

# **Convergence Behavior of RIP and OSPF Network Protocols**

By

Hubert Pun

B.A.Sc., University of British Columbia, 1998

PROJECT SUBMITTED IN PARTIAL FULFILLMENT OF  
THE REQUIREMENTS FOR THE DEGREE OF  
MASTER OF ENGINEERING  
IN THE SCHOOL OF ENGINEERING SCIENCE

© Hubert Pun 2001

SIMON FRASER UNIVERSITY

December 2001

All rights reserved. This work may not be  
reproduced in whole or in part, by photocopy  
or other means, without permission of the author.

# Approval

Name: Hubert Pun  
Degree: M.Eng.  
Title of project: Convergence Behavior of RIP and OSPF Network Protocols

Examining Committee:

Chair: Dr. R. Hobson

---

Dr. L. Trajkovic  
Senior Supervisor

---

Dr. W. Gruver  
Professor  
School of Engineering Science

Date approved: December 19, 2001

## Abstract

Routing is the heartbeat of the Internet. Several routing protocols exist nowadays but the most common ones are Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). The prime objectives of this project are to investigate the consequences of deploying RIP and OSPF simultaneously on a network and the performance improved by changing the timers of RIP.

This project will introduce the characteristics of IP addressing. The similarities and differences between Variable Length Subnet Mask (VLSM) and Classless Inter-Domain Routing (CIDR) will be reviewed. Moreover, the advantages of the classless, over the classful, nature of a routing protocol will be rationalized as well. The compositions of a routing table will also be discussed. The section will end with a detailed examination of RIP and OSPF.

Experiments involving seven Cisco routers will be performed. The three cases of interests are:

- impact of a failure Ethernet link to the OSPF convergence
- impact of a broken Frame Relay (FR) Virtual Circuit (VC) to the RIP convergence
- impact of a broken FR VC to the redistribution convergence

The RIP's timers will be changed in each of the three cases to inspect any performance improvement.

# Acronyms

<b>Abbreviation</b>	<b>Meaning</b>
AD	<b>Administrative Distance</b> The “trustworthy” of the routing protocol
BGP	<b>Border Gateway Protocol</b> The de facto routing protocol between Autonomous System
CIDR	<b>Classless Inter-Domain Routing</b> Ability to perform summarization beyond the IP address’s classful boundary
EIGRP	<b>Enhanced Interior Routing Protocol</b> A Cisco proprietary routing protocol. Enhancement of the IGRP
FR	<b>Frame Relay</b>
IETF	<b>Internet Engineering Task Force</b>
IGP	<b>Interior Gateway Protocol</b> Routing protocol that runs inside an Autonomous System
IGRP	<b>Interior Gateway Routing Protocol</b> A Cisco proprietary routing protocol that is designed to replace RIP
IP	<b>Internet Protocol</b> A routed protocol at OSI layer 3
IS-IS	<b>Intermediate System to Intermediate System</b> A link state routing protocol for the ISO’s Connectionless Network Protocol
ISP	<b>Internet Service Provider</b>
LAN	<b>Local Area Network</b>
LS	<b>Link State</b> This is a type of routing protocol that floods the link state throughout the entire area

OSPF	<b>Open Shortest Path First</b> A link state routing protocol that is implemented in most of the networks
RIP	<b>Routing Information Protocol</b> A distance vector routing protocol developed for TCP/IP
VC	<b>Virtual Circuit</b>
VLSM	<b>Variable Length Subnet Mask</b> The ability to support different subnet mask length
WAN	<b>Wide Area Network</b>

# Table of Contents

<b>1. Introduction</b>	1
<b>2. IP Addressing</b>	2
2.1 IP Address and Subnet Mask	2
2.2 VLSM and CIDR	5
2.3 Classful vs. Classless	6
<b>3. RIP and OSPF</b>	8
3.1 Routing Information Protocol (RIP)	8
3.2 Open Shortest Path First (OSPF)	10
3.3 Enhanced Interior Gateway Routing Protocol (EIGRP)	12
<b>4. Routing Tables</b>	14
4.1 Administrative Distance	14
4.2 Route Selection Process	15
4.3 Route-Selection-Process Example	16
<b>5. Experiments</b>	18
5.1 Impact of Hub Link on the OSPF Convergence	20
5.2 Impact of FR Cloud on the RIP Convergence	22
5.3 Impact of FR Cloud on the Redistribution Convergence	24
5.4 Changing RIP Timers	27
5.5 Discussion	33
<b>6. Conclusions</b>	34
<b>7. References</b>	35
<b>Appendix A – Code Listing</b>	36
<b>Appendix B – Routing Tables</b>	45

## List of Figures

**Figure 1.:** Experiment Setup - seven routers (R1, R2, R3, R4, R5, R6, R7), two routing protocols (OSPF, RIP), two ISPs (SIP#1, ISP#2). 19

## List of Tables

<b>Table 1.:</b> Network and host octet for difference classes of IP address.	3
<b>Table 2.:</b> Example of an IP address and subnet mask.	4
<b>Table 3.:</b> Example of IP address summarization.	6
<b>Table 4.:</b> EIGRP metric's bandwidth and delay value.	13
<b>Table 5.:</b> Administrative distance of different routing protocols.	15
<b>Table 6.:</b> Example of route selection process.	17
<b>Table7.:</b> Router specifications for the experiment setup.	18
<b>Table 8.:</b> Summary of the ten experimental results for OSPF convergence. Total 100 test ping packets. On average, 93 packets are received, 7 packets are lost. The convergence time is 14 seconds.	22
<b>Table 9.:</b> Summary of the ten experimental results for RIP convergence. Total 100 test ping packets. On average, 47.2 packets are received, 52.8 packets are lost. The convergence time is 105.6 seconds.	24
<b>Table 10.:</b> Summary of the ten experimental results for redistribution convergence. Total 1000 test ping packets. On average, 764 packets are received, 236 packets are lost. The convergence time is 472 seconds.	26
<b>Table 11.:</b> Summary of the ten experimental results for OSPF convergence with 5 sec update timer, 10 sec invalid timer, 10 sec holddown timer, 30 sec flush timer. Total 100 test ping packets. On average, 94.1 packets are received, 5.9 packets are lost. The convergence time is 11.8 seconds.	28
<b>Table 12.:</b> Summary of the ten experimental results for RIP convergence with 5 sec update timer, 10 sec invalid timer, 10 sec holddown timer, 30 sec flush timer. Total 100 test ping packets. On average, 40.1 packets are received, 59.9 packets are lost. The convergence time is 119.8 seconds.	29
<b>Table 13.:</b> Summary of the ten experimental results for redistribution convergence with 5 sec update timer, 10 sec invalid timer, 10 sec holddown timer, 30 sec flush timer. Total 1000 test ping packets. On average, 930.9 packets are received, 69.1 packets are lost. The convergence time is 138.2 seconds.	30
<b>Table 14.:</b> Summary of the ten experimental results for OSPF convergence	



with 60 sec update timer, 360 sec invalid timer, 360 sec holddown timer, 480 sec flush timer. Total 100 test ping packets. On average, 94.2 packets are received, 5.8 packets are lost. The convergence time is 11.6 seconds. 31

**Table 15.:** Summary of the ten experimental results for RIP convergence

with 60 sec update timer, 360 sec invalid timer, 360 sec holddown timer, 480 sec flush timer. Total 100 test ping packets. On average, 39.7 packets are received, 60.3 packets are lost. The convergence time is 120.6 seconds. 32

**Table 16.:** Summary of the ten experimental results for redistribution convergence

with 60 sec update timer, 360 sec invalid timer, 360 sec holddown timer, 480 sec flush timer. Total 1000 test ping packets. On average, 566.1 packets are received, 433.9 packets are lost. The convergence time is 867.8 seconds. 32

# 1. Introduction

The goal of this project is to investigate the behavior of routing convergence. It begins with an explanation of IP addressing. The report includes topics such as Variable Length Subnet Mask (VLSM), Classless Inter-Domain Routing (CIDR) and classful versus classless. Next, the report discusses the two routing protocols: Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) into great detail. The report then examines the structure of a routing table and the route selection process.

In order to be practical in the investigation of the routing convergence, we perform an experiment that involved seven Cisco routers. It is assumed that an end customer requires redundancy for its Wide Area Network (WAN) connection. The customer purchases WAN connectivity from two different ISPs that are, unfortunately, running two different routing protocols; hence, routing information must be redistributed. We conduct the experiment such that network convergences under different failure situation are examined. We will also modify the timers of RIP to inspect any improvement.

The Appendix contains the router codes that the author wrote and the routing table that was generated as a result of this project.

## **2. IP Addressing**

In order to understand routing protocol, one must have a deep understanding of IP addressing. Hence, we include a brief discussion of the IP addressing scheme. Next, we cover the concepts of Variable Length Subnet Mask (VLSM) and Classless Inter-Domain Routing (CIDR). They are techniques for making IP addressing more efficient. They are similar, yet with a subtle difference. Finally, we discuss the bases of classless and classful behavior of a routing protocol.

### **2.1 *IP Address and Subnet Mask***

The IP addressing space in North America is administered by the America Registry for Internet Number (ARIN). An IP address is 32 bits in length with two parts: network number and host number. The length of the network number is different for different classes [11].

#### **2.1.1 IP Address Classes**

IP address is defined in five classes as shown in Table 1. They differ in the number of hosts that can be attached to the network. The network number is assigned by ARIN while the host address is chosen by the network administrator.

When referring to the network address, the typical nomenclature is to put a “0” in the host address locations, e.g., 10.0.0.0. When referring to the host address, the convention is to use the complete address as the host address, e.g., 10.12.42.123.

	31 – 24 bit	23 – 16 bit	15 – 8 bit	7 – 0 bit
<b>Class A</b>	Network number	Host number	Host number	Host number
<b>Class B</b>	Network number	Network number	Host number	Host number
<b>Class C</b>	Network number	Network number	Network number	Host number
<b>Class D</b>	Reserved for Multicast			
<b>Class E</b>	Reserved for experimental purposes			

**Table 1.: Network and host octet for difference classes of IP address.**

Private address allows the users to create their own networking address schemes; these addresses must not interface to the public Internet directly. Their ranges are 10.0.0.0 – 10.255.255.255, 172.16.0.0 – 172.32.255.255, and 192.168.0.0 – 192.168.255.255.

