

# CCIE Warm-Up Advice and Learning Labs

Written by Thomas P. Larus  
CCIE 10,014



Distributed by  
IPexpert, Inc. ([www.ipexpert.com](http://www.ipexpert.com))

For technical support  
please visit [www.CertificationTalk.com](http://www.CertificationTalk.com)

Copyright© 2003 by Thomas P. Larus. All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the author.

This book is not sponsored by or endorsed by or affiliated with Cisco Systems, Inc. Cisco®, CCNA™, CCNP™, CCDP™, CCIP™, CCIE™, and the CCIE logo are trademarks or registered trademarks of Cisco Systems, Inc., in the United States and in certain other countries. All other products and services mentioned in this book are the trademarks or service marks of their respective companies or organizations.

Information contained in this book is intended for educational purposes only, and no guarantee is made as to the correctness or completeness of any information contained herein. The author will not be responsible for any errors, omissions, or damages arising out of the use of this information. The author is supplying information, not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.

Mr. Larus can be reached at [tlarus@ipexpert.com](mailto:tlarus@ipexpert.com). Mr. Larus welcomes feedback related to this book, and is available to prepare and deliver customized training on Cisco networking.

For specific questions relating to the scenarios, please go to [www.certificationtalk.com](http://www.certificationtalk.com).

## Table of Contents

Preface	iv
Chapter 1 Advice on CCIE Lab Preparation	1
Chapter 2 CCIE Practice Lab Equipment	15
Chapter 3 Scenario One	19
Chapter 4 Scenario Two	57
Chapter 5 Scenario Three	85
Chapter 6 Scenario Four	114
Chapter 7 Scenario Five	138
Appendix A General Information and Opinion about the CCIE	171
Appendix B Annotated Bibliography	177

## Preface

One of my goals in pursuing the CCIE was to be able to write a book about how to prepare for the CCIE Lab. This is the book, with practical advice and practice scenarios meant to impart to the reader many of the important lessons I learned during CCIE Lab preparation.

The intended audience is CCIE candidates in the first six months or so of focused preparation for the CCIE Lab exam. The CCIE Lab I took was in the Routing and Switching Track, but many of the central lessons I learned should be useful in other tracks. The labs in this book focus heavily on the routing protocols BGP, IS-IS, OSPF, EIGRP, and RIP version 2, and covers numerous Quality of Service features such as priority queueing, custom queueing, class-based weighted-fair queueing, frame-relay traffic-shaping, PPP multilink fragmentation and interleaving, and rate-limiting. The scenarios also address IP Multicast, IPSEC, IOS firewall, and basic DSLW+. The reader will configure a Catalyst 3550-EMI as a switch and as a router, but the labs focus more on routing than on switching.

The advice and scenarios emphasize lessons that candidates preparing for the CCIE Lab Exam need to learn. This book is not intended for use in preparation for the CCIE written qualifier exam, which is a very different test, covering much material that Cisco has expressly and publicly removed from the CCIE Lab Exam, such as IPX, ATM LAN Emulation, and Token Ring.

The scenarios would, incidentally, serve as excellent lab material for the student preparing for the Building Scalable Cisco Internetworks (BSCI) Exam, which is required for the CCNP, CCDP, and CCIP certifications, as well as for the Routing and Switching Specialization for Cisco Partner resellers. The student could simply ignore the QOS tasks and other tasks that are obviously beyond the scope of the BSCI Exam, and enjoy some very challenging routing labs.

I designed the five practice scenarios in this book to teach crucial CCIE lessons, not to resemble the CCIE Lab Exam. Since I use an affordable lab topology that leaves out expensive ATM gear, and includes only one Catalyst 3550-EMI, it will never be said of any of these labs that, "it was just like the real lab exam." I want readers to know this up front so that they will not expect scenarios that are "dry runs" of the Lab exam. My tasks involve very little cryptic language, so you will usually know *what* you need to do. You just may not know *how* to do it until you do some research or read my solution configuration scripts and explanations. Much of the value of this book lies in the explanations of the tasks that require explanations.

The lessons run the gamut from efficient CLI practices, diagramming techniques, and technology-specific tricks and traps, to subtleties of route redistribution. I am obsessed with route redistribution. I found it challenging while preparing for the CCIE Lab Exam, and still find things to learn about it. If you do not know how to do route redistribution properly, it will not matter that you know how to configure ATM, Voice, ISDN, and IP Multicast perfectly, because your network will be a non-functioning mess if any route redistribution is required.

The scenarios in this book should, ideally, be done after you have already completed numerous short, technology-specific exercises, but before you attempt the long and involved multiprotocol scenarios found in expensive commercial CCIE practice lab workbooks. My lab scenarios are intended to be a warm up to the more advanced commercial scenarios, not to replace them.

The CCIE lab "warm-up" provided by this book serves several purposes. First, it permits a CCIE candidate to start down the path of CCIE Lab preparation without spending a lot of money on equipment and commercial lab scenarios. Some people pass the CCIE written qualifier exam, buy a practice rack, and then find out that the lab practice is not as engaging as they had hoped. It is better to find this out before you have spent \$500 on a set of commercial scenarios and \$7,000-\$10,000 on a first-rate CCIE practice lab. Second, it may be better to learn some of the more difficult and crucial lessons from a book that focuses on explaining those lessons before you dive into the troubled waters of first-rate commercial practice scenarios. If you learn some of the advanced lessons before you begin the advanced commercial labs, you may be able absorb more of the fine points of those advanced commercial labs.

There is no reason that you should have to wait until you have purchased \$500 lab scenarios or attended a \$4,000 CCIE boot camp to be exposed to some of the more interesting and subtle points we all have to learn during CCIE Lab preparation. This book's practice scenarios are not intended to be full-blown workouts like the advanced labs in the [www.ipexpert.com](http://www.ipexpert.com) CCIE Lab Workbooks.

These are relatively lean lab scenarios, and I worked hard to insure that they can be done using a practice lab that is far less expensive than those required for the more involved commercial labs, and without any re-cabling between scenarios. My expectation is that several rental lab companies will offer lab racks on which to practice the scenarios in this book, thereby providing a cost effective alternative to purchasing equipment.

I would like to add just a quick note on equipment, because it will be discussed more fully Chapter Two. CCIE candidates today can purchase lab equipment at much lower prices than I paid in the years 2000-2002, when I built my first Cisco practice lab. You can now purchase a lab of seven routers, a frame relay switch, an ISDN simulator, and a closet switch capable of VLANs, for less than \$3,000. My lab scenarios are designed to be implemented in a lab of four 2501s, two 2503s, an ISDN simulator, and a Catalyst 3550-EMI. The Catalyst 3550-EMI switch adds a lot to the price tag, but is so important an addition to the official CCIE Lab Exam Equipment List, and is such a powerful device, that I have incorporated one into these scenarios.

For readers who simply cannot afford the 3550-EMI switch, a cheap 4500 or 4700 router with at least four ethernet interfaces will fulfill the Catalyst 3550-EMI's routing functions perfectly, and a cheap VLAN-capable closet switch will handle the Cat 3550's basic switch functions well. You still need to rent some time on a Cat 3550-EMI at some point to learn its specific features and syntax. Furthermore, it is also possible to do these labs with only six routers and a frame switch, if one has a spare serial interface on the frame switch so that it can be used in lieu of the seventh router. If you have a Cisco 2511 terminal server, that will function perfectly as R7.

The value of the CCIE certification lies primarily in what you learn from pursuing it. Try to enjoy the journey.

# Chapter One

Advice on CCIE  
Lab Preparation

## Obsessive Interest Helps

Two keys to success are persistence and interest. These are especially important in pursuing the CCIE, because the candidate is following a self-directed study regimen with a final exam that most candidates fail more than once. I recommend that the CCIE candidate have an obsessive interest in learning the CCIE material. You may wonder why I am including any pep talk or general advice in this book at all, but you will learn, if you have not already, that for most students getting the CCIE involves much more than learning technical content or developing technical skills. Persistence comes into play when you may be tempted to give up the chase after seeing the toll CCIE preparation and test-taking takes on your social life, your finances, your family, or your emotions. The stories of people who passed on their 5th attempt are inspiring, because of the persistence and determination the candidates showed.

Take CCIE study material with you everywhere you go where there is a chance of getting in five minutes of reading. Read before you go to bed. Read when you wake up on Saturday, if you can. Read at the car repair shop, and on the train. (Do *not* read while driving). I used to study before bed, because it would help me to sleep to know that I was doing something that would further my progress toward CCIE. You may not want to take it quite so far, but at least adopt the approach of weaving CCIE preparation into your everyday life and work toward the goal every chance you get.

Do not become a complete social outcast. Reading while out at a restaurant with your family could be carrying things a bit far, but I feel certain that some successful candidates have done this. The idea is to become as engaged in this course of study as you have ever been in anything in your life. Some people spend considerable time and energy following sports teams, or playing fantasy games in which they are team owners and have constant communication with other fantasy team owners. Some people become obsessed with a new love in their life. Channel your energy and time and focus into preparing for the CCIE as you would for these other pursuits.

You should not treat the CCIE the way some people treat night courses. They attend when they feel like attending, study when they feel like studying, and make slow progress overall. If you approach the CCIE as a subject you can pick up for a while, and then put down for a while, the material covered by the CCIE Lab could change so that the material will get away from you. I am sure that by the time some candidates mastered the token ring switch and ATM LANE, Cisco took these items off the Lab Exam. I have heard of people who bought expensive, first-rate lab scenarios, but never even unwrapped them before a new version of the scenario workbook came out much later.

At the risk of sounding extreme, at some point you need to decide whether you are going to get the CCIE or not, and once you decide to get it you should never let up (except for brief respites) until you get it.

## A lab notebook can help you to digest lessons learned



One thing I did from the start, and highly recommend, is keeping a lab notebook. In this notebook you put a general description of what you worked on each day, how much time you spent configuring routers, and the main lessons you learned. You might also want to record time you spent reading, along with a few salient points from the reading. I would not pretend that my lab notebooks are perfect, but even writing down a one-sentence statement of what I learned from several hours of wrestling with a problem helped me digest and remember the lesson.

