfaces of Zebra routers: eth0 from Zebra1 in port1, eth0 from Zebra2 in port 2, eth0 from Zebra3 in port 3, eth0 from Zebra4 in port 4.

• In Switch2 are connected: e0 interface of Cisco2511 router on port 1, e0 interface of Cisco2513 router on port 2, e0 interface of Cisco2501B router on port 3, PC1_L10 host on port 4, PC1_L20 host on port 5, PC1_L30 host on port 6, PC1_L40 host on port 7, PC1_L100 host on port 8, PC1_L110 host on port 9, PC2_L10 host on port 11, PC2_L20 host on port 12, PC2_L30 host on port 13, PC2_L40 host on port 14, PC2_L100 host on port 15, PC2_L110 host on port 16. The two switches are interconnected together through a trunk link, a link which carries multiple VLANs.

• GATEWAY router is a computer which runs Fedora Core 5. Eth0 interface of this router is connected to ISP via a cable modem.

The routing protocols used are listed in Table 1 and the VLANs used are listed in Table 2:

Table 1: Routing protocols

Router	**Routing protocol**
Cisco2501A	RIP
Cisco2511	RIP
Cisco2513	RIP, EIGRP, OSPF
Cisco2501B	EIGRP
Zebra1	OSPF
Zebra2	OSPF
Zebra3	OSPF

Table 2: VLAN definition

VLAN ID	**Subnet address / Netmask**
1	VLAN implicit
10	172.16.10.0/24
20	172.16.20.0/24
30	172.16.30.0/24
40	172.16.40.0/24
50	172.16.14.0/24
60	192.168.18.0/24
70	172.168.17.0/24
80	192.168.15.0/24
90	192.168.19.0/24
100	192.168.100.0/24
110	192.168.110.0/24

I. Theoretical Issues

Redistribution [1], [4] means using a routing protocol to advertise routes that are learned by other means, such as: another routing protocol, static routes or directly connected routes.

It is recommended that a single routing protocol be used in an IP network. On the other hand, multi-protocol routing is required in a number of situations, such as: merging companies, multiple departments managed by many network administrators and network media which use different vendor models. Running multiple routing protocols is a requirement needed since the planning phase of a network. Therefore, the existence of a communication medium running multiple routing protocols makes redistribution mandatory. The differences between characteristics of routing protocols, such as metrics, administrative distance, classfull and classless capabilities, may affect redistribution. For correctly configuring redistribution, all these aspects should be taken into consideration.

When you do route redistribution from one protocol to another, the metrics in each protocol plays an important role in redistribution. It is known that each protocol uses different metrics. We remind you that RIP metric is based on hop-count, while IGRP and EIGRP uses a composite metric based on bandwidth, delay, reliability, load and MTU, where by default bandwidth and delay are the only factors taken into consideration. When the routes are redistributed, a metric must be defined, which the protocol which receives those routes must understand. There are two methods of defining the metrics when you do route redistribution: either you can define a metric for that particular redistribution, or you can use the same metric as a default metric for all redistribution. By using the **default-metric** command you can save some time since there is no need to define separate metrics for each redistribution.

A rule of redistribution in a Cisco router [2], [4] states that the route which will be redistributed must be present in the routing table. The default behavior of RIP, IGRP and EIGRP routing protocols is to advertise directly connected networks when a **network** type statement is made inside the configuration of a routing protocol for subnets belonging to directly connected interface. There are two methods [3] of obtaining a directly connected route:

1. An interface is configured with an IP address and a subnet mask and the resulting subnet is considered a directly connected route;

2. A static route is configured with an exit interface only and no next-hop IP address. This kind of route is also considered a directly connected route.

Router Cisco

```
Router# configure terminal
Router(config)# ip route 172.16.20.0
255.255.255.0 ethernet 0/0
Router(config)# end

Router# show ip route static
   172.16.20.0/24 is subnetted, 1 subnets
   S 172.16.20.0 is directly connected,
Ethernet0/0
```

The **network** command used under EIGRP, RIP or IGRP configuration, which has as a parameter any of these directly connected route types, enables the advertising of the respective route.

For example, if an interface has assigned the 172.16.20.1 IP address with a 255.255.255.0 subnet mask, then the 172.16.20.0/24 subnet will be advertised by these routing protocols whenever a **network** type statement is configured as follows:

```
router rip | igrp # | eigrp #
   network 172.16.0.0
```

The static route 172.16.20.0/24 is also advertised by these routing protocols, because it is a directly connected route and is “covered” by the **network** statement.