### 2.1.2 Subnet

Subnet is a powerful concept that extends the network number one step further. Let's say a network administrator is given a class of IP address block. It is required to divide the hosts into different networks in order to separate the traffic streams. By using the concept of subnet, the network administrator can decide on the size of the subnet block according to the needs.

A subnet mask is used to identify the subnet boundary. It uses binary ones to denote the network and subnet bits, and binary zeros to denote the host bits. For the host address 172.16.2.4 with a subnet of 172.16.2.0, the subnet mask is:

11111111.11111111.11111111.00000000 (or 255.255.255.0 in decimal)

	Binary representation				=	Decimal equivalent
	31–24 bit	23–16 bit	15–8 bit	7–0 bit		
<b>IP address:</b>	10101100	00010000	00000010	00000100	=	172.16.2.4
<b>Subnet address:</b>	10101100	00010000	00000010	00000000	=	172.16.2.0
<b>Subnet mask:</b>	11111111	11111111	11111111	00000000	=	255.255.255.0

**Table 2.: Example of an IP address and subnet mask.**

A short form is used to show all the information in a more concise style. The number of ones in the subnet mask is written after the IP address, proceeding with a slash. The above example has 24 ones in the subnet mask. Consequently, the IP address and its subnet mask can be written as **172.6.2.4/24**. This can be referred to as a “bit mask”.

## **2.2 VLSM and CIDR**

There is a shortage of IP address. The main reason is its pre-set subnet mask. Variable Length Subnet Mask (VLSM) is designed to solve this problem. Moreover, an issue that routers are facing is the large size of routing table – up to 90,000 entries. Classless Inter-Domain Routing (CIDR) is a methodology that would summarize entries within the routing table. Both are similar in a way that they modify the predefined subnet mask. However, there exists a subtle difference: VLSM divides the standard class into smaller subnets while CIDR summarizes several subnets into an aggregated entry.

### 2.2.1 Variable Length Subnet Mask (VLSM)

Suppose that it is required to divide a class C network into three subnets: one consists of 100 hosts and two consists of 50 hosts each. Even though a “class C” network has 254 host addresses available, this cannot be done by simply divide the address space into two 127-host networks or four 63-host networks. The only solution is to split the entire address space into two big blocks, each with 127 host addresses, and further divide one of the blocks into two smaller blocks, each with 63 host addresses.

This method of dividing the IP address block into different sizes is called Variable Length Subnet Mask (VLSM). It is flexible and can suite the different requirement. It also reduces the waste of IP address [8].

### 2.2.2 Classless Inter-Domain Routing (CIDR)

The Internet routing table currently consists of more than 90,000 entries. In order to summarize some of the redundant information, Classless Inter-Domain Routing (CIDR) is acquainted. Consider the example in Table 3 of summarizing networks from 192.168.8.0/24 to 192.168.15.0/24.

	Binary representation				Decimal equivalent
	31–24 bit	23–16 bit	15–8 bit	7–0 bit	
<b>First IP Address:</b>	11000000	10101000	00001000	00000000	= 192.168.8.0
<b>Last IP Address:</b>	11000000	10101000	00010000	11111111	= 192.168.16.255
<b>Summarized:</b>	11000000	10101000	000xxxxx	xxxxxxxx	= 192.168.0.0/19

**Table 3.: Example of IP address summarization.**

Normally, the routing table would have eight entries. By deploying CIDR, these entries can be summarized as 192.168.0.0/19. However, one must be aware that no “hole” in the address space is allowed; otherwise, black hole will result in packet loss.

### **2.3 Classful vs. Classless**

The classful/classless nature of a routing protocol indicates whether or not the concept of subnet is allowed. If a routing protocol is classful, it automatically assumes that no subnet exists. For example, only the standard network address “10.0.0.0” is passed for the routing entry “10.1.1.0/24”. No subnet or subnet mask are transmitted. Then, when another router receives this routing entry, it uses the normal mask, namely “/8” or “255.0.0.0”. The information of subnet is lost.

On the other hand, if the routing protocol is classless, the routing entry will consist both the network and the subnet mask. In the above example, the routing entry includes both the network address 10.1.1.0 and the subnet mask 255.255.255.0 pair. This contains the complete information.

## **3. RIP and OSPF**

In today's commercial networks, Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) are the most widely used routing protocols. In this section, we examine both RIP and OSPF. We will also discuss another routing protocol, Enhanced Interior Gateway Routing Protocol (EIGRP) briefly for comparison purpose.

### **3.1 Routing Information Protocol (RIP)**

Routing Information Protocol (RIP) is one of the first widely deployed routing protocols. It uses a distance vector algorithm. It is simple to program, but has a number of disadvantages.

#### **3.1.1 Algorithm**

Routers pass periodic copies of their routing table to neighboring routers and accumulate cost. RIP uses hop count as the metric for each link. For example, consider three adjacent routers, A, B and C connected in a straight line. Router A passes its routing table to Router B; Router B adds one to the metric and passes the routing table to its other neighbor, Router C. The same step-by-step process occurs in all directions between direct-neighbor routers [7].

#### **3.1.2 Topology Change**

The routing table must be updated whenever the inter-network topology changes. A table update requires each router to send its routing table to each of the adjacent neighbors.

When a router receives an update, it compares the update with its routing table. It adds



the metric of reaching the neighbor router to the path metric reported by the neighbor to establish a new metric.

### **3.1.3 Problems and Solutions**

There are a number of issues relevant to RIP. First, the slow convergence may cause inconsistent routing entries, occasionally results in routing loops. When there is a link failure, other routers cannot receive the failure notification before sending their own updates. Consequently, the network bounces the incorrect routing table and increments the metric. The metric can eventually approach to infinity.

In order to correct this problem, combinations of solutions have been implemented [5]. By defining 15 to be the maximum number of hops, the infinite looping problem can be prevented. A second solution uses “split horizon”, which forbids the router from sending information about a route back in the direction from which the original packet arrived. Moreover, a hold-down timer can be used. It instructs the router to delay any changes that involves the defected routes. Finally, the router can send messages as soon as it notices a change in their routing table (triggered update).

### **3.1.4 Disadvantages**

There are several disadvantages to RIP. The network is restricted to the size of 15 hops due to the solution to the “count to infinity” problem. In addition, the periodic broadcast of the routing table consumes bandwidth. The convergence is slow too.

## **3.2 Open Shortest Path First (OSPF)**

Open Shortest Path First (OSPF) was developed by the Internet Engineering Task Force (IETF) as a replacement of the problematic RIP in RFC 2328 [9]. This is a non-proprietary routing protocol for the TCP/IP protocol family with many advantages over RIP [2].

### **3.2.1 Algorithm**

OSPF generates link-state packets that contain local information for each router. Each router exchanges local and external link state information and generates a shortest path tree. Each router uses this exact topology to calculate the shortest path to each destination. Recalculation occurs only if there are any changes.

### **3.2.2 Topology Changes**

Each router keeps track of the link states of its neighbors. Whenever there is a change, router notifies other routers by sending a link-state packet. Other routers then reconstruct a complete map of the inter-network.

### **3.2.3 Problems and Solutions**

Unsynchronized updates and inconsistent path decisions are the main problems of OSPF. Routers cannot determine the most recent update when two different link-state updates arrive at approximately the same time. If the link-state packet is not correctly distributed to all routers, invalid routing entries will be resulted. This problem is relatively minor when comparing to the problem encountered by RIP. This can be solved easily by

coordinating the updates. Time stamps, update numbering and counters can be used to show the sequence of the update [4].

### **3.2.4 Advantages and Disadvantages**

OSPF has both advantages and disadvantages. Some advantages of OSPF are:

- It is the highest-performance open standard routing protocol.
- It is a classless routing protocol.
- It provides shortest path routing and is fast to fault-discovery and rerouting.
- It consumes minimal link overhead when the network is in steady state.
- It has been endorsed by the IETF and implemented by many vendors.

Some disadvantages of OSPF are:

- It demands a higher processing and memory requirement than RIP.
- It consumes a large bandwidth at the initial link-state packet flooding.

### **3.3 *Enhanced Interior Gateway Routing Protocol (EIGRP)***

EIGRP is a third generation distance vector routing protocol that negotiates neighbor relation like a link state routing protocol. It combines the advantages of both the distance vector and link state routing protocol [1,10]. The calculation algorithm that EIGRP deploys is called the Diffusing Update Algorithm (DUAL) which is proprietary to Cisco.

The metric value is formulated by:

$$metric = [k_1 \times BW_{EIGRP(\min)} + \frac{k_2 \times BW_{EIGRP(\min)}}{256 - LOAD} + k_3 \times DLY_{EIGRP(\text{sum})}] \times \frac{k_5}{RELI + k_4} \times 256$$

where:

$BW_{EIGRP(\min)} = 10^7$  / minimum BW along the path(in kbps)

$DLY_{EIGRP(\text{sum})} =$  Total Delay along the path (in  $\mu\text{s}$ ) / 10

LOAD = how the link is loaded (out of 255)

RELI = how reliable the link is (out of 255)

By default, the k-values are:

$$\begin{bmatrix} k_1 \\ k_2 \\ k_3 \\ k_4 \\ k_5 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \\ RELI \end{bmatrix}$$

This simplifies the formula into:

$$metric = (BW_{EIGRP(\min)} + DLY_{EIGRP(\text{sum})}) \times 256$$

Table 4 lists the bandwidth and delay of different media and how EIGRP algorithm

interprets the  $BW_{EIGRP}$  and  $DLY_{EIGRP}$  values:

<b>Media</b>	<b>BW</b>	<b><math>BW_{EIGRP}</math></b>	<b>Delay</b>	<b><math>DLY_{EIGRP}</math></b>
<i>Fast Ethernet</i>	100,000k	<b>100</b>	100 $\mu$ s	<b>10</b>
<i>FDDI</i>	100,000k	<b>100</b>	100 $\mu$ s	<b>10</b>
<i>Ethernet</i>	10,000k	<b>1000</b>	1000 $\mu$ s	<b>100</b>
<i>T1</i>	1544k	<b>6476</b>	20000 $\mu$ s	<b>2000</b>
<i>DS0</i>	64k	<b>156250</b>	20000 $\mu$ s	<b>2000</b>
<i>56k</i>	56k	<b>178571</b>	20000 $\mu$ s	<b>2000</b>

**Table 4.: EIGRP metric's bandwidth and delay value.**

### 3.3.1 Advantages and Disadvantages

There are some outstanding advantages using EIGRP. First, it supports multi-network layer routed protocols, namely, IP, Inter-network Packet Exchange (IPX) and AppleTalk (AT). This is a huge advantage for the non-TCP/IP oriented networks. In addition, the convergence time for EIGRP is very fast.