I used those sewn notebooks that children use in grade school, and that adults use for personal journals. When I started, I wrote down practically everything I did, in too much detail. Along with some CCIE-level lessons, there were basic CCNA-level lessons about DTE versus DCE, needing "no shut" on every physical router interface after configuration, and the like. This is the stage in CCIE lab preparation where people with even a year or two of solid hands-on experience with Cisco routers have a big edge over the newcomers.

Unless you have considerable hands-on configuration experience, you will need to take some time with these baby steps, and it helps to record what you learn if you want to digest the lessons well. My lab journal reveals that I spent many days working to get "ip ospf demand-circuit" to work right. If you look at the [www.groupstudy.com](http://www.groupstudy.com) ccielab list archives you will see that I am not alone in struggling with this subject. It seems simple once you have learned it, but in the course of learning it you get countless hours of practice with ISDN and dialer and PPP debugging commands.

One practice I ended up dropping was printing out my configuration scripts and stapling them into my lab notebook. I seldom had any need to refer to my old configurations for basic scenarios I made up myself, and scenarios from books and commercial workbooks come with solution sets that will serve well enough. You still might want to print or save noteworthy configs, of course.

The lab notebook shows me mistakes I made. One day, I wrote, "I have the OSPF scenario wired so that I don't need the LOUD 2828 switch. I need to go back and practice the switch-based items later." That was a serious error. Because I did not take every opportunity to practice carving out VLANs on a switch, I was much too slow at that basic task when it came time for my first lab exam. What makes it worse is that I had a Catalyst 5000 with 10BaseT blades that were awkward to use with the used patch panel I bought, as well as fiber-optic fast-ethernet blades that had limited usefulness with my 2500-series routers. I should have bought a 10/100 blade like the WS-X5213A, and taken every opportunity to practice basic switch configuration tasks like carving out VLANs.

This is only the most egregious mistake I spotted in my lab notebooks. It is interesting to watch people coming after me going through the same learning process. In this book, I hope to help readers to avoid making some of the mistakes I and many others made.

## **View the CCIE as an Exam to pass, not as a “Lifetime Achievement Award,” and Use the IOS Documentation and the Equipment List to Set Boundaries to CCIE Material**

Approach the CCIE as a single, discrete goal to achieve, a single battle to win, a single hill to take, and do what it takes to achieve it. Too many people view the CCIE as a Nobel Prize or “lifetime achievement award” for networking professionals. This view could lead some to decide that they are not worthy to have the CCIE until they know everything there is to know about computer networking. When the Allies planned the D-Day invasion, I would imagine that they did not practice jungle warfare. Jungle warfare may be good to learn, and might well contain lessons that would apply to an invasion of Europe through France, but it makes more sense to focus limited time on developing the skills most pertinent to achieving the goal.

Therefore, if you think that you should not attain the CCIE until you have become an expert at Linux and shell scripting and PERL, or whatever other subjects you think a CCIE “should” know about, lose that attitude. If you want to learn these other things for a job, or because you think these are more marketable skills than the ones covered by the CCIE (and they may well be), then go ahead. Just do not distract yourself from CCIE studies because of some fear that you will be unworthy of the CCIE if you do not know all these other subjects.

The CCIE Lab covers a vast, but finite, set of material. You should be very well acquainted with the IOS documentation set. As this book goes to press, 12.2 has officially become the current IOS version in the CCIE R&S Lab. This IOS documentation set, many thousands of pages long, gives you a rough idea of the borders of the CCIE world, if you skim it with the CCIE equipment list in mind. It covers a lot of material, and there are details of theory that are not taught in the pages of the IOS documentation set, so you still need to read Doyle, Halabi, *et al.* The 12.2 IOS documentation gives you a great map of the material to cover. When 12.1 was the IOS version to know, it was a comforting exercise to flip through the 12.1 Multiservice Configuration Guide, and use the CCIE Lab Equipment list to eliminate material from the scope of my CCIE preparation. When I got to a long section that dealt with cable head-end routers, I understood that this was outside the scope of material that I must study for the CCIE Lab.

Some will protest that this is the wrong approach to take, that this is merely studying for the test, not becoming a CCIE. A CCIE is someone who has passed the CCIE written qualifier exam and the Lab exam, nothing more and nothing less. You may indeed want to learn about cable head-end routers and SOHO cable routers, but you should do this for reasons having to do with a job, or a burning desire to learn about those routers. Do not think for a minute that you should become well versed in every matter relating to Cisco equipment before you deserve to pass the CCIE, or you may never pass, or it will take so long that it really could amount to a “lifetime achievement award.”

Some will protest that you could be asked to work with this or that technology on a job someday, and as a CCIE you should be competent with everything. The universe of things that you could be asked to do in a network engineering job is vast, ever expanding, and without boundaries, and has little relation to the changing but bounded

universe of things that are tested on the CCIE Lab Exam. Cisco took IPX, DECNET, and Appletalk off the Lab Exam. Should you spend precious time mastering IPX that could have been spent on mastering the intricacies of OSPF or IS-IS, just so that you can feel like a "real CCIE" when you finally pass? Not if you want to pass before you are ready to retire. While you are mastering Appletalk, new technologies will come out that need to be learned. You may never master "everything," so concentrate on what needs to be learned for your work and for the exam.

## **Be realistic about your strengths and weaknesses as a CCIE candidate**

Before you set out to achieve this goal (at least before you plunk down \$7,000 for lab equipment), you should think through what it could cost in time, effort, money, and strain on family and social relations. You should also assess your ability to learn what you will need to learn. These days, you will also want to assess the job market and the benefit that you personally would gain from the pursuing and achieving the CCIE. Remember to include both the benefits of the certification itself as well as what you will have learned in the process. The costs are enormous.

I did not decide that I was going for the CCIE until I was well on my way through the CCNP program. I had always read ahead to the next level, and I finally had read enough that I knew that there was no difficult math or other cryptic material that could keep me from achieving the CCIE. I knew that there was a lot to learn, but nothing that I could not learn. Most readers of this book have become comfortable with binary numbers. You will need to get very comfortable with binary and comfortable with hexadecimal, but there is no need to worry about any math more complicated than that. There are no derivatives, or integrals, or anything like that. We are configuring routers, not designing algorithms or routing protocols. This is not rocket science, but there is a lot of material to learn well, and you need to be able to analyze a complex scenario, spot the issues that are raised, plan your solution, and execute it with few errors.

The conventional wisdom is that the CCIE certification is for network engineers with many years of experience building and troubleshooting production networks, and that you should not pursue the CCIE without at least two years of solid work experience with Cisco routers and switches. This is a reasonable guideline, and real-world experience will certainly help in the early stages of CCIE Lab preparation, and it will, just as importantly, help you to get a job and to perform well in that job.

Nevertheless, it is another matter entirely whether you can learn what you need to learn to pass the CCIE Lab Exam. To state an unpopular truth, many CCIE candidates do not configure Cisco routers or switches on a daily basis as part of their work. Many of them are pursuing higher Cisco certifications in order to change the nature of the work they do everyday, so that they will be able to do more work with Cisco routers and switches.

Too many people sell themselves short and worry overly much about their dearth of on-the-job experience. I am a career-changer with a background in law, and I can assure

you that the analytical skills I developed in law school, college, a judicial clerkship, and law practice, as well as the experience of preparing for and taking two-day Bar Exams in Virginia and West Virginia helped me to pass the CCIE. Even if you are not yet a seasoned network engineer, you may have skills and experience that will help you to attain the CCIE, and which will make you valuable in a job involving computer networking. Perhaps your education honed your analytical skills, or your civilian or military work experience prepared you to perform under pressure.

Conversely, you may be a seasoned network engineer or network technician, but not like learning from books. Your experience and knowledge will help you, but you need to recognize that the path to CCIE may be just as hard for you as it is for some people with much less hands-on experience. For other experienced engineers, the unrealistic designs, demands, and constraints of CCIE practice scenarios may be annoying. Some experienced and highly educated engineers may dislike focusing on one vendor's implementation and syntax, rather than on theory and RFCs. To attain the CCIE, all of these experienced networkers need to embrace the CCIE material, if only for a time.

## **Set a reasonable daily quota for time configuring routers and switches in the lab**

Some CCIEs talk about having studied eight hours or more each day for nine months to a year, and I know I never could have managed that. I recommend setting a realistic quota of hours that you need to spend on the routers each day, factoring in a few days off. I found that when I got many hours in one day, the next day sometimes was less fruitful. Toward the end, I was more efficient, and built some momentum, so that I had many long days in a row. For me, the quota was only two hours on the routers each day at the beginning of my study, while I got six to eight hours each day toward the end. The quota covers only time configuring routers, and does not include reading and on-line research and discussion.

I recommend using the lab equipment as often as you can, even if for only 45 minutes at a time, rather than waiting until you have long stretches of time available. Configuring routers and switches needs to be something you do every day, liking making coffee or toast, or driving to work, or catching the bus. That way, when you sit down to take the lab exam, you will be doing something that you do every day. You are less likely to have stage fright if you configure Cisco equipment every day.

Everyone has a different job, commute, and personal life. Some can do many lab scenarios on the job, while others have to make real sacrifices to squeeze in ten or twelve hours each week of solid configuration. Just remember that ten hours per week for a year adds up to 520 hours, which could be just enough practice to permit some candidates to pass the CCIE Lab exam.

## **Use an efficient approach to the Command Line Interface early on in your study**

I made serious mistakes in this department, and deeply regret the time and energy I wasted. When I speak of getting in longer days toward the end of my CCIE Lab preparation than I got in at the beginning, I believe this was possible because of the physical and mental energy I saved by adopting more efficient CLI techniques after attending a CCIE boot camp. When I arrived at the class, I was switching between terminal server sessions using "resume" and the router names. The instructor insisted that we set our terminal server sessions so that they corresponded to router numbers.