II. Experimental Results

In this case study we run through different stages: primary configuration of network equipments (hostnames, name servers, passwords, interfaces, static routes, etc.); VLANs were implemented for Zebra routers – vconfig utility and system scripts; RIP protocol configuration on Cisco2501A, Cisco2511, Cisco2513, Internet server routers; EIGRP protocol configuration – Cisco2513 and Cisco2501B routers; OSPF protocol configuration – Cisco2513 router.
For starting OSPF routing on Zebra1 we used zebra and ospfd daemons. Router-ids are listed in table 3. We used OSPF with authentication, with md5 hash password being cisco. The Zebra2, Zebra3 and Zebra4 routers were configured in the same manner.

Table 3: Router-IDs

Router	Router-ID
Zebra1	192.168.250.1
Zebra2	192.168.250.2
Zebra3	192.168.250.3
Zebra4	192.168.250.4

Zebra1

```
[root@zebra1 root]# telnet 127.0.0.1 2604
Trying 127.0.0.1...
Connected to 127.0.0.1.
Escape character is '^]'.

Hello, this is zebra (version 0.93b).
Copyright 1996-2002 Kunihiro Ishiguro.


User Access Verification

Password:
zebra1-ospfd> enable
zebra1-ospfd# configure terminal
zebra1-ospfd(config)# interface eth0.10
zebra1-ospfd(config-if)# ip ospf message-
digest-key 1 md5 cisco
zebra1-ospfd(config-if)# interface
eth0.70
zebra1-ospfd(config-if)# ip ospf message-
digest-key 1 md5 cisco
zebra1-ospfd(config-if)# interface
eth0.80
zebra1-ospfd(config)# router ospf
zebra1-ospfd(config-router)# ospf router-
id 192.168.250.1
zebra1-ospfd(config-router)# redistribute
connected
zebra1-ospfd(config-router)# network
172.16.10.0/24 area 0
zebra1-ospfd(config-router)# network
192.168.15.0/24 area 0
zebra1-ospfd(config-router)# network
192.168.17.0/24 area 0
zebra1-ospfd(config-router)# area 0
authentication message-digest
```

The results clearly shows each successful stage: redistribution configuration, static route redistribution in RIP, bidirectional route redistribution between RIP and EIGRP, bidirectional route redistribution between RIP and OSPF, bidirectional route redistribution between OSPF and EIGRP.

III. Conclusion

This study case used on the lab platform allows configuring routers and switches, connectivity testing and troubleshooting, basic configuration of RIP, EIGRP, OSPF and configuring route redistribution between RIP, EIGRP and OSPF routing protocols. The particularity of this case study is that we can link normal computers running Zebra / Quagga together with Cisco dedicated equipments.
Interrupting various links between routers for studying convergence time as well as shutting down various routers to study network flexibility led to results according to routing protocols specifications.
Multi-protocol routing is a requirement in a number of situations, such as: merging companies, multiple departments managed by many network administrators and network media which use different vendor models. Running multiple routing protocols is a requirement needed since the planning phase of a network. Therefore, the existence of a communication medium running multiple routing protocols makes redistribution mandatory. This case study outlined the success of route redistribution between different routing protocols.

Appendix

1. Routing table listing from all routers

Router Cisco2501A

```
Cisco2501A#show ip route
Codes: C - connected, S - static, I - IGRP, R
- RIP, M - mobile, B – BGP, D - EIGRP, EX -
EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA
external type 2, E1 - OSPF external type 1, E2
- OSPF external type 2, E – EGP, i - IS-IS, L1
- IS-IS level-1, L2 - IS-IS level-2, * -
candidate default, U - per-user static route,
o - ODR
Gateway of last resort is 192.168.1.100 to
network 0.0.0.0

R 192.168.110.0/24 [120/2] via 10.0.1.2,
00:00:16, Serial0
10.0.0.0/24 is subnetted, 3 subnets
R 10.0.3.0 [120/2] via 10.0.1.2, 00:00:17,
Serial0
C 10.0.1.0 is directly connected, Serial0
R 10.0.4.0 [120/1] via 10.0.1.2,00:00:17,S0
R 192.168.100.0/24 [120/1] via 10.0.1.2,
00:00:17, Serial0
R 192.168.12.0/24 [120/2]via10.0.1.2,
00:00:17, Serial0
R 192.168.14.0/24 [120/2] via 10.0.1.2,
00:00:17, Serial0
R 192.168.15.0/24 [120/2] via 10.0.1.2,
00:00:17, Serial0
C 192.168.1.0/24 is directly connected, Eth0
R 192.168.17.0/24 [120/2] via 10.0.1.2,
00:00:17, Serial0
R 192.168.18.0/24 [120/2] via 10.0.1.2,
00:00:17, Serial0
R 192.168.19.0/24 [120/2] via 10.0.1.2,
00:00:17, Serial0
R 172.16.0.0/16 [120/2] via 10.0.1.2,
00:00:17, Serial0
S* 0.0.0.0/0 [1/0] via 192.168.1.100
```