On the other hand, the drawback of EIGRP is its proprietary nature. Network manager hesitates to commit to a pure Cisco environment; any network with one non-Cisco router would not be able to deploy EIGRP.

## 4. Routing Tables

Routers exchange updates according to the specific protocol to locate the most efficient route. Then a routing table is used to determine which next-hop (route) to use in order to send a packet to a specific destination. In this section, we discuss the concept of administrative distance and explain the route selection process.

### 4.1 Administrative Distance

Administrative distance (AD) is the rating of the trustworthiness of a routing protocol which is expressed as an integer between 0 and 255. The lower the value, the more trustworthy the information is. Table 5 lists the default ADs that are implemented in a Cisco router.

Route Source	AD
Static entry	1
Internal EIGRP	90
OSPF	110
RIP	120
Unknown source	255

Table 5.: Administrative distance of different routing protocols.

### 4.2 Route Selection Process

There are three steps in order to determine which next hop (route) to use for a certain destination. First, the longest-prefix rule is used. This rule states that the most precise

entry should be used. For example consider a routing table with two entries, next hop of 1.1.1 for destination network 10.1.1.0 and subnet mask of 255.255.255.0, and next hop of 2.2.2 for destination network 10.1.1.0 and subnet mask of 255.255.255.128. The first entry specifies IP addresses ranging from 10.1.1.0 to 10.1.1.255 while the second entry only specifies from 10.1.1.0 to 10.1.1.127. The second routing entry is more concise. Consequently, if a packet has a destination address of 10.1.1.1, it would use 2.2.2.2 as the next hop.

The next step is to check the entry's AD value. A lower numerical AD source of information is favored. For example, routing entries are sourced from both RIP and OSPF. The entry learnt from OSPF would be preferred because OSPF has a lower AD than RIP (110 vs. 120).

The final step considers the metric value. In this case, a lower metric is preferred. If there is no unique decision resulted after these steps, the traffic will be load-balanced [3].

### ***4.3 Route-Selection-Process Examples***

Assume the several sources of routing entries of a particular router are listed in Table 6. Four examples will be used to illustrate the procedure of how this router chooses the next hop (route).

	Source	Route	AD	Metric
1	OSPF	10.1.1.0/24	110	20
2	RIP	10.0.0.0/8	120	4
3	EIGRP	10.1.1.0/24	90	185324
4	OSPF	10.1.1.1/32	110	120
5	EIGRP	10.1.1.0/24	90	512

**Table 6.: Example of route selection process.**

- If the destination is 10.1.1.1, it would use entry #4 because of the longest-prefix rule.
- If the destination is 10.1.1.2, it would use entry #5 because EIGRP (entry #5) has the lower AD than OSPF (entry #4) and metric “512” (entry #5) is smaller than metric “185324” (entry #3).
- If the destination is 10.2.3.4, it would use entry #2 because it is the only entry that consists of the IP address 10.2.3.4.
- If the destination is 172.16.2.3, the packet will be dropped because no entry includes this IP address.



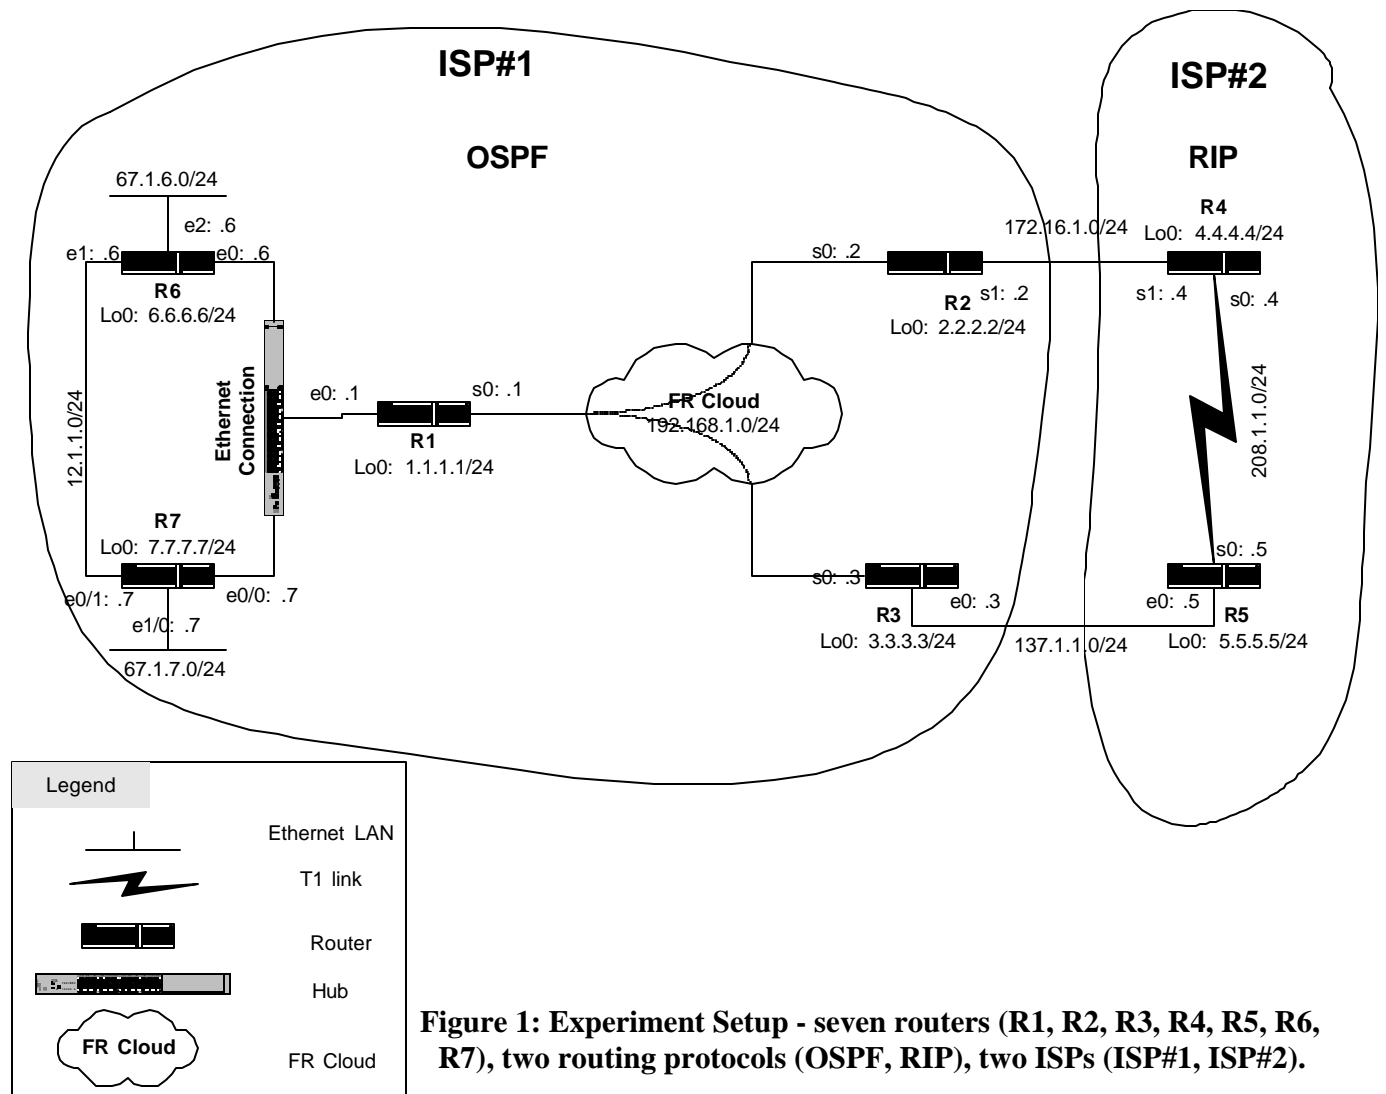
## 5. Experiments

The experiment consists of eight Cisco routers. The specifications of the routers are listed in Table 7. A Cisco 2522-DC router, **FR**, is configured as a Frame Relay switch by using the command “frame-relay switching” to simulate a Frame Relay cloud. This capability is usually not used in production network. Rather, it is mainly intended for use in test beds and experimental networks.

Name	IOS Version	RAM	Model	Serial Number
<b>R1</b>	igs-d-l.111-12	8192k	2503	25191623
<b>R2</b>	C2500-ds-l_113-3_T.bin	2048k	2513	250432533
<b>R3</b>	C2500-ds-l_113-3_T.bin	2048k	2524	25795218
<b>R4</b>	C2500-ds-l_113-3_T.bin	4096k	2515	25201810
<b>R5</b>	C1600-nr2y-l.112-10a.P	1536k	1601	JAB033530WD
<b>R6</b>	C4500-js-mz_112-23.bin	16384k	4000	45575012
<b>R7</b>	C3640-js56i-mz_120-10.bin	24576k	3640	JAB040180KW
<b>FR</b>	C2500-ds-l_113-3_T.bin	2048k	2522-DC	250330094

**Table7.: Router specifications for the experiment setup.**

The setup of the experiment is shown in Figure 1. It involves two Internet Service Providers (ISP) with seven routers. The first Internet service provider, **ISP#1**, runs OSPF as the routing protocol while the other Internet service provider, **ISP#2**, runs RIP. **R2** and **R3** are the border routers that interface with ISP#1 and ISP#2, and redistribute the routing information between the two domains.