This is what it looks like when you set up your sessions this way:

```
termserver>r1
Translating "r1"... domain server (192.168.0.1)
(192.168.0.1)Trying r1 (192.168.0.10, 2001)... Open

õ
termserver>r2
Translating "r2"... domain server (192.168.0.1)
(192.168.0.1)Trying r2 (192.168.0.10, 2002)... Open

õ
termserver>r3
Translating "r3"... domain server (192.168.0.1)
(192.168.0.1)Trying r3 (192.168.0.10, 2003)... Open

õ
termserver>r4
Translating "r4"... domain server (192.168.0.1)
(192.168.0.1)Trying r4 (192.168.0.10, 2004)... Open

õ
termserver>r5
Translating "r5"... domain server (192.168.0.1)
(192.168.0.1)Trying r5 (192.168.0.10, 2005)... Open

õ
termserver>r6
Translating "r6"... domain server (192.168.0.1)
(192.168.0.1)Trying r6 (192.168.0.10, 2006)... Open

õ
termserver>r7
Translating "r7"... domain server (192.168.0.1)
(192.168.0.1)Trying r7 (192.168.0.10, 2007)... Open

õ
termserver>cat
Translating "cat"... domain server (192.168.0.1)
(192.168.0.1)Trying cat (192.168.0.10, 2008)... Open
```

cat>en

Actually, you only see the "o" with the tilde over it when you cut and paste the output to Word. That symbol appears where you press "control-shift-6-x" to return to the terminal server. You may also notice that I do not move into privileged exec mode on the terminal server. This prevents me from configuring anything on the terminal server. Rental rack companies will not give you privileged exec access to the terminal server, so you may as well refrain from getting used to features like the "send" command, which permits you to send the same command to numerous routers at once.

I also made more aggressive use of aliases and I recommend aliases to anyone who is a poor typist. At first, I resisted using many aliases, thinking that they involved yet another layer of memorization. The compromise I reached was to use many aliases, but to use aliases that were easy to memorize. Almost all of them consist of the first letter of each word being abbreviated.

Here is a list of aliases that I used while doing practice labs. You may improve them by shortening them further if you are good at memorization:

```
alias configure rr router rip
alias configure ro router ospf
alias configure rb router bgp
alias configure ri router igrp
alias configure re router eigrp
alias exec sir sh ip rou
alias exec sib sh ip bgp
alias exec s sh runn
alias exec c conf t
alias exec sio sh ip ospf
alias exec sb sh runn | begi n
alias exec siib sh ip interfac brief
alias exec sri sh runn interf
alias exec sx sh ipx rou
alias exec sis sh ipx servers
alias exec cir clear ip rout
alias exec cib clear ip bgp
```

You will not need the alias involved in IPX, of course, since IPX is no longer on the Lab Exam. If you are a great typist, you may not need to trouble yourself with aliases. They save me a lot of time and mental energy, though. I prepared Scenario Two without aliases (except on Cat), and I missed most keenly the "s" for "show runn," "c" for "config t," "sib" for "show ip bgp" and "cib" for "clear ip bgp."

Along the same lines, cutting and pasting can save you some time, but you need to recognize the limits of cutting and pasting. In your preparation, you will have occasion to cut and paste entire configuration files, and you had better watch to make sure what parts are accepted by the router and what parts are being rejected. You cannot cut and paste the OSPF router configuration without having at least one OSPF interface "up." This illustrates one more practical benefit of using loopback interfaces, which are up by default, with OSPF. A big problem for cutting and pasting is route-maps. The router will reject a statement applying a route-map until the route-map has been

configured. Therefore, you may want to develop a habit of cutting and pasting the route-maps and access-lists first, then going back and cutting and pasting earlier parts of the configuration script.

## **Find a study partner, and use websites like [www.groupstudy.com](http://www.groupstudy.com) and [www.certificationtalk.com](http://www.certificationtalk.com)**

I did not find an individual study partner until I took a CCIE preparation class. For two months after class ended, I bounced ideas off him and asked him advice about problems I was having with scenarios, and told him about lab findings. He shared his insights with me and told me of his lab findings, and gave me encouragement. You can get this same benefit from participating actively in on-line certification-oriented study groups, and you have a lot more knowledge and experience to draw on than one study partner would provide.

I found the "CCIE Lab" mailing list provided by [www.groupstudy.com](http://www.groupstudy.com) to be of enormous value. Numerous CCIE candidates scoff at the groupstudy ccielab mailing list because it is not a very exclusive virtual community, and fills their inbox with too much e-mail. The amount of e-mail is staggering, and not all messages are helpful, but if I had not used the list, I would have had a longer and lonelier journey to CCIE.

Some new list members ask questions that suggest that they have never seen a router before they just received one from an ebay merchant that day, and that may be true for some. Some answers are not perfect. It is like a classroom. Some of us raise our hands and say, "I'll take a stab at it," with an implicit (or explicit) "Correct me if I'm wrong." There are thousands of people around to correct you if you are wrong, too, so it is unlikely that too many wrong answers will go uncorrected or undisputed. The richness of the resource is inestimable. As for experts, you have countless CCIEs and highly knowledgeable non-CCIEs that read the list carefully and answer questions. If you ask a question on the groupstudy ccielab list, you could get a response from one of many CCIE instructors and authors of network engineering textbooks. All for free. It does not get any better than that.

I used [www.certificationtalk.com](http://www.certificationtalk.com) back when it was an excellent support site for IPExpert products, but it has become so much more than a support site. For specific questions about scenarios in IPExpert products, this is really the best place to go. The IPExpert materials are now so popular that people ask questions on the various lists at [www.groupstudy.com](http://www.groupstudy.com). At [www.certificationtalk.com](http://www.certificationtalk.com), your question can be answered by someone who wrote the scenario, or, by a fellow student or CCIE who struggled with the same issue and is dying to share what he or she learned.

When I checked recently, [www.certificationtalk.com](http://www.certificationtalk.com) had over 9000 registered members, many messages and a wealth of unexpected features, such as a web-based IRC chat room and a calendar where people could post their upcoming lab date and test location. This last feature makes getting together socially the night before the exam a lot easier, and you might be surprised how many people want to do this. Of course, there

would not be 9000 active participants at any given time, but it is a very active and content-rich forum.

Look out for other Cisco-related IRC chat rooms, and check out the internet newsgroups related to Cisco certification. There is also a subscription site, [www.certificationzone.com](http://www.certificationzone.com), with many in-depth technical articles and other resources.

## CCIE Prep Courses can be well worth the money

I attained the CCNA and CCNP certifications entirely through a self-directed study program. When it comes to the CCIE, however, it can definitely help to attend a good CCIE preparation class or "boot camp." All CCIE boot camps are different, and each student needs or gets different benefits from a given boot camp based on his level of preparation going in.

For me, it was therapeutic to go through each technology in a structured manner with an authoritative expert who could point out some of the landmines that can arise with a particular technology. In addition, a CCIE preparation class offers some advanced lessons on theory or little nuggets of syntax that you might not run across in your own studies. I had done considerable reading about DLSW+, but I did not remember reading about the commands "dlsw disable" and "no dlsw disable" until I attended a boot camp. The first command is necessary to reset the DLSW+ process so that your filters can take effect. Together, these commands do for you in DLSW+ what the crucial command "clear ip bgp \*" and its less disruptive variations, "clear ip bgp \* soft in," and "clear ip bgp \* soft out" do for you in BGP.

This point brings me to one of the great benefits of a prep course. When you have completed a good prep course you should feel that you have covered almost all of the technologies that need to be covered, and been exposed to many of the neat little tricks and gotchas and features. You need to go further and continue to study the nooks and crannies of the IOS documentation thoroughly, but going through a prep course allows you to feel that you have met "the usual suspects," in terms of common issues and landmines that can torpedo a lab scenario. This builds confidence so you can be calm in your Lab Exam and can deal with the issues at hand.

There are numerous well-established boot camp programs and excellent new ones are formed from time to time. I hesitate to mention any, because of the dynamic nature of this very specialized industry, where a new provider can gain reputation and market share quickly. I mention a few organizations, with the caveat that you should follow up with your own research. First, [www.ipexpert.com](http://www.ipexpert.com) offers CCIE Bootcamps in San Jose, New York, NY, Washington, DC, Chicago, and Dallas. To my knowledge, IPExpert is the only organization that employs as their CCIE Boot Camp Instructor a Quad CCIE (working on his fifth CCIE). IPExpert offers training for all four CCIE tracks (R&S, Security, Service Provider, and Voice). IPExpert offers its CCIE boot camps as both Instructor-Led Training and Distance Learning vClasses.



Other CCIE boot camp providers include [www.netmasterclass.net](http://www.netmasterclass.net) in Herndon, Virginia; [www.ccbootcamp.com](http://www.ccbootcamp.com) in Rochester Hills, Michigan, in Marysville, Michigan, and [www.cyscoexpert.com](http://www.cyscoexpert.com) in Lincolnwood, Illinois. Global Knowledge ([www.globalknowledge.com](http://www.globalknowledge.com)) also offers various classes aimed at CCIE preparation. Heinz Ulm ([www.heinzulm.com](http://www.heinzulm.com)) has been offering a multi-week boot camp in Augsburg, Germany for many years.

As you do your internet research, and read reviews, keep in mind that each CCIE candidate will usually attend only one company's CCIE boot camp(s), and upon passing the Lab Exam that candidate will have a high regard for the boot camp he attended, and will have no personal experience of any of the other boot camps. Therefore, a rave review of one provider does not mean that others would receive a less favorable review, and you need to examine the specific strengths and weaknesses of each course, and how those strengths and weaknesses map to your needs. For example, one candidate may benefit most from a class that covers a broad array of technologies, while another may need one that quickly finds and then focuses on his weak areas.

## Use Commercial CCIE Scenario Workbooks

I mentioned before that this book is not intended to replace the commercial lab workbooks. It is intended as a warm-up for these commercial labs, so that you can absorb more of their fine points when you work through them. You should research the various offerings and choose at least one. Make sure you have access to equipment that will support the topology of the commercial lab scenarios you choose.

I used an earlier version of the Routing and Switching CCIE workbook from [www.ipexpert.com](http://www.ipexpert.com), and found their advanced labs extremely challenging and content-rich. IPExpert offers Workbooks for all four CCIE tracks. I am told that these workbooks are updated at least once each quarter with enhanced content and technical updates. IPExpert also offers free vLectures for all Workbook customers. It also offers boot camps, and rental racks, and prides itself on its rich on-line content with e-scenarios for various Cisco certifications ranging from CCNA to CCIE.