Router Cisco2511

```
Cisco2511#show ip route
Codes: C - connected, S - static, I - IGRP, R
- RIP, M - mobile, B – BGP, D - EIGRP, EX -
EIGRP external, O - OSPF, IA - OSPF inter
area, N1 - OSPF NSSA external type 1, N2 -
OSPF NSSA external type 2, E1 - OSPF external
type 1, E2 - OSPF external type 2, E – EGP, i
- IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-
2, * - candidate default, U - per-user static
route, o - ODR
Gateway of last resort is 10.0.1.1 to network
0.0.0.0

R 192.168.110.0/24 [120/1] via 10.0.4.1,
00:00:03, Serial1
10.0.0.0/24 is subnetted, 3 subnets
R 10.0.3.0 [120/1] via 10.0.4.1, 00:00:03, S1
C 10.0.1.0 is directly connected, S0
C 10.0.4.0 is directly connected, S1
C 192.168.100.0/24 is directly connected, Eth0
R 192.168.12.0/24 [120/1] via 10.0.4.1,
00:00:03, S1
R    192.168.14.0/24 [120/1] via 10.0.4.1,
00:00:03, Serial1
R    192.168.15.0/24 [120/1] via 10.0.4.1,
00:00:03, Serial1
R    192.168.1.0/24 [120/1] via 10.0.1.1,
00:00:00, Serial0
R    192.168.17.0/24 [120/1] via 10.0.4.1,
00:00:03, Serial1
R    192.168.18.0/24 [120/1] via 10.0.4.1,
00:00:04, Serial1
R    192.168.19.0/24 [120/1] via 10.0.4.1,
00:00:04, Serial1
R    172.16.0.0/16 [120/1] via 10.0.4.1,
00:00:04, Serial1
R*   0.0.0.0/0 [120/1]via10.0.1.1,00:00:02,S0
```

Router Cisco2513

```
Cisco2513#show ip route
Codes: C - connected, S - static, I - IGRP, R
- RIP, M - mobile, B – BGP, D - EIGRP, EX -
EIGRP external, O - OSPF, IA - OSPF inter
area, N1 - OSPF NSSA external type 1, N2 -
OSPF NSSA external type 2, E1 - OSPF external
type 1, E2 - OSPF external type 2, E – EGP, i
- IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-
2, ia - IS-IS inter area, * - candidate
default, U - per-user static route, o – ODR, P
- periodic downloaded static route
Gateway of last resort is 10.0.4.2 to network
0.0.0.0

C    192.168.12.0/24 is directly connected,
Ethernet0
O    192.168.14.0/24 [110/20] via
192.168.12.2, 00:01:03, Ethernet0
O    192.168.15.0/24 [110/20] via
192.168.12.2, 00:01:03, Ethernet0
D    192.168.110.0/24 [90/1787392] via
10.0.3.1, 05:31:46, Serial0
     172.16.0.0/24 is subnetted, 4 subnets
O       172.16.40.0 [110/30] via 192.168.12.2,
00:01:03, Ethernet0
O E2    172.16.30.0 [110/20] via 192.168.12.2,
00:01:03, Ethernet0
O       172.16.20.0 [110/20] via 192.168.12.2,
00:01:03, Ethernet0
O       172.16.10.0 [110/30] via 192.168.12.2,
00:01:04, Ethernet0
     192.168.250.0/32 is subnetted, 1 subnets
C       192.168.250.5 is directly connected,
Loopback0
     10.0.0.0/24 is subnetted, 3 subnets
C       10.0.3.0 is directly connected,
Serial0
R       10.0.1.0 [120/1] via 10.0.4.2,
00:00:05, Serial1
C       10.0.4.0 is directly connected,
Serial1
O    192.168.17.0/24 [110/30] via
192.168.12.2, 00:01:04, Ethernet0
R    192.168.1.0/24 [120/2] via 10.0.4.2,
00:00:05, Serial1
O    192.168.19.0/24 [110/20] via
192.168.12.2, 00:01:04, Ethernet0
R    192.168.100.0/24 [120/1] via 10.0.4.2,
00:00:05, Serial1
O    192.168.18.0/24 [110/30] via
192.168.12.2, 00:01:04, Ethernet0
R*   0.0.0.0/0 [120/2] via 10.0.4.2, 00:00:05,
Serial1
```