**Figure 1: Experiment Setup - seven routers (R1, R2, R3, R4, R5, R6, R7), two routing protocols (OSPF, RIP), two ISPs (ISP#1, ISP#2).**

The end customer has purchased Ethernet and Frame Relay (FR) connectivity from **ISP#1**. **R1** is the main site. **R6** and **R7** are the two Ethernet sites, and a redundant connection exists between these two sites. **R2** and **R3** are the two FR sites with 128k and 256k FR virtual circuit (VC) connected to the main sites respectively. This customer requires WAN redundancy for the FR site. Consequently, a T1 link is purchased from **ISP#2** between the two sites. **R4** is located at the same physical location as **R2**, and **R5** is

located in the same physical location as **R3**. **R2** and **R4** are interconnected with a T1 serial link while **R3** and **R5** are interconnected with an Ethernet link.

We investigate three convergence behaviors after link failure. First, we loose the hub link and examine the convergence time. Next, we remove the FR VC between **R1** and **R2** and measure the convergence time. In addition, we investigate the RIP/OSPF convergence behavior while the FR VC remains removed. Finally, we modify the RIP timers in order to examine any improvement in the convergence time.

## ***5.1 Impact of Hub Link on the OSPF Convergence***

The normal path for traffic flow from **R4** to **R7** is:

**R4 → R5 → R3 → R1 → R7**

In order to verify this traffic pattern, the command “traceroute” is issued at **R4**’s command line interface; we trace the route to **R7**’s Ethernet link (67.1.7.7). There are four hops; each represented by a line. After stating the line number, the next hop’s IP address is indicated. For example, the second hop is at **R3** and the IP address 137.1.1.3.

Three trace route packets are sent between each hop; the traveling times for these three packets are recorded in the last three numbers. In the case of line 2, the three numbers, 20ms, 24ms and 24ms, indicate the traveling time for the three trace-route packets from the first hop **R5** (208.1.1.5) to the second hop **R3** (137.1.1.3). If the traveling time is a “\*”, it means that the packet never arrives.

```
R4# traceroute 67.1.7.7
```

```
Type escape sequence to abort.
```

```
Tracing the route to 67.1.7.7
```

```
 1 208.1.1.5 32 msec 28 msec 24 msec          ← R5
 2 137.1.1.3 20 msec 24 msec 24 msec          ← R3
 3 192.168.1.1 48 msec 64 msec 60 msec         ← R1
 4 10.1.1.7 56 msec * 56 msec                 ← R7
```

We disconnect **R7**'s Ethernet interface module 0, slot 0 (e0/0) to simulate a broken

Ethernet link. The traffic is shifted to the backup Ethernet connection between **R6** and

**R7**. The resulting route according to the "traceroute" command, is:

**R4 → R5 → R3 → R1 → R6 → R7**

```
R4# traceroute 67.1.7.7
```

```
Type escape sequence to abort.
```

```
Tracing the route to 67.1.7.7
```

```
 1 208.1.1.5 28 msec 28 msec 20 msec          ← R5
 2 137.1.1.3 20 msec 24 msec 24 msec          ← R3
 3 192.168.1.1 56 msec 60 msec 60 msec         ← R1
 4 10.1.1.6 56 msec 60 msec 56 msec           ← R6
 5 12.1.1.7 56 msec * 56 msec                 ← R7
```

One hundred ping packets, each with a timeout of two seconds, are transmitted from **R4**

by using the command "ping". **R7**'s e0/0 is disconnected sometime during the "ping"

command is issued. The destination IP address is 67.1.7.7 (**R7**); the number of ping

packets is 100. The default values are chosen for the protocol, datagram size, timeout,

extended commands and sweep range of sizes. Each successful ping packet is

represented by a "!" and each failure ping packet is represented by a ".". In this

particular ping command, 94% of the packets successfully arrive to the destination. The

minimum, average and maximum ping-packet-round-trip times are 88ms, 95ms and 324ms respectively.

```

R4# ping
Protocol [ip]:
Target IP address: 67.1.7.7                ← R7
Repeat count [5]: 100                       ← 100 packets
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 94 percent (94/100), round-trip min/avg/max = 88/95/324 ms

```

The first 23 ping packets travel the normal path between **R4** and **R7**. Then **R7**'s e0/0 is disconnected. During the transient time when the routing protocol is converging, six ping packets are lost. After the convergence of OSPF, the last 71 ping packets travel the backup path between **R4** and **R7**. This ping test is repeated for ten times, and the results are listed in Table 8. On average, there are seven missing packets during the transient period; it indicates that OSPF requires 14 seconds to converge in this topology.

	Packet received	Traffic loss	Convergence time
	92	8	16
	92	8	16
	92	8	16
	92	8	16
	92	8	16
	93	7	14
	93	7	14
	94	6	12
	94	6	12
	96	4	8
<b>Average</b>	<b>93</b>	<b>7</b>	<b>14</b>

**Table 8.: Summary of the ten experimental results for OSPF convergence.**

**Total 100 test ping packets. On average, 93 packets are received, 7 packets are lost.**

**The convergence time is 14 seconds.**

## **5.2 Impact of FR Cloud on the RIP Convergence**

The normal path for the traffic flow from **R7** to **R4** is:

**R7 → R1 → R2 → R4**

It can be verified by issuing the command “traceroute 4.4.4.4” at **R7**.

```
R7# traceroute 4.4.4.4
```

```
Type escape sequence to abort.
```

```
Tracing the route to r4 (4.4.4.4)
```

```
 1 10.1.1.1 0 msec 4 msec 8 msec          ← R1
 2 192.168.1.2 32 msec 36 msec 36 msec     ← R2
 3 172.16.1.4 48 msec * 48 msec          ← R4
```

We disconnect **R2**'s serial interface slot 0 (s0) to simulate a broken FR link. It can represent a loss of VC or a loosen connection. The traffic ends up using the other FR link between **R1** and **R3**, and ISP#2's T1 link.

**R7 → R1 → R3 → R5 → R4**

```
R7# traceroute 4.4.4.4
```

```
Type escape sequence to abort.
```

```
Tracing the route to r4 (4.4.4.4)
```

```
 1 10.1.1.1 4 msec 8 msec 4 msec          ← R1
 2 192.168.1.3 32 msec 36 msec 32 msec     ← R3
 3 137.1.1.5 36 msec 32 msec 32 msec     ← R5
 4 208.1.1.4 52 msec * 48 msec          ← R4
```

Again, one hundred ping packets, each with a timeout period of two seconds, are sent from **R7**. The settings are similar to the ones from Section 5.1. A sample ping command output is:

```
R7# ping
Protocol [ip]:
Target IP address: 4.4.4.4                ← R4
Repeat count [5]: 100                     ← 100 packets
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
.....!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 45 percent (45/100), round-trip min/avg/max = 84/86/92 ms
```

We repeat the ping test for ten times, and the results are listed in Table 9. On average, there are 47.2 missing packets during the transient period; it indicates that, RIP requires 105.6 seconds to converge in this topology. It is considerably longer than OSPF's convergence time of 14 seconds in Section 5.1.

	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence time</b>
	43	57	114
	45	55	110
	47	53	106
	47	53	106
	47	53	106
	47	53	106
	48	52	104
	49	51	102
	49	51	102
	50	50	100
<b>Average</b>	<b>47.2</b>	<b>52.8</b>	<b>105.6</b>

**Table 9.: Summary of the ten experimental results for RIP convergence. Total 100 test ping packets. On average, 47.2 packets are received, 52.8 packets are lost. The convergence time is 105.6 seconds.**

### 5.3 Impact of FR Cloud on the Redistribution Convergence

Normally, traffic would flow a direct path from **R7** to **R2**'s loopback link (2.2.2.2):

**R7 → R1 → R2**

It can be verified by issuing the command “traceroute 2.2.2.2” at **R7**.

```
R7# traceroute 2.2.2.2
```

```
Type escape sequence to abort.
```

```
Tracing the route to r2 (2.2.2.2)
```

```
 1 10.1.1.1 4 msec 8 msec 4 msec          ← R1
 2 192.168.1.2 36 msec * 32 msec         ← R2
```

Similar to Section 5.2, we disconnect **R2**'s serial interface slot 0 (s0) to simulate a broken

FR link. Traffic can still be transmitted between the two sites with an indirect route:

**R7 → R1 → R3 → R5 → R4 → R2**

```
R7# traceroute 2.2.2.2
```

```
Type escape sequence to abort.
```

```
Tracing the route to r2 (2.2.2.2)
```

```
 1 10.1.1.1 0 msec 4 msec 8 msec          ← R1
 2 192.168.1.3 36 msec 36 msec 32 msec     ← R3
 3 137.1.1.5 32 msec 36 msec 36 msec      ← R5
 4 208.1.1.4 48 msec 52 msec 48 msec      ← R4
 5 172.16.1.2 64 msec * 64 msec          ← R2
```

This time, we send 1000 ping packets from **R7**. Other settings remain the same.





	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence time</b>
	759	241	482
	759	241	482
	760	240	480
	761	239	478
	762	238	476
	765	235	470
	767	233	466
	767	233	466
	769	231	462
	771	229	458
<b>Average</b>	<b>764</b>	<b>236</b>	<b>472</b>

**Table 10.: Summary of the ten experimental results for redistribution convergence. Total 1000 test ping packets. On average, 764 packets are received, 236 packets are lost. The convergence time is 472 seconds.**

However, there is a stability problem in the RIP domain. In order to achieve convergence in the RIP domain, the time specified in the flush timer must be waited in order to clear the RIP's routing table. In the transient stage, **R5** still uses **R3** to reach to **R2**; it advertises this false information to other routers, including **R3**. **R3** then mislead the OSPF domain with a false path; saying that **R5** has a way (which has next hop of **R3**) to reach **R2**. This is a routing loop. The metric of RIP continues to rise until it hits the predefined infinity, which is 16. Then **R5** signals to routers in RIP domain that **R2** cannot be reached by using route **R3**. At last the correct routing information is sent and the convergence begins [6]. We show in the next section that this result is extremely sensitive to the variation of the RIP timers.