Another major provider of lab scenarios is [www.ccbootcamp.com](http://www.ccbootcamp.com), which also offers workbooks, boot camps, and on-line rental racks designed for their lab scenarios.

I strongly recommend that you follow up with your own research, asking friends and members of on-line forums what materials they recommend, and why. I cannot emphasize enough that assessments of quality can be very subjective, different people have different needs and styles, product offerings change, and individual products tend to improve through revisions over time.

## Eliminate opportunities for human error

The common advice is to make your own detailed diagram of the lab IP addresses, routing protocols, etc. That is good advice for most people, but be careful that you are accurate and get all the interface numbers and IP addresses correct on your diagram. If you work from an incorrect diagram, you could have big troubles. If any diagrams are provided, and you are error-prone, you might want to consider relying on printed diagrams and reserve making your own logical diagrams for tasks requiring them, in areas such as BGP and DLSW+. This is controversial advice, so consider carefully whether it would suit you.

The theme of eliminating opportunities for human error also underlies my advice to use aliases if you are not a great typist, and to avoid using distribute-lists for preventing route-feedback unless you like them. For prevention of route feedback, I prefer route tagging and route-maps because it seems more elegant to me, because it sometimes allows for greater redundancy, and because it involves fewer opportunities to make a typographical error in an IP addresses or wildcard mask. On the other hand, if you are a very accurate person, you may feel more comfortable "nailing things down" using distribute-lists.

## **Nail everything down**

"Nail everything down" means manually to configure such things as frame relay maps, even though inverse ARP may be available. Inverse ARP can work, but you are wise not to count on it. We all know about "frame-relay map ip 172.16.1.1 102 broadcast," but many of us do not remember "frame-relay map bridge 102 broadcast" when doing transparent bridging, or "frame-relay map cls 102 broadcast" for IS-IS routing.

Nailing down OSPF router IDs is important, particularly in practice scenarios where the authors later tell you to add a loopback address on a router as part of the BGP section. Nailing down the port speed and duplex on a Fast Ethernet interface is a good idea, too. As most of you already know, duplex mismatches due to problems with auto-negotiation on 10/100 ethernet interfaces are a big "real world" problem. A duplex mismatch problem often goes undetected for some time, because there is some basic connectivity, albeit with extremely brief interruptions or degraded performance. There are exceptions to the principle that you should nail everything down, but the general rule is a good one.

## **Use basic network monitoring tools in your practice lab**

Ping scripts that you run to test connectivity from time to time are good, but I recommend going further and using software tools that ping your devices at intervals of much less than a minute and also try implementing SNMP in your network with an SNMP trap watching program.

Some network problems are intermittent, and you need to be notified when they occur. Installing an SNMP management program that monitors connectivity by pinging the

devices often and an SNMP trap-watching program can be a real eye opener, because it exposes intermittent outages that might otherwise go unnoticed. In addition, since you can keep track of the time as you watch the devices lose connectivity or regain connectivity, the time intervals could give you a clue about what the problem is.

I do not recommend that you use these tools all the time and for every scenario, but you should use these tools in at least a few of your practice scenarios. They might be particularly helpful in a scenario involving complex route redistribution, or tricky access-lists that could threaten to break connectivity. You also get experience configuring SNMP on your equipment.

## **Always be conscious of what router you are working on, and be careful with placement of access lists**

Configuring the right thing on the wrong router was the most frustrating mistake I made in practice labs. Do not make this mistake in the exam.

Similarly, you need carefully to consider where to place access-lists. Think about the flow of traffic, the direction in which different types of traffic will be moving, which are the inside interfaces and outside interfaces, etc. Keep in mind the source and destination UDP or TCP ports, if applicable. Also, remember not to break your network with an access-list that does not let through routing protocol traffic or other overhead traffic essential to keep the network running.

## **Look out for cabling problems in your home practice lab**

Because the scenarios in this book require no re-cabling between them, you may not run into many cabling problems, and any should be resolved during the first scenario. However, in the course of your study, you may well do a lot of cabling. I had some problems from one bent pin on back-to-back DTE/DCE cables. Sometimes you will simply connect one end of a cable to the wrong interface. If you using a Cisco 2514, be aware of which interface is e0 and which is e1.

Too many of us use UTP cables in our lab that we would never use in a production network. If a Category 5 UTP cable has an RJ-45 plug that is broken so that the cable can slip out of a router or switch, replace it.

## **Be conscious of point values during the exam**

This last piece of advice involves the exam itself, rather than preparation for it. The goal of the test is to get a certain number of points. If you have a three point item that no other task depends on, and you are not sure how to implement it, be prepared to jettison it and concentrate on configuring other tasks and testing your configuration. If you manage to figure out how to do that three point task, but it takes you so long that

you have no time to check your other work, you may regret it. Just as importantly, if you do not fully understand the material covered by the three-point task, you could misconfigure it and break your network so that you lose many more than three points.

Concentrate on doing what you know you can do, then go back and work on figuring out the other tasks. You will be more relaxed and confident doing research using the CD if you know you have already completed many tasks correctly. If you have gotten comfortable with the CD, figuring out how to do something new could be feasible, but you will need to be wary of problems and test carefully to insure that you do not break anything.

# Chapter 2

## CCIE Practice Lab Equipment

## Difficult decisions to make regarding lab equipment

When preparing for the CCIE Lab exam, you have some hard decisions to make regarding equipment. To build a perfect CCIE preparation lab means having at least 12 routers, two with ISDN BRI interfaces, and two with ATM interfaces, and two with voice modules, an ISDN simulator, a Cisco LightStream 1010 ATM switch or functional equivalent, and two Catalyst 3550-EMI switches. You would want at least one of the routers to have a 10/100 interface for trunking. To own this perfect lab is simply out of the question for most of us. Perhaps some of us could afford to finance such a lab for a few months and then sell it immediately after passing the Lab exam, but CCIE preparation can take well over a year.

Compromise is necessary. For the early to intermediate stages of CCIE lab preparation, my advice would be to buy (or rent) seven or eight routers, one with at least four serial interfaces so it can serve as a frame relay switch, and one VLAN-capable switch. Many of the hardest lessons CCIE candidates have to learn involve routing, and can be learned in a lab consisting only of routers with one VLAN-capable layer-two switch.

Cisco 2500 and 4500 and 4700 routers are very inexpensive right now. Even though Cisco token ring routers have gotten very cheap, I would not waste my time with token ring now, since Ethernet has dropped so much in price and you will want to be able to connect your routers to an Ethernet switch to practice configuring VLANs and other layer-two switch functions.

All the scenarios in this book are designed to be done with a frame switch that has only four serial interfaces, but some commercial scenarios require a frame switch with more than four serial interfaces, such as a 2522 or a 4000-series router with a couple of NP-4T interfaces. I should warn the reader that Cisco 4000-series routers are loud. They are not jet engine loud like the old Cisco AGS routers, but one 4500-M can be louder than seven 2500s and a Catalyst 3550 put together. The advantage of using a 4000-series router as your frame switch is that you can have eight full T1 synchronous serial interfaces, while a 2522 will have more serial interfaces (ten), but eight will be low-speed serial (128 kbps max). By the way, you do not need a lot of flash or DRAM, or a recent IOS version, in your frame switch unless you also plan to use it as a router in your scenarios.

If you can afford more than a basic lab of six or seven routers and a frame switch, I would add ISDN and one 3550-EMI switch (with routing capability), which is what I have done for this book. ISDN is a particularly slippery technology, because it is so easy to think you have it working just right when there are latent problems, so it is beneficial to get as much practice with it as often as possible, and as early as possible. ISDN may never be your friend, but it can be the enemy you know well. It is worth noting that Cisco seems to have made ISDN work a bit more smoothly in IOS 12.2. The Catalyst 3550 is such an interesting and powerful device, with so much potential to alter the entire shape of CCIE lab logical diagrams that it would certainly help to have at least one in your lab.

I purchased voice equipment for my lab, but I am not sure if it was necessary. Voice over IP, standing alone, is not that complicated a subject. Voice over IP is a subset of the larger and more difficult subject of IP telephony, which involves Voice over IP, the Cisco Call Manager, the Cisco Unity server, IP phones, QOS, and other matters. It is useful to

have voice equipment in the home lab to be able to test voice quality in addition to mere connectivity. While testing voice quality could help you in testing QOS, it is not going to work for every kind of QOS method, since some frame relay QOS methods are not going to be completely functional with a Cisco router acting as a frame relay switch. The ability to test voice quality does not justify spending well over \$1000 for voice equipment, not to mention the additional cost using 2600, 3600 or 3700 series routers that can take the voice modules. On the other hand, having used the voice equipment extensively in my home lab did give me a certain comfort and familiarity with voice that may have permitted me to concentrate on other technologies. You can learn voice using 3810s, but these are not on the equipment list and are not exactly like 2600s and 3600s.

As for ATM, I relied on rental labs and I think that is a good approach for most people who are not rich. Let me run through my thinking on this subject, since the idea of studying for the CCIE Lab without ATM in our home labs has troubled many CCIE candidates. To study ATM in your home lab, you should get a Lightstream 1010 switch and ATM modules that will work on a router that will run a recent version of IOS. Do not waste your time with the older Lightstream 100, whose features are so limited that you might as well hook your ATM interfaces to each other back-to-back. There are other ATM switches made by other vendors, but you will need to research the precise features offered by each switch, and compare them to the features that you see used on Lightstream 1010s in practice scenarios. On Ebay, the Lightstream 1010 switches are becoming a bit more affordable, and are being offered for around \$1,500.

You can get Cisco 7000-series ATM modules and Cisco 7000-series routers fairly cheap -- less than \$1,000 for two routers with ATM modules. The problem is that these routers will not run a recent version of IOS unless you add a special RSP7000 card, which can cost as much the router. There are two kinds of Cisco ATM syntax, an old style of ATM syntax and a new style of ATM syntax. With older 7000 series routers, you will only be able to practice using the old syntax. I just cannot see the sense in buying two huge, loud routers, just to practice ATM using the old syntax. A better option would be 7500 series routers, which have come down a lot in price. The shipping will normally be a big part of the cost of any of the 7000 or 7500 series routers.