Router Cisco2501B

```
Cisco2501B#show ip route
Codes: C - connected, S - static, I - IGRP, R
- RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O -
OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF
external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-
IS level-2, * - candidate default
       U - per-user static route

Gateway of last resort is 10.0.3.2 to network
0.0.0.0

C    192.168.110.0/24 is directly connected,
Ethernet0
     10.0.0.0/24 is subnetted, 3 subnets
```

```
C       10.0.3.0 is directly connected,
Serial0
D EX    10.0.1.0 [170/2195456] via 10.0.3.2,
04:05:21, Serial0
D       10.0.4.0 [90/2681856] via 10.0.3.2,
05:42:03, Serial0
D EX 192.168.100.0/24 [170/2195456] via
10.0.3.2, 04:05:21, Serial0
D EX 192.168.12.0/24 [170/2195456] via
10.0.3.2, 03:56:29, Serial0
D EX 192.168.14.0/24 [170/2195456] via
10.0.3.2, 03:56:29, Serial0
D EX 192.168.15.0/24 [170/2195456] via
10.0.3.2, 03:56:29, Serial0
D EX 192.168.1.0/24 [170/2195456] via
10.0.3.2, 04:05:21, Serial0
D EX 192.168.17.0/24 [170/2195456] via
10.0.3.2, 03:56:29, Serial0
D EX 192.168.18.0/24 [170/2195456] via
10.0.3.2, 03:56:29, Serial0
D EX 192.168.19.0/24 [170/2195456] via
10.0.3.2, 03:56:29, Serial0
     172.16.0.0/24 is subnetted, 4 subnets
D EX    172.16.40.0 [170/2195456] via
10.0.3.2, 03:56:29, Serial0
D EX    172.16.30.0 [170/2195456] via
10.0.3.2, 00:01:44, Serial0
D EX    172.16.20.0 [170/2195456] via
10.0.3.2, 03:56:30, Serial0
D EX    172.16.10.0 [170/2195456] via
10.0.3.2, 03:56:30, Serial0
D*EX 0.0.0.0/0 [170/2195456] via 10.0.3.2,
04:05:23, Serial0
```

Zebra1

```
[root@zebra1 root]# route -n
Kernel IP routing table
Destination    Gateway    Flags Met Ref Use Intf
192.168.100.0  192.168.15.2 UG 1  0  0 eth0.80
172.16.20.0    192.168.15.2 UG 20 0  0 eth0.80
192.168.19.0   192.168.17.2 UG 20 0  0 eth0.70
192.168.18.0   192.168.17.2 UG 20 0  0 eth0.70
192.168.1.0    192.168.15.2 UG 1  0  0 eth0.80
192.168.17.0   0.0.0.0      U  0  0  0 eth0.70
192.168.15.0   0.0.0.0      U  0  0  0 eth0.80
172.16.30.0    192.168.17.2 UG 20 0  0 eth0.70
192.168.14.0   192.168.15.2 UG 20 0  0 eth0.80
192.168.110.0  192.168.15.2 UG 1  0  0 eth0.80
192.168.12.0   192.168.15.2 UG 20 0  0 eth0.80
172.16.10.0    0.0.0.0      U  0  0  0 eth0.10
172.16.40.0    192.168.17.2 UG 20 0  0 eth0.70
127.0.0.0      0.0.0.0      U  0  0  0 lo
0.0.0.0        192.168.15.2 UG 1  0  0 eth0.80
```