## 5.4 Changing of RIP Timers

There are four timers in RIP. The **Update timer** dictates the time interval between successive sending of the routing table. Each time when an entry is updated, the **invalid timer** is reset. If a route is not received in an update for the time specified by the invalid timer, it is declared unusable and will trigger the **holddown timer**: this entry will still be used to route packets; however it will not be announced in the routing updates. Finally, the route will be removed after the **flush timer** expired. The default settings of the four timers are 30s (update timer), 180s (invalid timer), 180s (holddown timer) and 240s (flush timer).

We repeat the tests performed in Section 5.1 to 5.3 in this section twice. The procedures and settings are identical to the ones before. The first set of tests uses a shorter timer values and the second set of tests uses a long timer values. We examine the effect of the convergence based on these tests.

### 5.4.1 Timer Basic 5 10 10 30

The RIP timers are decreased in this sub-section. It means that the update will be sent more frequently and routing table would converge faster. On the other hand, there are more overhead routing traffics sent between sites. This setting is only suitable for a small network. It is because for a large network, the routing entry may be flushed away before the update arrives.

#### 5.4.1.1 Impact of Hub Link on the OSPF Convergence

The nature and settings of this ping test are similar to the ones in Section 5.1. The ping command is repeated for ten times with 100 ping packets each. The results are listed in Table 11. On average, 5.9 packets are lost during the transient period. The network topology requires 11.8 seconds to converge. This is roughly equal to the result obtained by the default timers setting in Section 5.1 (14s convergence time). It is reasonable because changing RIP's timers should have minimal effect to the OSPF convergence.

	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence Time</b>
	93	7	14
	93	7	14
	93	7	14
	93	7	14
	93	7	14
	95	5	10
	95	5	10
	95	5	10
	95	5	10
	96	4	8
<b>Average</b>	<b>94.1</b>	<b>5.9</b>	<b>11.8</b>

**Table 11.: Summary of the ten experimental results for OSPF convergence with 5 sec update timer, 10 sec invalid timer, 10 sec holddown timer, 30 sec flush timer. Total 100 test ping packets. On average, 94.1 packets are received, 5.9 packets are lost. The convergence time is 11.8 seconds.**

#### 5.4.1.2 Impact of FR Cloud on the RIP Convergence

RIP has a mechanism of “triggered update”. Any changes occurred would be detected immediately; the updated routing table would be redistributed to OSPF simultaneously too. Because of this triggered update mechanism, the results obtained from this subsection should be similar to the ones from Section 5.2. In this subsection, the settings of the test are identical to the ones in Section 5.2 except for the change in the RIP timers.

The ping command is repeated for ten times with 100 ping packets each. According to Table 12, there are 119.8s traffic loss during the convergence. This is approximately equal to the default case in Section 5.2 (105.6s).

	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence time</b>
	37	63	126
	38	62	124
	38	62	124
	38	62	124
	39	61	122
	39	61	122
	40	60	120
	41	59	118
	45	55	110
	46	54	108
<b>Average</b>	<b>40.1</b>	<b>59.9</b>	<b>119.8</b>

**Table 12.: Summary of the ten experimental results for RIP convergence with 5 sec update timer, 10 sec invalid timer, 10 sec holddown timer, 30 sec flush timer. Total 100 test ping packets. On average, 40.1 packets are received, 59.9 packets are lost. The convergence time is 119.8 seconds.**

#### 5.4.1.3 Impact of FR Cloud on the Redistribution Convergence

As discussed in the previous section, the redistribution convergence time is very sensitive to the RIP timer. By reducing the timer values, the test is repeated for ten times with 1000 ping packets each. The settings of the ping command are identical to the ones in Section 5.3. According to Table 13, there are only 69.1 missing packets during the convergence. Hence, the convergence time improves from a 472s traffic loss in Section 5.3 to only 138.2s traffic loss. This advancement suggests that for any dynamic network with acceptable bandwidth between sites, the RIP timer should be set to smaller values.

	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence time</b>
	925	75	150
	927	73	146
	927	73	146
	927	73	146
	928	72	144
	930	70	140
	935	65	130
	936	64	128
	937	63	126
	937	63	126
<b>Average</b>	<b>930.9</b>	<b>69.1</b>	<b>138.2</b>

**Table 13.: Summary of the ten experimental results for redistribution convergence with 5 sec update timer, 10 sec invalid timer, 10 sec holddown timer, 30 sec flush timer. Total 1000 test ping packets. On average, 930.9 packets are received, 69.1 packets are lost. The convergence time is 138.2 seconds.**

#### **5.4.2 Timer Basic 60 360 360 480**

The tests from Section 5.1 to 5.3 are repeated again in this sub-section, except that the RIP timers are doubled. Consequently, updates are sent sparsely. This setting is suitable for either a large network, a static network, or a network with limited bandwidth between sites. Nonetheless, the routing tables converge slower.

##### **5.4.2.1 Impact of Hub Link on the OSPF Convergence**

As in Section 5.1, the ping command is repeated for ten times with 100 ping packets each. The settings of the ping command are identical to the ones from Section 5.1. The results are listed in Table 14. On average, there are 5.8 missing packets during the transient period; i.e., OSPF requires 11.6 seconds for this network topology to converge.

The results obtained from the various timer settings are roughly the same (14s in Section

5.1, 11.8s in Section 5.4.1.1 and 11.6s in Section 5.4.2.1) because changing RIP's timers should have an insignificant impact to the OSPF convergence.

	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence time</b>
	92	8	16
	92	8	16
	93	7	14
	93	7	14
	94	6	12
	94	6	12
	96	4	8
	96	4	8
	96	4	8
	96	4	8
<b>Average</b>	<b>94.2</b>	<b>5.8</b>	<b>11.6</b>

**Table 14.: Summary of the ten experimental results for OSPF convergence with 60 sec update timer, 360 sec invalid timer, 360 sec holddown timer, 480 sec flush timer. Total 100 test ping packets. On average, 94.2 packets are received, 5.8 packets are lost. The convergence time is 11.6 seconds.**

#### 5.4.2.2 Impact of FR Cloud on the RIP Convergence

We repeats the test from Section 5.2 but with a different RIP timers. The results are listed in Table 15. On average, there are 60.3 missing packets during the transient period; i.e., RIP requires 120.6 seconds for this network topology to converge. This result is more or less the same as the ones obtained from the previous two RIP timer settings (105.6s and 119.8s).

	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence time</b>
	38	62	124
	38	62	124
	38	62	124
	39	61	122
	39	61	122
	39	61	122
	41	59	118
	41	59	118
	42	58	116
	42	58	116
<b>Average</b>	<b>39.7</b>	<b>60.3</b>	<b>120.6</b>

**Table 15.: Summary of the ten experimental results for RIP convergence with 60 sec update timer, 360 sec invalid timer, 360 sec holddown timer, 480 sec flush timer. Total 100 test ping packets. On average, 39.7 packets are received, 60.3 packets are lost. The convergence time is 120.6 seconds.**

#### 5.4.2.3 Impact of FR Cloud on the Redistribution Convergence

Increasing the RIP timers should increase the time redistribution required to converge.

This sensitivity can be shown. The ping command is repeated for ten times with 1000 ping packets each. The settings of the ping command are identical to the ones in Section 5.3. According to Table 16, there are now 867.8s of traffic loss, which almost double the 472s traffic loss from Section 5.3. This setup is not recommended unless the bandwidth between links is valuable and the network is stable.



	<b>Packet received</b>	<b>Traffic loss</b>	<b>Convergence time</b>
	545	455	910
	556	444	888
	560	440	880
	565	435	870
	566	434	868
	567	433	866
	572	428	856
	572	428	856
	573	427	854
	585	415	830
<b>Average</b>	<b>566.1</b>	<b>433.9</b>	<b>867.8</b>

**Table 16.: Summary of the ten experimental results for redistribution convergence with 60 sec update timer, 360 sec invalid timer, 360 sec holddown timer, 480 sec flush timer. Total 1000 test ping packets. On average, 566.1 packets are received, 433.9 packets are lost. The convergence time is 867.8 seconds.**

## **5.4 Discussion**

The main goal is to examine the convergence behavior for OSPF, RIP and redistribution between the two. First, we break the hub link to simulate an Ethernet failure. Ten measurements are recorded for the three RIP timers cases. The result shows that the amount of traffic loss during the OSPF convergence is very close in each of the cases (14s in Section 5.1, 11.8s in Section 5.4.1.1, 11.6s in Section 5.4.2.1). This insensitive variation to the RIP timer is expected.

Next, we loose the connection to the R2's serial port to simulate a loss to the FR VC.

Once again, the RIP convergence time is about two minutes for the three different cases (105.6s in Section 5.2, 119.8s in Section 5.4.1.2, 119.6 in Section 5.4.2.2). "Triggered

update”, a standard mechanism of RIP, causes an immediate update when any routing changes is detected. Hence, the result is independent from the RIP timers.

Finally, we examine the convergence behavior of the redistribution between OSPF and RIP. According to the experimental result, the convergence time is sensitive to the setting of the RIP timers. Larger timer values cause a slower convergence, and vice versa. The routing tables converge in 472s for the default timers. Decreasing these values yields a faster convergence time (138.2s). Doubling these values result a slower convergence time (8678.8s). For a company that has a dynamic network, it is recommended to set the timers to smaller values to improve the convergence time.