Another option is to purchase 4500 or 4700 routers with ATM multimode or single-mode modules. The routers are very cheap, and the modules are currently available for less than \$600 each on ebay. If you can find an ATM switch with a DS-3 interface, and a DS-3 ATM interface for a 4500 or 4700 router, you could save a lot of money. The prices and availability of used ATM equipment vary widely from week to week, so do your own careful research and you may find that some option has become much cheaper than when this book went to press. Also, look into issues regarding the availability of new IOS versions for whatever router you seek to use for ATM. The 3600-series and 3700-series will be great in this respect, but they and their ATM modules are expensive.

Over the long term, you will want to own or otherwise gain access to lab equipment that fits the commercial lab scenarios you will be doing. If you can afford to do so, buy first-rate commercial practice scenarios and then buy or rent equipment that can be easily adapted to those scenarios.

I bought two routers at full "used" price of around \$800 each when I started studying for the CCNA in early 2000, but after that, I tended to buy equipment as it became available at an affordable price. This flexible "bargain hunter" approach may work out fine if you start buying equipment long before you pass the CCIE Written Exam, as I did, but once you pass the Written Exam, you should get the equipment you need without delay. I learned many important real-world lessons and gained a lot of hands-on experience from building my own lab and having to make do with the wrong equipment, but many readers will not need this sort of experience as much as I did. My time-consuming detour with fiber-optic blades and fiber converters for my Catalyst 5000 gave me good material to talk about during an interview when I had done well on a quiz given by the hiring manager. I was going to get a job with the company, but staying and talking with the hiring manager about my home lab experience with fiber-optic equipment may have gotten me a higher position than I would otherwise have gotten based on my work experience at that time.

For too long, I practiced ISDN using a Cisco 2504 at one end of the ISDN link and a 2501 attached to an Adtran terminal adapter at the other. This was a mistake. ISDN is hard enough to learn and implement, without adding another layer of complexity. Now that routers with ISDN are much cheaper, I recommend using routers with built-in ISDN interfaces.

As I have said before, it is a good idea to buy or rent lab equipment that goes well with your practice scenarios. I did not do this, and it hindered my progress. Research the relative quality CCIE Workbooks, and buy at least one of them. Then buy or rent equipment to match the topology of those labs as closely as possible. Unless you do this, you will end up adapting your lab to handle the scenarios, and sometimes this can change the way the scenario will work. If you want flexibility to handle many practice scenarios, buy routers like 2514s that have two Ethernet interfaces. Modular 2600 series routers are extremely flexible, but also expensive. I really enjoyed moving the ISDN or WIC-1T interfaces around as needed to fit a given scenario, but this was a costly luxury.

Some people have the idea that one needs a protocol analyzer in their CCIE practice lab, or that you need to have lots of experience with protocol analyzers to succeed in the CCIE Lab. It is not necessary to have one in your lab, but it cannot hurt to have used one, so you know what kinds of things hosts and routers and switches are saying to each other. Watching the SMB traffic for Microsoft file sharing, or watching the STP traffic among switches, can be instructive. Sometimes, it is good to watch the traffic simply to see how chatty some protocols are.

As a practical matter, you can usually see the traffic that you need to see in your lab network by careful use of IOS debugging tools. In short, experience with a protocol analyzer is helpful but not essential.

Please remember not to put a protocol analyzer on a production network unless you have a legal right to do it. You could easily find yourself in violation of a computer crime law or privacy law if you decide on your own to put a protocol analyzer on your company network.



# Chapter Three

## Scenario One

## Staging of Scenarios

The e0 interface of each router, except R7, will be connected to the Catalyst 3550 as follows:

Router	Interface	Switch port # fastethernet0/x
R1	e0	1
R2	e0	2
R3	e0	3
R4	e0	4
R5	e0	5
R6	e0	6

This cabling is straightforward, yet permits dramatic changes in lab topologies without any re-cabling. At times, we will connect additional devices to the switch for testing purposes.

These labs are geared to IOS version 12.2, the version that is used in the CCIE Routing and Switching Lab as this book goes to press. In all the scenarios, only one physical interface on R7 will be needed, so the reader can use a "frame switch" router or terminal server as R7.

I used the following frame relay switch configuration script on my Cisco 4500-M router:

```
frame#sh run
Building configuration...
```

```
Current configuration:
!
version 11.0
service udp-small-servers
service tcp-small-servers
!
hostname frame
!
!
frame-relay switching
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
```

```

no ip address
encapsulation frame-relay
clockrate 1300000
frame-relay intf-type dce
frame-relay route 102 interface Serial1 201
frame-relay route 103 interface Serial2 301
frame-relay route 104 interface Serial3 401
!
interface Serial1
no ip address
encapsulation frame-relay
clockrate 1300000
frame-relay intf-type dce
frame-relay route 201 interface Serial0 102
frame-relay route 203 interface Serial2 302
frame-relay route 204 interface Serial3 402
!
interface Serial2
no ip address
encapsulation frame-relay
clockrate 1300000
frame-relay intf-type dce
frame-relay route 301 interface Serial0 103
frame-relay route 302 interface Serial1 203
frame-relay route 304 interface Serial3 403
!
interface Serial3
no ip address
encapsulation frame-relay
clockrate 1300000
frame-relay intf-type dce
frame-relay route 401 interface Serial0 104
frame-relay route 402 interface Serial1 204
frame-relay route 403 interface Serial2 304
!
!
line con 0
line aux 0
transport input all
line vty 0 4
login
!
end

frame#

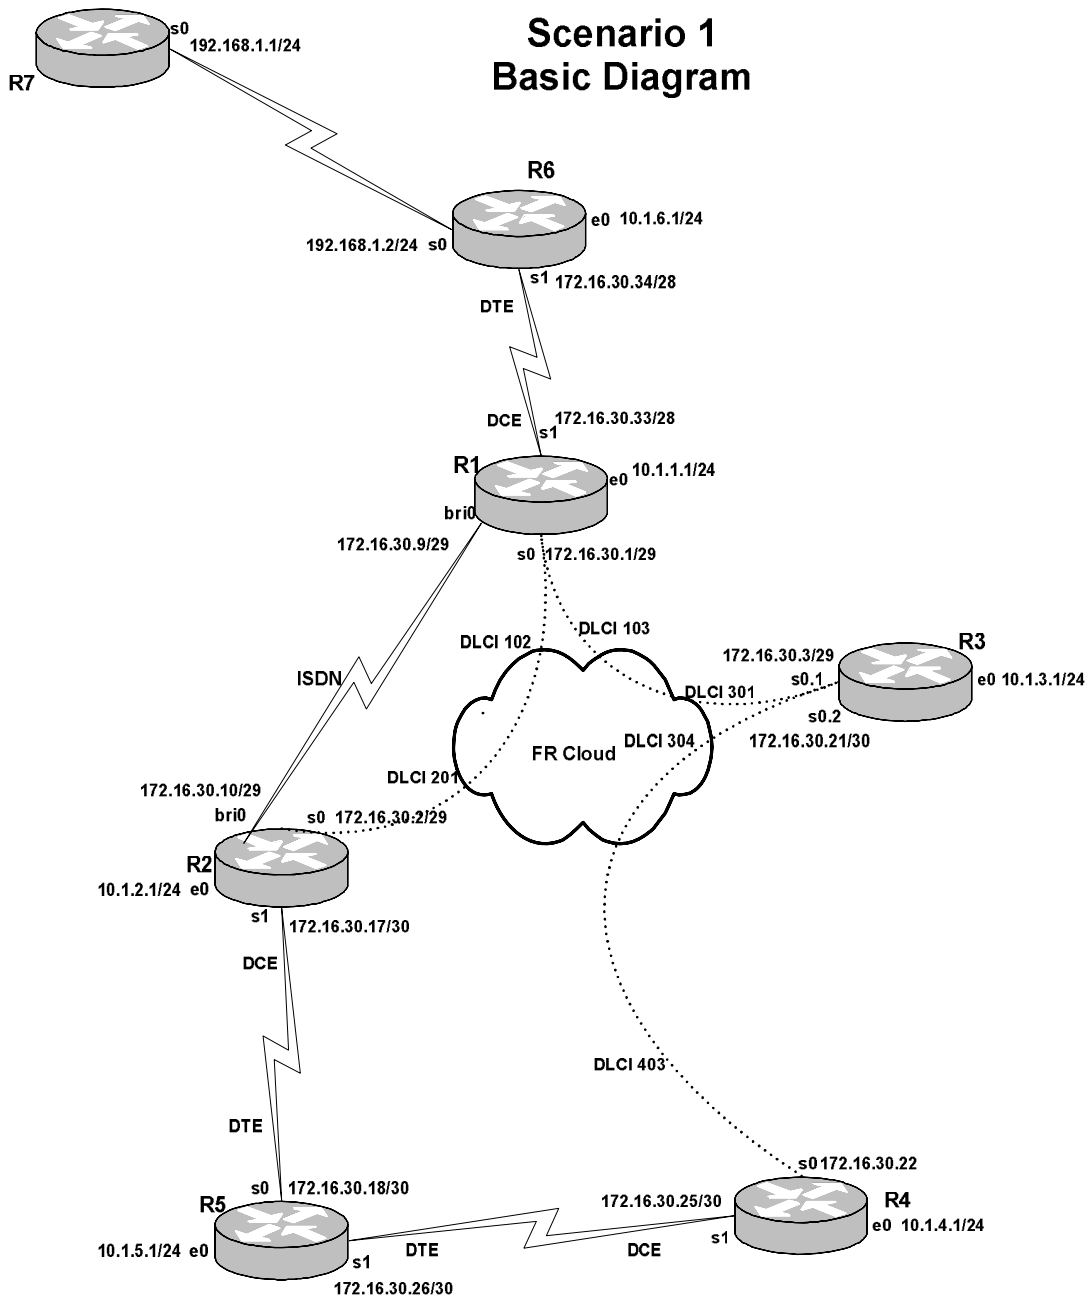
```

As you can see from the "version 11.0" above, if the router is only being used as a frame relay switch, an old IOS version will do fine.