Zebra2

```
[root@zebra2 root]# route -n
Kernel IP routing table
Destination    Gateway    Flags Met Ref Use Intf
192.168.100.0 192.168.12.1 UG   1 0  0 eth0.120
172.16.20.0   0.0.0.0      U    0 0  0 eth0.20
192.168.19.0  0.0.0.0      U    0 0  0 eth0.90
192.168.18.0  192.168.14.2 UG 20 0  0 eth0.50
192.168.1.0   192.168.12.1 UG   1 0  0 eth0.120
192.168.17.0  192.168.15.1 UG 20 0  0 eth0.80
192.168.15.0  0.0.0.0      U    0 0  0 eth0.80
172.16.30.0   192.168.14.2 UG 20 0  0 eth0.50
192.168.14.0  0.0.0.0      U    0 0  0 eth0.50
192.168.110.0 192.168.12.1 UG   1 0  0 eth0.120
192.168.12.0  0.0.0.0      U    0 0  0 eth0.120
172.16.10.0  192.168.15.1  UG 20  0  0 eth0.80
172.16.40.0  192.168.19.2  UG 20  0  0 eth0.90
127.0.0.0    0.0.0.0       U  0   0  0 lo
0.0.0.0      192.168.12.1  UG 1   0  0eth0.120
```

Zebra3

```
[root@zebra3 root]# route -n
Kernel IP routing table
Destination   Gateway    Flags Met Ref Use Intf
192.168.100.0 192.168.14.1 UG 1  0  0  eth0.50
172.16.20.0   192.168.14.1 UG 20 0  0  eth0.50
192.168.19.0  192.168.14.1 UG 20 0  0  eth0.50
192.168.18.0  0.0.0.0      U  0  0  0  eth0.60
192.168.1.0   192.168.14.1 UG 1  0  0  eth0.50
192.168.17.0  192.168.18.2 UG 20 0  0  eth0.60
192.168.15.0  192.168.14.1 UG 20 0  0  eth0.50
172.16.30.0   0.0.0.0      U  0  0  0  eth0.30
192.168.14.0  0.0.0.0      U  0  0  0  eth0.50
192.168.110.0 192.168.14.1 UG 1  0  0  eth0.50
192.168.12.0  192.168.14.1 UG 20 0  0  eth0.50
172.16.10.0   192.168.14.1 UG 30 0  0  eth0.50
172.16.40.0   192.168.18.2 UG 20 0  0  eth0.60
127.0.0.0     0.0.0.0      U  0  0  0  lo
0.0.0.0       192.168.14.1 UG 1  0  0  eth0.50
```

Zebra4

```
[root@zebra4 root]# route -n
Kernel IP routing table
Destination   Gateway    Flags Met Ref Use Intf
192.168.100.0 192.168.19.1 UG 1  0  0  eth0.90
172.16.20.0   192.168.19.1 UG 20 0  0  eth0.90
192.168.19.0  0.0.0.0      U  0  0  0  eth0.90
192.168.18.0  0.0.0.0      U  0  0  0  eth0.60
192.168.1.0   192.168.19.1 UG 1  0  0  eth0.90
192.168.17.0  0.0.0.0      U  0  0  0  eth0.70
192.168.15.0  192.168.17.1 UG 20 0  0  eth0.70
172.16.30.0   192.168.18.1 UG 20 0  0  eth0.60
192.168.14.0  192.168.18.1 UG 20 0  0  eth0.60
192.168.110.0 192.168.19.1 UG 1  0  0  eth0.90
192.168.12.0  192.168.19.1 UG 20 0  0  eth0.90
172.16.10.0   192.168.17.1 UG 20 0  0  eth0.70
172.16.40.0   0.0.0.0      U  0  0  0  eth0.40
127.0.0.0     0.0.0.0      U  0  0  0  lo
0.0.0.0       192.168.19.1 UG 1  0  0  eth0.90
[root@zebra4 root]#
```

REFERENCES

[1] Moy J. - OSPF. Complete Implementation, Addison-Wesley, October 2000.

[2] Cisco Systems. Configuring OSPF, http://www.cisco.com/univercd/cc/ td/doc/product/software/ ios122/122cgcr/fipr_c/ipcprt2/1cfospf.htm

[3] Cisco Systems. Interior Gateway Routing Protocol (IGRP). Cisco TAC Technology pages, http://www.cisco.com/pcgibin/Support/PSP/psp_view.pl?p =Internetworking:IGRP

[4] Doyle J., Carroll J. - Routing TCP/IP, Cisco Press 2001

[5] *** Quagga Routing Software Suite, http://www.quagga.net/

[6] *** GNU Zebra, http://www.zebra.org