## 6. Conclusions

Routing tables across the entire network should converge in minimum time in order to avoid excessive traffic loss. This is the main interest of this project. We also discussed IP addressing and the various components of the routing table. Next, we conducted an experiment that examines the behavior of the routing protocols, RIP and OSPF. We investigated the routing convergence under three different situations and concluded that OSPF converges in about 10 seconds when there is a broken Ethernet connection, while RIP converges in a minute when there is a failure FR VC.

We also observed that the RIP timers impact the RIP/OSPF redistribution convergence behavior significantly. We recommend that a shorter RIP timers be programmed whenever possible. Otherwise, the end devices (PC, work station or servers) should be programmed with a longer idle time-out whenever the WAN connection involves a multi-routing protocol because the convergence requires up to several hundred seconds.

## 7. References

[1] H. Benjamin and T. Thomas, *Routing and Switching*. Scottsdale, AZ: Coriolis, 2000, pp. 192-193.

[2] A. Caslow, *Cisco Certification: Bridges, Routers & Switches for CCIEs*. Upper Saddle River, NJ: Prentice Hall PTR, 1998, pp. 373-410.

[3] Certification Zone documentation on OSPF: <http://www.certificationzone.com> (14 Oct. 2001).

[4] Cisco System, OSPF documentation: <http://www.cisco.com/warp/public/104/1.html> (14 Oct. 2001).

[5] Cisco System, *Introduction Cisco Router Configuration*. Indianapolis, IN: Cisco Press, 1998.

[6] L. Chappell, *Advanced Cisco Router Configuration*. Indianapolis, IN: Cisco Press, 1999, pp. 280-296.

[7] J. Doyles, *Routing TCP/IP, Vol. 1*. Indianapolis, IN: Cisco Press, 1998, pp. 191-228.

[8] B. Halabi, *Internet Routing Architectures 2<sup>nd</sup> Edition*. Indianapolis, IN: Cisco Press, 2000, pp. 57-64.

[9] J. Moy, "OSPF Version 2," RFC 2328, Apr. 1998: <http://www.ietf.org/rfc/rfc2328.txt> (14 Oct. 2001).

[10] M. Saterlee and S. Hutnik, *Cisco CCIE All-in-One Lab Study Guide*. New York, NY: Oracle Press, 1999, pp. 456-461.

[11] J. Walrand and P. Varaiya, *High-Performance Communication Networks*. San Francisco, CA: Morgan Kaufmann Publisher, 2000, pp. 163-165.

# Appendix A

## *Code Listing*

The configurations of a Cisco router are divided into different sections. Lines that follow the “interface” command (e.g., “interface serial 1”) are the configuration for that particular interface. Lines that follow the “router ospf” or “router rip” commands are the settings for the OSPF and RIP protocols.

### Interface Commands:

“encapsulation frame-relay” – use Frame Relay as a layer 2 protocol to encapsulate the packets

“ip address 1.1.1.1” – set the IP address for that particular interface

“frame-relay interface-dlci 102” – use FR VC 102 for that interface

“frame-relay traffic-rate 128000” – set the bandwidth for the specified VC to be 128kbps

“frame-relay traffic-shaping” – enable FR traffic shaping

“no keepalive” – assume the link is always operational without sending any keepalive

### Routing Protocol Commands:

“timers basic 60 360 360 480” – set the RIP timers

“redistribute ospf 7” – redistribute OSPF process 7 into RIP

“network 4.0.0.0” – include interfaces with IP address 4.0.0.0 to 4.255.255.255 to RIP

“network 192.168.1.0 0.0.0.255 area 0” – include all the interfaces with IP address between 192.168.1.0 to 192.168.1.255 to OSPF’s Area 0

Cisco routers have empty configurations originally. We configured the routers such that they are suitable for the needs of this network protocol convergence project. The configurations of the seven Cisco routers plus the Frame Relay Switch that produce the experiment result are listed. Lines begin with an “!” denote comments.

---

---

### **Router FR**

! This router behaves like a FR switch.

! There are two PVCs, one from R1 to R2 and one from R1 to R3.

! Enable this router to be a FR switch

```
frame-relay switching
```

! Connected to the main site R1

! One PVC to R2 and another PVC to R3

```
interface Serial1
  encapsulation frame-relay
  clockrate 64000
  frame-relay intf-type dce
  frame-relay route 102 interface Serial2 201
  frame-relay route 103 interface Serial3 301
```

! Connected to the site R2 with a PVC to R1

```
interface Serial2
  encapsulation frame-relay
  clockrate 64000
  frame-relay intf-type dce
  frame-relay route 201 interface Serial1 102
```

! Connected to the site R3 with a PVC to R1

```
interface Serial3
  encapsulation frame-relay
  clockrate 64000
  frame-relay intf-type dce
  frame-relay route 301 interface Serial1 103
```

## Router R1

! this is the main office's routers and connect between the ATM network and the FR network

! Loopback interface for management purpose

```
interface Loopback0
  ip address 1.1.1.1 255.255.255.0
```

! connected to the fully-meshed ATM link

```
interface Ethernet0
  ip address 10.1.1.1 255.255.255.0
```

! FR interfaces; One 128k VC to R2 and one 256k VC to R3

```
interface Serial0
  ip address 192.168.1.1 255.255.255.0
  encapsulation frame-relay
  ip ospf network point-to-multipoint
  frame-relay traffic-shaping
  frame-relay interface-dlci 102
    class VC_128k
  frame-relay interface-dlci 103
    class VC_256k
```

! OSPF process; FR interface is in Area 0; ATM interface is in Area 1

! Perform summarization of IP address for network 67.1.0.0/16 and 192.168.1.0/24

```
router ospf 7
  redistribute connected subnets route-map loopback
  network 10.1.1.0 0.0.0.255 area 1
  network 192.168.1.0 0.0.0.255 area 0
  area 0 range 192.168.1.0 255.255.255.0
  area 1 range 67.1.0.0 255.255.0.0
```

! access-list and route-map for redistributing the loopback interface to the OSPF routing process

```
ip access-list standard loopback
  permit 1.1.1.0 0.0.0.255
route-map loopback permit 10
  match ip address loopback
```

! Traffic Shaping the FR VC to the desired Bandwidth

```
map-class frame-relay VC_128k
  frame-relay traffic-rate 128000
map-class frame-relay VC_256k
  frame-relay traffic-rate 256000
```

## Router R2

! R2 is a leg for the FR site. It also connects to ISP#2 using a T1 link.

! Loopback interface for management purpose

```
interface Loopback0
  ip address 2.2.2.2 255.255.255.0
```

! FR interfaces; 128k VC to R1

```
interface Serial0
  ip address 192.168.1.2 255.255.255.0
  encapsulation frame-relay
  ip ospf network point-to-multipoint
  frame-relay traffic-shaping
  frame-relay class VC_128k
```

! T1 connection to ISP#2

```
interface Serial1
  ip address 172.16.1.2 255.255.255.0
```

! OSPF Routing process; FR interface is in Area 0

! Redistributing between the RIP and OSPF

! ISP#2 (RIP)'s ntwk that is redistruted back to OSPF would have AD=152

```
router ospf 7
  redistribute connected subnets route-map loopback
  redistribute rip subnets
  network 192.168.1.0 0.0.0.255 area 0
  distance 152 0.0.0.0 255.255.255.255 3
```

! RIP routing process for the S1 Interface

! Also, change the timers and redistributed from the OSPF process.

! The OSPF network that is redistributed to RIP would have AD=153

```
router rip
  timers basic 60 360 360 480
  redistribute connected metric 10 route-map loopback
  redistribute ospf 7
  network 172.16.0.0
  default-metric 7
  distance 152 0.0.0.0 255.255.255.255 7
```

! access-list and route-map for redistributing the loopback interface to the OSPF routing process

```
ip access-list standard loopback
  permit 2.2.2.0 0.0.0.255
route-map loopback permit 10
  match ip address loopback
```

! Traffic Shaping the FR VC to the desired Bandwidth

```
map-class frame-relay VC_128k
```



```
frame-relay traffic-rate 128000
no frame-relay adaptive-shaping
```

```
! Access List that consists of the RIP network
! For redistributing between OSPF and RIP
access-list 3 permit 4.0.0.0 0.255.255.255
access-list 3 permit 172.16.0.0 0.0.255.255
access-list 3 permit 208.1.1.0 0.0.0.255
access-list 3 permit 5.0.0.0 0.0.0.255
access-list 3 permit 137.1.0.0 0.0.255.255
! Access List that consists of the RIP network
! For redistributing between OSPF and RIP
access-list 7 permit 1.1.1.0 0.0.0.255
access-list 7 permit 2.2.2.0 0.0.0.255
access-list 7 permit 3.3.3.0 0.0.0.255
access-list 7 permit 6.6.6.0 0.0.0.255
access-list 7 permit 7.7.7.0 0.0.0.255
access-list 7 permit 10.1.1.0 0.0.0.255
access-list 7 permit 12.1.1.0 0.0.0.255
access-list 7 permit 67.1.6.0 0.0.0.255
access-list 7 permit 67.1.7.0 0.0.0.255
access-list 7 permit 192.168.1.0 0.0.0.255
```

---

---

### Router R3

! R3 is a leg for the FR site. It also connects to ISP#2 using a Ethernet link.

```
! Loopback interface for management purpose
interface Loopback0
 ip address 3.3.3.3 255.255.255.0
```

```
! Ethernet connection to ISP#2
interface Ethernet0
 ip address 137.1.1.3 255.255.255.0
```

```
! FR interfaces; 128k VC to R1
interface Serial0
 ip address 192.168.1.3 255.255.255.0
 encapsulation frame-relay
 ip ospf network point-to-multipoint
 frame-relay traffic-shaping
 frame-relay class VC_256k
```

```
! OSPF Routing process; FR interface is in Area 0
! Redistributing between the RIP and OSPF
! ISP#2 (RIP)'s ntwk that is redistruted back to OSPF would have AD=153
router ospf 7
```

```
redistribute connected subnets route-map loopback
redistribute rip
network 192.168.1.0 0.0.0.255 area 0
distance 153 0.0.0.0 255.255.255.255 3
```