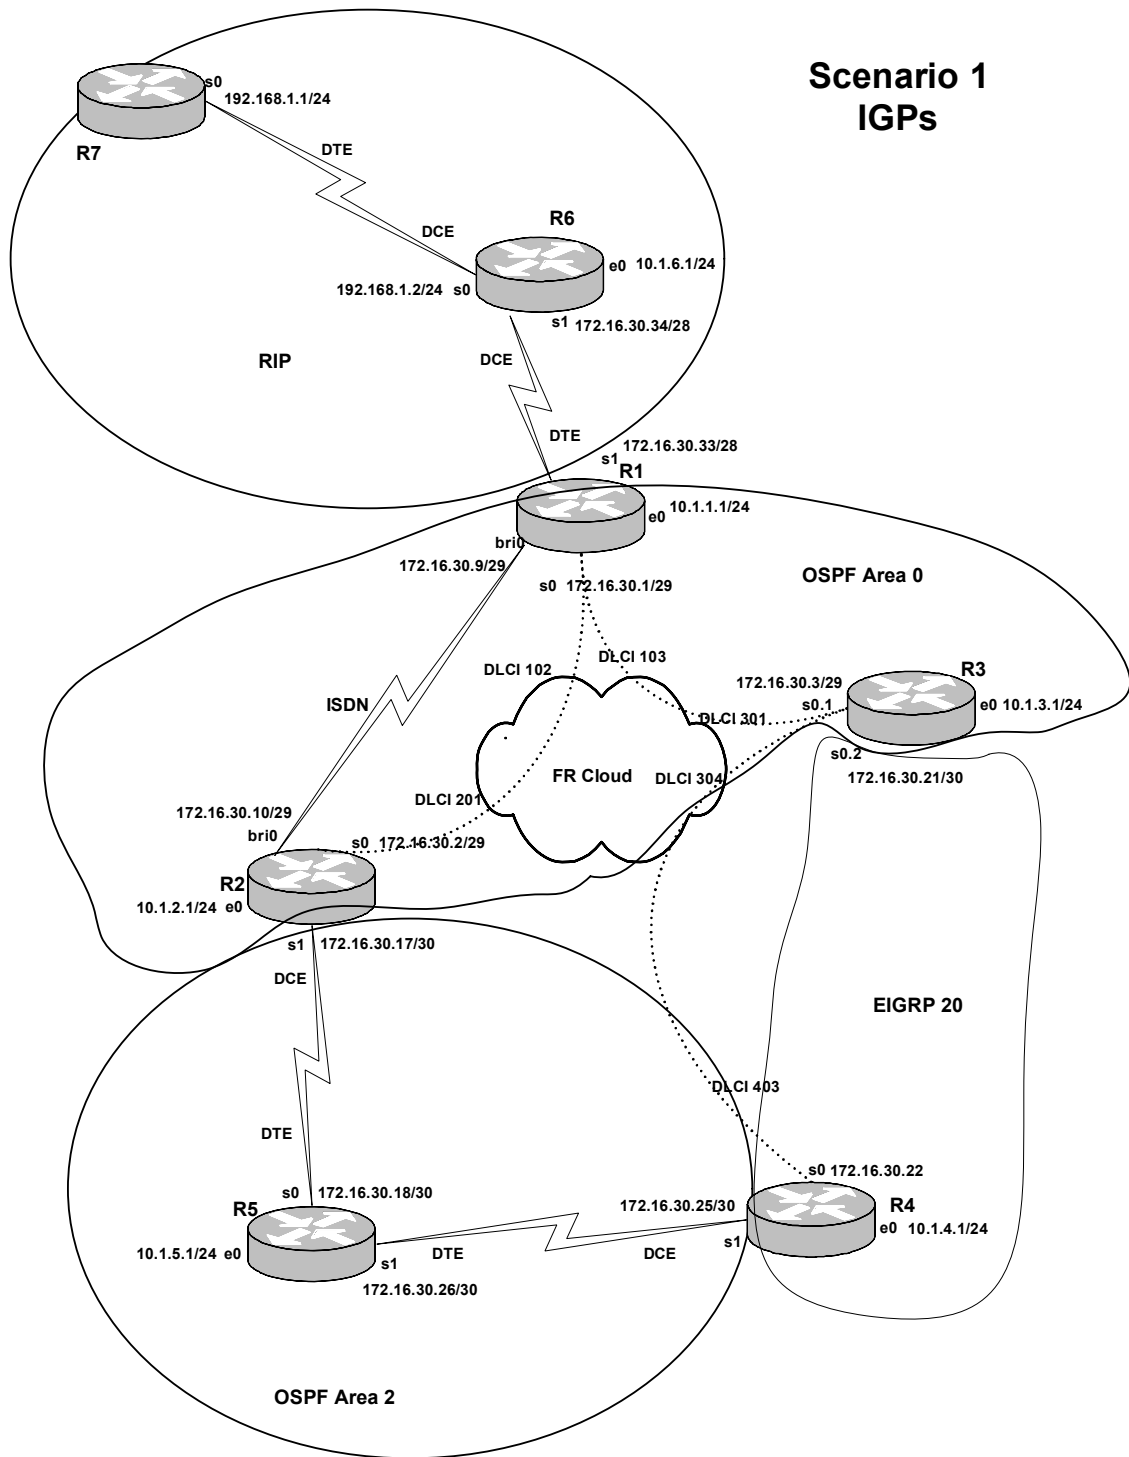
In all scenarios besides Scenario Three, there is more than one logical diagram. Where I refer to "the diagram," please refer to the appropriate diagram, which will be clear from the context of the task. For readability, I refer to the routers as R1, R2, R3, and so on, in

the text, but they are r1, r2, r3, and so on, in the router configuration scripts and in the terminal server.

Scenario One has a few features that give it a "real-world network" look and feel. The routers connect to each other using their serial interfaces and the ethernet interfaces are reserved for servers and end-users. I even refrained from having a NAT task in this scenario so that you will see in each routing table routes to all of the Ethernet subnets on all routers except R7. The only unrealistic aspect is that too many routing protocols are running in this small network. All scenarios after this one belong to the alternate universe of CCIE practice labs.



### Scenario 1 IGPs



## Tasks

1. Configure R1 as the hub in a hub-and-spoke frame relay network with R2 and R3. Use the DLCIs and IP addresses shown in the diagram. Only R3 can use subinterfaces. Make sure each frame relay interface can be pinged from any other interface on the same router.
2. Configure the Ethernet interfaces, and configure the VLANs on the Cat 3550 as shown in the table below.

Router	Interface	VLAN #	Switch port # (fastether 0/x)
R1	e0	10	1
R2	e0	20	2
R3	e0	30	3
R4	e0	40	4
R5	e0	50	5
R6	e0	60	6

3. Configure the switch to dramatically shorten the Spanning Tree Protocol delay on each port when a device is booted up or connected to the port, or when the link is otherwise reset.
4. Configure a frame relay point-to-point link between R3 and R4 as shown in the diagram.
5. Configure PPP links between R7 and R6, R6 and R1, R4 and R5, and R5 and R2. On the PPP link between R7 and R6, use CHAP authentication, with "customer" as the CHAP hostname and "ISP" as the password on R6, and "provider" as the CHAP hostname on R7.
6. Configure R1 e0 and s0 in OSPF area 0, as well as R2 e0 and s0, and R3 e0 and s0.1. Implement clear text authentication in area 0 using the password "area0."
7. Configure R2 s1, R4 s1, and all of R5 in OSPF area 2.
8. Configure a demand-circuit ISDN link between R1 and R2 that participates in OSPF area 0.
9. Configure R4 s0 and e0, and R3 s0.2 in EIGRP AS 20.
10. R6 e0, s0 and s1, R1 s1, and R7 s0 all participate in RIP routing.
11. Configure route redistribution throughout the network, using route tagging and route maps to prevent route feedback. Hypothetical users on all LANs in the

network should be able to ping all other LANs shown here, even if the frame relay link between R3 and R4 fails. Make any adjustments to administrative distance necessary to make the VPN in Task 13 work as specified.

12. Configure R5 as a DHCP server for devices on the e0 LAN interface, handing out IP addresses from the range 10.1.5.10/24-10.1.5.62/24. Set R5 e0 as the default gateway and 10.1.4.7 and 10.1.1.7 as the DNS servers.
13. Configure a VPN for traffic between R4 e0 and R1 e0. Use the pre-shared key "cisco" and DES encryption with DES authentication using an MD5 hash. Configure IPSEC using s0 as the outgoing interface on R1 and s0 as the outgoing interface on R4. The VPN does not have to be able to make use of redundant network paths.
14. Configure DLSW+ between R1 e0 and R5 e0, using a reliable transport.
15. Configure Low Latency Queueing (Class-Based Weighted Fair Queueing) on the serial link between R6 and R1, giving 100 kbps reserved bandwidth to all IP traffic with IP precedence of 5, and giving that traffic top priority. On the same link, configure multilink PPP fragmentation and interleaving so that packets will not have greater than 10 ms serialization delay. Configure RTP header compression and TCP header compression on the same link.