! RIP routing process for the E0 Interface  
! Also, change the timers and redistributed from the OSPF process.  
! The OSPF network that is redistributed to RIP would have AD=153

```
router rip
timers basic 60 360 360 480
redistribute connected route-map loopback
redistribute ospf 7
network 137.1.0.0
default-metric 3
distance 153 0.0.0.0 255.255.255.255 7
```

! access-list and route-map for redistributing the loopback interface to the OSPF routing process

```
ip access-list standard loopback
permit 3.3.3.0 0.0.0.255
route-map loopback permit 10
match ip address loopback
```

! Traffic Shaping the FR VC to the desired Bandwidth

```
map-class frame-relay VC_256k
frame-relay traffic-rate 256000
no frame-relay adaptive-shaping
```

! Access List that consists of the RIP network

! For redistributing between OSPF and RIP

```
access-list 3 permit 4.0.0.0 0.255.255.255
access-list 3 permit 172.16.0.0 0.0.255.255
access-list 3 permit 208.1.1.0 0.0.0.255
access-list 3 permit 5.0.0.0 0.0.0.255
access-list 3 permit 137.1.0.0 0.0.255.255
```

! Access List that consists of the OSPF network

! For redistributing between OSPF and RIP

```
access-list 7 permit 1.1.1.0 0.0.0.255
access-list 7 permit 2.2.2.0 0.0.0.255
access-list 7 permit 3.3.3.0 0.0.0.255
access-list 7 permit 6.6.6.0 0.0.0.255
access-list 7 permit 7.7.7.0 0.0.0.255
access-list 7 permit 10.1.1.0 0.0.0.255
access-list 7 permit 12.1.1.0 0.0.0.255
access-list 7 permit 67.1.6.0 0.0.0.255
access-list 7 permit 67.1.7.0 0.0.0.255
access-list 7 permit 192.168.1.0 0.0.0.255
```

## Router R4

! R4 located inside ISP#2. The interconnection between the ISP#2 is a T1 link.  
! It also connects to ISP#1 R2 with a T1 link.

! Loopback interface for management purpose

```
interface Loopback0
  ip address 4.4.4.4 255.255.255.0
```

! T1 connection between ISP#2

```
interface Serial0
  ip address 208.1.1.4 255.255.255.0
  clockrate 64000
```

! T1 connection to ISP#1

```
interface Serial1
  ip address 172.16.1.4 255.255.255.0
  clockrate 64000
```

! RIP routing process for all interfaces

! Also, change the timers

```
router rip
  timers basic 60 360 360 480
  network 4.0.0.0
  network 172.16.0.0
  network 208.1.1.0
```

---

## Router R5

! R5 located inside ISP#2. The interconnection between the ISP#2 is a T1 link.  
! It also connects to ISP#1 R3 with an Ethernet link.

! Loopback interface for management purpose

```
interface Loopback0
  ip address 5.5.5.5 255.255.255.0
```

! Ethernet connection to ISP#1

```
interface Ethernet0
  ip address 137.1.1.5 255.255.255.0
```

! T1 connection between ISP#2

```
interface Serial0
  ip address 208.1.1.5 255.255.255.0
```

! RIP routing process for all interfaces

! Also, change the timers

```
router rip
```

```
timers basic 60 360 360 480
network 137.1.0.0
network 5.0.0.0
network 208.1.1.0
```

---

## Router R6

! This is a leg of the fully-meshed ATM link.

! It also has a connection to the other ATM site by an ethernet link.

! Loopback interface for management purpose

```
interface Loopback0
 ip address 6.6.6.6 255.255.255.0
```

! connected to the fully-meshed ATM link

```
interface Ethernet0
 ip address 10.1.1.6 255.255.255.0
```

! direct connection between the two ATM sites

```
interface Ethernet1
 ip address 12.1.1.6 255.255.255.0
```

! user can be located on this ethernet interface

```
interface Ethernet2
 ip address 67.1.6.6 255.255.255.0
 no keepalive
```

! OSPF process; all interfaces are in Area 1

```
router ospf 7
 redistribute connected subnets route-map loopback
 network 12.1.1.0 0.0.0.255 area 1
 network 10.1.1.0 0.0.0.255 area 1
 network 67.1.6.0 0.0.0.255 area 1
```

! access-list and route-map for redistributing the loopback interface to the OSPF routing process

```
ip access-list standard loopback
 permit 6.6.6.0 0.0.0.255
 route-map loopback permit 10
 match ip address loopback
```

## Router R7

! This is a leg of the fully-meshed ATM link.

! It also has a connection to the other ATM site by an ethernet link.

! Loopback interface for management purpose

```
interface Loopback0
  ip address 7.7.7.7 255.255.255.0
```

! connected to the fully-meshed ATM link

```
interface Ethernet0/0
  ip address 10.1.1.7 255.255.255.0
```

! direct connection between the two ATM sites

```
interface Ethernet0/1
  ip address 12.1.1.7 255.255.255.0
```

! user can be located on this ethernet interface

```
interface Ethernet1/0
  ip address 67.1.7.7 255.255.255.0
  no keepalive
```

! OSPF process; all interfaces are in Area 1

```
router ospf 7
  redistribute connected subnets route-map loopback
  network 10.1.1.0 0.0.0.255 area 1
  network 12.1.1.0 0.0.0.255 area 1
  network 67.1.7.0 0.0.0.255 area 1
```

! access-list and route-map for redistributing the loopback interface to the OSPF routing process

```
ip access-list standard loopback
  permit 7.7.7.0 0.0.0.255
route-map loopback permit 10
  match ip address loopback
```

## Appendix B

### *Routing Tables*

The routing table of the routers with all links operational are listed in this section.

#### **Router R1**

1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 1.0.0.0/8 [110/123] via 192.168.1.2, 00:00:52, Serial0  
C 1.1.1.0/24 is directly connected, Loopback0  
O E2 137.1.0.0/16 [110/123] via 192.168.1.2, 00:00:52, Serial0  
2.0.0.0/24 is subnetted, 1 subnets  
O E2 2.2.2.0 [110/20] via 192.168.1.2, 00:00:52, Serial0  
3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 3.0.0.0/8 [110/123] via 192.168.1.2, 00:00:52, Serial0  
O E2 3.3.3.0/24 [110/20] via 192.168.1.3, 00:00:52, Serial0  
O E2 4.0.0.0/8 [110/123] via 192.168.1.2, 00:00:52, Serial0  
O E2 5.0.0.0/8 [110/123] via 192.168.1.2, 00:00:52, Serial0  
6.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 6.6.6.0/24 [110/20] via 10.1.1.6, 00:00:52, Ethernet0  
O E2 6.0.0.0/8 [110/123] via 192.168.1.2, 00:00:52, Serial0  
O E2 172.16.0.0/16 [110/321] via 192.168.1.3, 00:00:53, Serial0  
67.0.0.0/8 is variably subnetted, 3 subnets, 2 masks  
O E2 67.0.0.0/8 [110/123] via 192.168.1.2, 00:00:53, Serial0  
O 67.1.6.0/24 [110/20] via 10.1.1.6, 00:00:53, Ethernet0  
O 67.1.7.0/24 [110/20] via 10.1.1.7, 00:00:53, Ethernet0  
7.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 7.7.7.0/24 [110/20] via 10.1.1.7, 00:00:53, Ethernet0  
O E2 7.0.0.0/8 [110/123] via 192.168.1.2, 00:00:53, Serial0  
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 10.0.0.0/8 [110/123] via 192.168.1.2, 00:00:53, Serial0  
C 10.1.1.0/24 is directly connected, Ethernet0  
O E2 208.1.1.0/24 [110/123] via 192.168.1.2, 00:00:53, Serial0  
12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O 12.1.1.0/24 [110/20] via 10.1.1.6, 00:00:53, Ethernet0  
[110/20] via 10.1.1.7, 00:00:53, Ethernet0  
O E2 12.0.0.0/8 [110/123] via 192.168.1.2, 00:00:53, Serial0  
192.168.1.0/24 is variably subnetted, 3 subnets, 2 masks  
C 192.168.1.0/24 is directly connected, Serial0  
O 192.168.1.3/32 [110/64] via 192.168.1.3, 00:12:27, Serial0  
O 192.168.1.2/32 [110/64] via 192.168.1.2, 00:12:27, Serial0

## Router R2

1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 1.1.1.0/24 [110/20] via 192.168.1.1, 00:00:57, Serial0  
R 1.0.0.0/8 [120/5] via 172.16.1.4, 00:00:06, Serial1  
R 137.1.0.0/16 [120/2] via 172.16.1.4, 00:00:06, Serial1  
2.0.0.0/24 is subnetted, 1 subnets  
C 2.2.2.0 is directly connected, Loopback0  
3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 3.3.3.0/24 [110/20] via 192.168.1.1, 00:00:57, Serial0  
R 3.0.0.0/8 [120/5] via 172.16.1.4, 00:00:07, Serial1  
R 4.0.0.0/8 [120/1] via 172.16.1.4, 00:00:07, Serial1  
R 5.0.0.0/8 [120/2] via 172.16.1.4, 00:00:07, Serial1  
6.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 6.6.6.0/24 [110/20] via 192.168.1.1, 00:00:57, Serial0  
R 6.0.0.0/8 [120/5] via 172.16.1.4, 00:00:07, Serial1  
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks  
O E2 172.16.0.0/16 [152/321] via 192.168.1.1, 00:00:58, Serial0  
C 172.16.1.0/24 is directly connected, Serial1  
67.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O IA 67.1.0.0/16 [110/84] via 192.168.1.1, 00:00:58, Serial0  
R 67.0.0.0/8 [120/5] via 172.16.1.4, 00:00:07, Serial1  
7.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O E2 7.7.7.0/24 [110/20] via 192.168.1.1, 00:00:58, Serial0  
R 7.0.0.0/8 [120/5] via 172.16.1.4, 00:00:07, Serial1  
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O IA 10.1.1.0/24 [110/74] via 192.168.1.1, 00:00:53, Serial0  
R 10.0.0.0/8 [120/5] via 172.16.1.4, 00:00:07, Serial1  
R 208.1.1.0/24 [120/1] via 172.16.1.4, 00:00:07, Serial1  
12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
O IA 12.1.1.0/24 [110/84] via 192.168.1.1, 00:00:58, Serial0  
R 12.0.0.0/8 [120/5] via 172.16.1.4, 00:00:07, Serial1  
192.168.1.0/24 is variably subnetted, 3 subnets, 2 masks  
O 192.168.1.1/32 [110/64] via 192.168.1.1, 00:12:37, Serial0  
C 192.168.1.0/24 is directly connected, Serial0  
O 192.168.1.3/32 [110/128] via 192.168.1.1, 00:12:37, Serial0

## Router R3

- 1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 1.0.0.0/8 [110/123] via 192.168.1.1, 00:01:05, Serial0
- O E2 1.1.1.0/24 [110/20] via 192.168.1.1, 00:01:05, Serial0
- 137.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
- O E2 137.1.0.0/16 [153/123] via 192.168.1.1, 00:01:05, Serial0
- C 137.1.1.0/24 is directly connected, Ethernet0
- 2.0.0.0/24 is subnetted, 1 subnets
- O E2 2.2.2.0 [110/20] via 192.168.1.1, 00:01:05, Serial0
- 3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 3.0.0.0/8 [110/123] via 192.168.1.1, 00:01:05, Serial0
- C 3.3.3.0/24 is directly connected, Loopback0
- R 4.0.0.0/8 [120/2] via 137.1.1.5, 00:00:57, Ethernet0
- R 5.0.0.0/8 [120/1] via 137.1.1.5, 00:00:57, Ethernet0
- 6.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 6.6.6.0/24 [110/20] via 192.168.1.1, 00:01:05, Serial0
- O E2 6.0.0.0/8 [110/123] via 192.168.1.1, 00:01:05, Serial0
- R 172.16.0.0/16 [120/2] via 137.1.1.5, 00:00:57, Ethernet0
- 67.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O IA 67.1.0.0/16 [110/84] via 192.168.1.1, 00:01:05, Serial0
- O E2 67.0.0.0/8 [110/123] via 192.168.1.1, 00:01:05, Serial0
- 7.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 7.7.7.0/24 [110/20] via 192.168.1.1, 00:01:05, Serial0
- O E2 7.0.0.0/8 [110/123] via 192.168.1.1, 00:01:05, Serial0
- 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O IA 10.1.1.0/24 [110/74] via 192.168.1.1, 00:01:05, Serial0
- O E2 10.0.0.0/8 [110/123] via 192.168.1.1, 00:01:05, Serial0
- R 208.1.1.0/24 [120/1] via 137.1.1.5, 00:00:57, Ethernet0
- 12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O IA 12.1.1.0/24 [110/84] via 192.168.1.1, 00:01:05, Serial0
- O E2 12.0.0.0/8 [110/123] via 192.168.1.1, 00:01:05, Serial0
- 192.168.1.0/24 is variably subnetted, 3 subnets, 2 masks
- O 192.168.1.1/32 [110/64] via 192.168.1.1, 00:12:44, Serial0
- C 192.168.1.0/24 is directly connected, Serial0
- O 192.168.1.2/32 [110/128] via 192.168.1.1, 00:12:44, Serial0



## Router R4

```
R 1.0.0.0/8 [120/4] via 208.1.1.5, 00:00:12, Serial0
R 137.1.0.0/16 [120/1] via 208.1.1.5, 00:00:12, Serial0
R 2.0.0.0/8 [120/4] via 208.1.1.5, 00:00:12, Serial0
R 3.0.0.0/8 [120/4] via 208.1.1.5, 00:00:12, Serial0
  4.0.0.0/24 is subnetted, 1 subnets
C   4.4.4.0 is directly connected, Loopback0
R 5.0.0.0/8 [120/1] via 208.1.1.5, 00:00:12, Serial0
R 6.0.0.0/8 [120/4] via 208.1.1.5, 00:00:12, Serial0
  172.16.0.0/24 is subnetted, 1 subnets
C   172.16.1.0 is directly connected, Serial1
R 67.0.0.0/8 [120/4] via 208.1.1.5, 00:00:12, Serial0
R 7.0.0.0/8 [120/4] via 208.1.1.5, 00:00:12, Serial0
R 10.0.0.0/8 [120/4] via 208.1.1.5, 00:00:12, Serial0
C 208.1.1.0/24 is directly connected, Serial0
R 12.0.0.0/8 [120/4] via 208.1.1.5, 00:00:14, Serial0
R 192.168.1.0/24 [120/4] via 208.1.1.5, 00:00:14, Serial0
```

---

---

## Router R5

```
R 1.0.0.0/8 [120/3] via 137.1.1.3, 00:00:47, Ethernet0
R 2.0.0.0/8 [120/3] via 137.1.1.3, 00:00:47, Ethernet0
R 3.0.0.0/8 [120/3] via 137.1.1.3, 00:00:47, Ethernet0
R 4.0.0.0/8 [120/1] via 208.1.1.4, 00:00:36, Serial0
  5.0.0.0/24 is subnetted, 1 subnets
C   5.5.5.0 is directly connected, Loopback0
R 6.0.0.0/8 [120/3] via 137.1.1.3, 00:00:48, Ethernet0
R 7.0.0.0/8 [120/3] via 137.1.1.3, 00:00:48, Ethernet0
R 10.0.0.0/8 [120/3] via 137.1.1.3, 00:00:48, Ethernet0
R 12.0.0.0/8 [120/3] via 137.1.1.3, 00:00:48, Ethernet0
R 67.0.0.0/8 [120/3] via 137.1.1.3, 00:00:48, Ethernet0
R 192.168.1.0/24 [120/3] via 137.1.1.3, 00:00:48, Ethernet0
  137.1.0.0/24 is subnetted, 1 subnets
C   137.1.1.0 is directly connected, Ethernet0
R 172.16.0.0/16 [120/1] via 208.1.1.4, 00:00:37, Serial0
C 208.1.1.0/24 is directly connected, Serial0
```

## Router R6

- 1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 1.1.1.0/24 [110/20] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 1.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- 2.0.0.0/24 is subnetted, 1 subnets
- O E2 2.2.2.0 [110/20] via 10.1.1.1, 00:01:37, Ethernet0
- 3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 3.3.3.0/24 [110/20] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 3.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 4.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 5.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- 6.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 6.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- C 6.6.6.0/24 is directly connected, Loopback0
- 7.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 7.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 7.7.7.0/24 [110/20] via 12.1.1.7, 00:01:37, Ethernet1
- [110/20] via 10.1.1.7, 00:01:37, Ethernet0
- 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 10.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- C 10.1.1.0/24 is directly connected, Ethernet0
- 12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 12.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- C 12.1.1.0/24 is directly connected, Ethernet1
- 67.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
- O E2 67.0.0.0/8 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- C 67.1.6.0/24 is directly connected, Ethernet2
- O 67.1.7.0/24 [110/20] via 12.1.1.7, 00:01:37, Ethernet1
- [110/20] via 10.1.1.7, 00:01:37, Ethernet0
- O IA 192.168.1.0/24 [110/10] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 137.1.0.0/16 [110/123] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 172.16.0.0/16 [110/321] via 10.1.1.1, 00:01:37, Ethernet0
- O E2 208.1.1.0/24 [110/123] via 10.1.1.1, 00:01:37, Ethernet0

## Router R7

- 1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 1.1.1.0/24 [110/20] via 10.1.1.1, 00:01:42, Ethernet0/0
- O E2 1.0.0.0/8 [110/123] via 10.1.1.1, 00:01:42, Ethernet0/0
- O E2 137.1.0.0/16 [110/123] via 10.1.1.1, 00:01:42, Ethernet0/0
- 2.0.0.0/24 is subnetted, 1 subnets
- O E2 2.2.2.0 [110/20] via 10.1.1.1, 00:01:42, Ethernet0/0
- 3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 3.3.3.0/24 [110/20] via 10.1.1.1, 00:01:42, Ethernet0/0
- O E2 3.0.0.0/8 [110/123] via 10.1.1.1, 00:01:42, Ethernet0/0
- O E2 4.0.0.0/8 [110/123] via 10.1.1.1, 00:01:42, Ethernet0/0
- O E2 5.0.0.0/8 [110/123] via 10.1.1.1, 00:01:42, Ethernet0/0
- 6.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 6.0.0.0/8 [110/123] via 10.1.1.1, 00:01:42, Ethernet0/0
- O E2 6.6.6.0/24 [110/20] via 12.1.1.6, 00:01:42, Ethernet0/1
- [110/20] via 10.1.1.6, 00:01:42, Ethernet0/0
- O E2 172.16.0.0/16 [110/321] via 10.1.1.1, 00:01:43, Ethernet0/0
- 67.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
- O E2 67.0.0.0/8 [110/123] via 10.1.1.1, 00:01:43, Ethernet0/0
- O 67.1.6.0/24 [110/20] via 12.1.1.6, 00:01:43, Ethernet0/1
- [110/20] via 10.1.1.6, 00:01:43, Ethernet0/0
- C 67.1.7.0/24 is directly connected, Ethernet1/0
- 7.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 7.0.0.0/8 [110/123] via 10.1.1.1, 00:01:43, Ethernet0/0
- C 7.7.7.0/24 is directly connected, Loopback0
- 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 10.0.0.0/8 [110/123] via 10.1.1.1, 00:01:43, Ethernet0/0
- C 10.1.1.0/24 is directly connected, Ethernet0/0
- O E2 208.1.1.0/24 [110/123] via 10.1.1.1, 00:01:43, Ethernet0/0
- 12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
- O E2 12.0.0.0/8 [110/123] via 10.1.1.1, 00:01:43, Ethernet0/0
- C 12.1.1.0/24 is directly connected, Ethernet0/1
- O IA 192.168.1.0/24 [110/10] via 10.1.1.1, 00:01:43, Ethernet0/0