## Solution Configuration Scripts

### R1

```
r1#sh runn
Building configuration...
```

```
Current configuration : 3697 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r1
!
logging rate-limit console 10 except errors
!
username r2 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
class-map match-all prec5
  match ip precedence 5
!
!
policy-map llq
  class prec5
    priority 100
  class class-default
    fair-queue
!
isdn switch-type basic-ni
!
crypto isakmp policy 1
  hash md5
  authentication pre-share
  lifetime 3600
crypto isakmp key cisco address 172.16.30.22
!
!
crypto ipsec transform-set myset esp-des esp-md5-hmac
!
crypto map mymap 10 ipsec-isakmp
  set peer 172.16.30.22
  set transform-set myset
```

```

    match address 101
    !
    !
    dlsw local-peer peer-id 10.1.1.1
    dlsw remote-peer 0 tcp 10.1.5.1
    dlsw bridge-group 1
    !
    !
    interface Multilink1
    description multilink that works with serial 1
    ip address 172.16.30.33 255.255.255.240
    ip tcp header-compression iphc-format
    no ip mroute-cache
    service-policy output llq
    no peer neighbor-route
    no cdp enable
    ppp multilink
    ppp multilink fragment-delay 10
    multilink-group 1
    ip rtp header-compression iphc-format
    !
    interface Ethernet0
    ip address 10.1.1.1 255.255.255.0
    bridge-group 1
    !
    interface Serial0
    ip address 172.16.30.1 255.255.255.248
    encapsulation frame-relay
    ip ospf authentication
    ip ospf authentication-key area0
    ip ospf network point-to-multipoint
    no fair-queue
    frame-relay map ip 172.16.30.1 102 broadcast
    frame-relay map ip 172.16.30.2 102 broadcast
    frame-relay map ip 172.16.30.3 103 broadcast
    no frame-relay inverse-arp
    frame-relay lmi-type cisco
    crypto map mymap
    !
    interface Serial1
    no ip address
    encapsulation ppp
    no peer neighbor-route
    no fair-queue
    ppp multilink
    multilink-group 1
    !
    interface BRI0
    ip address 172.16.30.9 255.255.255.248
    encapsulation ppp
    ip ospf authentication
    ip ospf authentication-key area0
    ip ospf demand-circuit
    dialer idle-timeout 30

```

```

dialer map ip 172.16.30.10 name r2 broadcast 4082222222
dialer-group 1
isdn switch-type basic-ni
isdn spid1 4081111111 4081111111
isdn spid2 4081111112 4081111111
cdapi buffers regular 0
cdapi buffers raw 0
cdapi buffers large 0
no peer neighbor-route
ppp authentication chap
!
router ospf 64
log-adjacency-changes
area 0 authentication
redistribute rip metric-type 1 subnets route-map blocktag64
network 10.1.1.1 0.0.0.0 area 0
network 172.16.30.1 0.0.0.0 area 0
network 172.16.30.9 0.0.0.0 area 0
!
router rip
version 2
redistribute ospf 64 metric 3 route-map blocktag100
passive-interface BRI0
network 172.16.0.0
no auto-summary
!
ip kerberos source-interface any
ip classless
no ip http server
!
access-list 101 permit ip 10.1.1.0 0.0.0.255 10.1.4.0 0.0.0.255
dialer-list 1 protocol ip permit
route-map blocktag64 permit 10
  match tag 0
  set tag 100
!
route-map blocktag64 deny 20
  match tag 64
!
route-map blocktag100 deny 10
  match tag 100
!
route-map blocktag100 permit 20
  match tag 20
  set tag 20
!
route-map blocktag100 permit 30
  set tag 64
!
!
bridge 1 protocol ieee
!
line con 0
  exec-timeout 0 0

```

```

transport input none
line aux 0
line vty 0 4
  login
!
end

r1#

```

## R2

```

r2#sh runn
Building configuration...

Current configuration : 2244 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname r2
!
logging rate-limit console 10 except errors
!
username r1 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Ethernet0
 ip address 10.1.2.1 255.255.255.0
!
interface Serial0
 ip address 172.16.30.2 255.255.255.248
 encapsulation frame-relay
 ip ospf authentication
 ip ospf authentication-key area0
 ip ospf network point-to-multipoint
 frame-relay map ip 172.16.30.1 201 broadcast
 frame-relay map ip 172.16.30.2 201 broadcast
 frame-relay map ip 172.16.30.3 201 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco

```

```

!
interface Serial1
 ip address 172.16.30.17 255.255.255.252
 encapsulation ppp
 clockrate 1300000
!
interface BRI0
 ip address 172.16.30.10 255.255.255.248
 encapsulation ppp
 ip ospf authentication
 ip ospf authentication-key area0
 dialer idle-timeout 30
 dialer map ip 172.16.30.9 name r1 broadcast 4081111111
 dialer-group 1
 isdn switch-type basic-ni
 isdn spid1 40822222221 4082222222
 isdn spid2 40822222222 4082222222
 cdapi buffers regular 0
 cdapi buffers raw 0
 cdapi buffers large 0
 no peer neighbor-route
 ppp authentication chap
!
router ospf 64
 log-adjacency-changes
 area 0 authentication
 network 10.1.2.1 0.0.0.0 area 0
 network 172.16.30.2 0.0.0.0 area 0
 network 172.16.30.10 0.0.0.0 area 0
 network 172.16.30.17 0.0.0.0 area 2
!
 ip kerberos source-interface any
 no ip classless
 no ip http server
!
 dialer-list 1 protocol ip permit
!
!
 line con 0
   exec-timeout 0 0
   transport input none
 line aux 0
 line vty 0 4
   login
!
end

r2#

```

### R3

```
r3#sh runn
```

Building configuration...

```
Current configuration : 2428 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r3
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Ethernet0
 ip address 10.1.3.1 255.255.255.0
 no ip route-cache
 no ip mroute-cache
!
interface Serial0
 no ip address
 encapsulation frame-relay
 no ip route-cache
 no ip mroute-cache
 frame-relay lmi-type cisco
!
interface Serial0.1 multipoint
 ip address 172.16.30.3 255.255.255.248
 no ip route-cache
 ip ospf authentication
 ip ospf authentication-key area0
 ip ospf network point-to-multipoint
 no ip mroute-cache
 frame-relay map ip 172.16.30.1 301 broadcast
 frame-relay map ip 172.16.30.2 301 broadcast
 frame-relay map ip 172.16.30.3 301 broadcast
 no frame-relay inverse-arp
!
interface Serial0.2 point-to-point
 ip address 172.16.30.21 255.255.255.252
 no ip route-cache
 no ip mroute-cache
 frame-relay interface-dlci 304
!
interface Serial1
```

```

no ip address
no ip route-cache
no ip mroute-cache
shutdown
!
router eigrp 20
 redistribute ospf 64 metric 1500 20 255 1 1500 route-map otoa
 network 172.16.30.20 0.0.0.3
 no auto-summary
 no eigrp log-neighbor-changes
!
router ospf 64
 log-adjacency-changes
 area 0 authentication
 redistribute eigrp 20 metric-type 1 subnets route-map blocktag64
 network 10.1.3.0 0.0.0.255 area 0
 network 172.16.30.3 0.0.0.0 area 0
!
ip kerberos source-interface any
ip classless
no ip http server
!
access-list 50 permit 10.1.3.0 0.0.0.255
route-map otoa deny 10
 match tag 20
!
route-map otoa permit 20
 match tag 100
!
route-map otoa permit 30
 set tag 64
!
route-map blocktag64 deny 10
 match tag 64
!
route-map blocktag64 deny 20
 match tag 100
!
route-map blocktag64 permit 30
 set tag 20
!
!
line con 0
 exec-timeout 0 0
 transport input none
line aux 0
 transport input all
line vty 0 4
 login
!
end

r3#

```

## R4

```
r4#sh runn
Building configuration...

Current configuration : 2390 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname r4
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
crypto isakmp policy 1
  hash md5
  authentication pre-share
  lifetime 3600
crypto isakmp key cisco address 172.16.30.1
!
!
crypto ipsec transform-set myset esp-des esp-md5-hmac
!
crypto map mymap 10 ipsec-isakmp
  set peer 172.16.30.1
  set transform-set myset
  match address 101
!
!
!
!
interface Ethernet0
  ip address 10.1.4.1 255.255.255.0
!
interface Serial0
  ip address 172.16.30.22 255.255.255.252
  encapsulation frame-relay
  ip split-horizon
  no fair-queue
  frame-relay map ip 172.16.30.21 403 broadcast
  no frame-relay inverse-arp
  frame-relay lmi-type cisco
```



```

crypto map mymap
!
interface Serial1
ip address 172.16.30.25 255.255.255.252
encapsulation ppp
no peer neighbor-route
clockrate 1300000
!
interface TokenRing0
no ip address
shutdown
!
router eigrp 20
redistribute ospf 64 metric 1500 20 255 1 1500 route-map otoo
network 10.1.4.0 0.0.0.255
network 172.16.30.20 0.0.0.3
distance eigrp 70 105
no auto-summary
no eigrp log-neighbor-changes
!
router ospf 64
log-adjacency-changes
redistribute eigrp 20 metric-type 1 subnets route-map etoo
network 172.16.30.24 0.0.0.3 area 2
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 101 permit ip 10.1.4.0 0.0.0.255 10.1.1.0 0.0.0.255
route-map etoo deny 10
match tag 64
!
route-map etoo deny 20
match tag 100
!
route-map etoo permit 30
set tag 20
!
route-map otoo deny 10
match tag 20
!
route-map otoo permit 20
match tag 100
!
route-map otoo permit 30
set tag 64
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
transport input all

```

```
line vty 0 4
  login
  !
end

r4#
```

## R5

```
r5#sh runn
Building configuration...
```

```
Current configuration : 1866 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname r5
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
ip dhcp excluded-address 10.1.5.1 10.1.5.9
ip dhcp excluded-address 10.1.5.63 10.1.5.255
!
ip dhcp pool dhcppool
  network 10.1.5.0 255.255.255.0
  default-router 10.1.5.1
  dns-server 10.1.4.7 10.1.1.7
!
no ip dhcp-client network-discovery
!
!
dlsw local-peer peer-id 10.1.5.1
dlsw remote-peer 0 tcp 10.1.1.1
dlsw bridge-group 1
!
!
interface Ethernet0
  ip address 10.1.5.1 255.255.255.0
  no ip route-cache
  no ip mroute-cache
  bridge-group 1
```

```

!
interface Serial0
 ip address 172.16.30.18 255.255.255.252
 encapsulation ppp
 no ip route-cache
 no ip mroute-cache
 no peer neighbor-route
 no fair-queue
!
interface Serial1
 ip address 172.16.30.26 255.255.255.252
 encapsulation ppp
 no ip route-cache
 no ip mroute-cache
 no peer neighbor-route
!
router ospf 64
 log-adjacency-changes
 network 10.1.5.1 0.0.0.0 area 2
 network 172.16.30.18 0.0.0.0 area 2
 network 172.16.30.26 0.0.0.0 area 2
!
 ip kerberos source-interface any
 ip classless
 no ip http server
!
!
 bridge 1 protocol ieee
!
 line con 0
  exec-timeout 0 0
  transport input none
 line aux 0
  transport input all
 line vty 0 4
  login
!
end

r5#

```

## R6

```

r6#sh runn
Building configuration...

Current configuration : 1981 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption

```

```

!
hostname r6
!
logging rate-limit console 10 except errors
!
username provider password 0 ISP
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
class-map match-all prec5
  match ip precedence 5
!
!
policy-map llq
  class prec5
    priority 100
  class class-default
    fair-queue
!
!
!
!
!
interface Multilink1
  description multilink that works with serial 1
  ip address 172.16.30.34 255.255.255.240
  ip tcp header-compression iphc-format
  no ip mroute-cache
  service-policy output llq
  no peer neighbor-route
  no cdp enable
  ppp multilink
  ppp multilink fragment-delay 10
  multilink-group 1
  ip rtp header-compression iphc-format
!
interface Ethernet0
  ip address 10.1.6.1 255.255.255.0
!
interface Ethernet1
  no ip address
  shutdown
!
interface Serial0
  ip address 192.168.1.2 255.255.255.0
  encapsulation ppp
  no peer neighbor-route
  no fair-queue
  clockrate 1300000
  ppp authentication chap
  ppp chap hostname customer

```

```

    ppp chap password 7 0726127C
    !
interface Serial1
  no ip address
  encapsulation ppp
  no fair-queue
  clockrate 1300000
  ppp multilink
  multilink-group 1
  !
router rip
  version 2
  network 10.0.0.0
  network 172.16.0.0
  network 192.168.1.0
  no auto-summary
  !
ip kerberos source-interface any
ip classless
no ip http server
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  no login
!
end

r6#

```

## R7

```

r7#sh runn
Building configuration...

Current configuration : 1444 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
service password-encryption
!
hostname r7
!
logging rate-limit console 10 except errors
!
username customer password 7 002D2036
ip subnet-zero

```

```
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Ethernet0
  no ip address
!
interface Ethernet1
  no ip address
  no ip route-cache
  no ip mroute-cache
  shutdown
!
interface Serial0
  ip address 192.168.1.1 255.255.255.0
  encapsulation ppp
  no ip route-cache
  no ip mroute-cache
  no peer neighbor-route
  ppp authentication chap
  ppp chap hostname provider
  ppp chap password 7 09657D39
!
interface Serial1
  no ip address
  no ip route-cache
  no ip mroute-cache
  shutdown
!
router rip
  version 2
  network 192.168.1.0
  no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
line con 0
  transport input none
line aux 0
  transport input all
line vty 0 4
  login
!
end

r7#
```

## Cat

```
cat#sh runn
Building configuration...

Current configuration : 2217 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname cat
!
!
ip subnet-zero
no ip domain-lookup
!
!
spanning-tree portfast default
spanning-tree extend system-id
!
!
!
interface FastEthernet0/1
  switchport access vlan 10
  no ip address
!
interface FastEthernet0/2
  switchport access vlan 20
  no ip address
!
interface FastEthernet0/3
  switchport access vlan 30
  no ip address
!
interface FastEthernet0/4
  switchport access vlan 40
  no ip address
!
interface FastEthernet0/5
  switchport access vlan 50
  no ip address
!
interface FastEthernet0/6
  switchport access vlan 60
  no ip address
!
interface FastEthernet0/7
  no ip address
!
```

```
interface FastEthernet0/8
  no ip address
!
interface FastEthernet0/9
  no ip address
!
interface FastEthernet0/10
  no ip address
!
interface FastEthernet0/11
  switchport access vlan 10
  no ip address
!
interface FastEthernet0/12
  no ip address
!
interface FastEthernet0/13
  no ip address
!
interface FastEthernet0/14
  no ip address
!
interface FastEthernet0/15
  switchport access vlan 50
  no ip address
!
interface FastEthernet0/16
  no ip address
!
interface FastEthernet0/17
  no ip address
!
interface FastEthernet0/18
  no ip address
!
interface FastEthernet0/19
  no ip address
!
interface FastEthernet0/20
  no ip address
!
interface FastEthernet0/21
  no ip address
!
interface FastEthernet0/22
  no ip address
!
interface FastEthernet0/23
  no ip address
!
interface FastEthernet0/24
  no ip address
!
interface GigabitEthernet0/1
```



```
no ip address
!  
interface GigabitEthernet0/2  
no ip address  
!  
interface Vlan1  
no ip address  
shutdown  
!  
ip classless  
ip http server  
!  
!  
line con 0  
exec-timeout 0 0  
line vty 0 4  
login  
line vty 5 15  
login  
!  
end  
  
cat#
```

## Explanation

**1. Configure R1 as the hub in a hub-and-spoke frame relay network with R2 and R3. Use the DLCIs and IP addresses shown in the diagram. Only R3 can use subinterfaces. Make sure each frame relay interface can be pinged from any other interface on the same router.**

When setting up frame-relay it is a good practice to shut down the physical serial interface before you put frame relay encapsulation on it, to avoid having DLCIs available that you may not want available. You should also disable inverse ARP before you bring up the interface to avoid erroneous mappings. Along the same lines, it is safer to configure frame relay map statements than to rely on inverse-arp.

The requirement that every frame relay interface needs to be pingable from other interfaces on the same router calls out for a "frame-relay map" statement to the local frame relay interface. Otherwise, you will be able to ping the local interface from networks across the frame relay cloud, but not from this router or other networks on this side of the frame relay cloud.

Be prepared to set the encapsulation and lmi-types to some value other than the defaults, and make sure you are using the right encapsulation and lmi-type.

As a matter of habit, I often set the OSPF network type to point-to-multipoint. This eliminates the problem of making sure that your hub wins the DR/BDR election. If you have to use broadcast or nonbroadcast network types, set the OSPF priority of the spokes to 0 to make sure the spokes are not even in the running for DR or BDR. If you have to use nonbroadcast network type, remember that you set the unicast neighbors manually using the OSPF neighbor statement.

**2. Configure the Ethernet interfaces, and configure the VLANs on the Cat 3550 as shown in the table below.**

<i>Router</i>	<i>Interface</i>	<i>VLAN #</i>	<i>Switch port # (fastether 0/x)</i>
<i>R1</i>	<i>e0</i>	<i>10</i>	<i>1</i>
<i>R2</i>	<i>e0</i>	<i>20</i>	<i>2</i>
<i>R3</i>	<i>e0</i>	<i>30</i>	<i>3</i>
<i>R4</i>	<i>e0</i>	<i>40</i>	<i>4</i>
<i>R5</i>	<i>e0</i>	<i>50</i>	<i>5</i>
<i>R6</i>	<i>e0</i>	<i>60</i>	<i>6</i>

This task consists of setting up the VLANs and bringing up the ethernet interfaces on the routers. The switch's ethernet interfaces come on by default when an ethernet device is attached.

All that is necessary on each router's ethernet interface is to put the ip address on and put "no shut" on the interface to bring it up.

### **3. Configure the switch to dramatically shorten the Spanning Tree Protocol delay on each port when a device is booted up or connected to the port, or when the link is otherwise reset.**

The Catalyst 3550 permits you to enable the spanning tree portfast feature on all interfaces by default by simply typing "spanning-tree portfast default." This feature is wonderful because it allows you to mitigate the delay of spanning tree while removing any incentive to disable spanning tree entirely, which would be a very dangerous step in a production network. Spanning tree protocol is a good idea to keep in place in a production network because extra switches and hubs and redundant links find their way into your wiring closets and user workspaces, and can kill your network.

On an access-layer switch in a wiring closet, you might well want spanning tree portfast enabled on all ports connected to end-user equipment so that your users do not have to wait over 50 seconds to get DHCP addresses and access to other network services.

In a home CCIE lab, you will often connect routers to each other using ethernet links, and your routers should not have to wait 50 seconds to hear from a neighbor every time an ethernet link is reset.

### **4. Configure a frame relay point-to-point link between R3 and R4 as shown in the diagram.**

In practice labs, I usually saw an additional point-to-point frame relay subinterface on the hub router in a hub-and-spoke topology. I wanted to do something a little different and have a frame relay point-to-point link originate at a spoke router. Since it is a point-to-point subinterface on R3, you simply set the DLCI manually using "frame-relay interface-dlci 304."

### **5. Configure PPP links between R7 and R6, R6 and R1, R4 and R5, and R5 and R2. On the PPP link between R7 and R6, use CHAP authentication, with "customer" as the CHAP hostname and "ISP" as the password on R6, and "provider" as the CHAP hostname on R7.**

We set the CHAP hostname to be different from the global hostname for the router by using "ppp chap hostname customer" on R6 and "ppp chap hostname provider" on R7. Since the CHAP password is "ISP" on R6, we know it will be the same on R7, since CHAP passwords are always the same on both sides of the link.

**6. Configure R1 e0 and s0 in OSPF area 0, as well as R2 e0 and s0, and R3 e0 and s0.1. Implement clear text authentication in area 0 using the password "area0."**

For OSPF, you may want to pick a process number and stick with it where possible. I stick with process number 64 where I can. It is simpler to use the wildcard mask 0.0.0.0 for OSPF "network" statements than to calculate the correct wildcard mask for each "network" statement. Do not try this with EIGRP, though, and be careful in calculating your wildcard masks in EIGRP "network" statements.

**7. Configure R2 s1, R4 s1, and all of R5 in OSPF area 2.**

No explanation needed.

**8. Configure a demand-circuit ISDN link between R1 and R2 that participates in OSPF area 0.**

Configuring ISDN is pretty straightforward. Nevertheless, when doing it from memory, it is easy to forget some minor but crucial step, such as setting the ISDN switch-type, which can now be done in global configuration mode or on the interface, or specifying interesting traffic using a dialer-list and then applying that list to the interface using the dialer-group statement. Be careful typing the spids and dial numbers.

Get comfortable using "debug isdn status" and "debug ppp authentication." When specifying the username and password in global configuration mode, the username is that of the remote router, not the local router, and the password should be the same on both ends. I remember it as "username remote password cisco." I sure hope no one is using the password "cisco" in production networks in circumstances where authentication is actually needed for security purposes.

If your calls go through for a second but are then dropped, use "debug ppp authentication" and look for an authentication problem. Watch out for passwords that contain spaces at the end produced when one types a space and then uses the IOS interactive help feature to see what can be typed next. This "extra space" problem also crops up with passwords used for routing protocol authentication.

A big problem with OSPF demand circuit is that the link may come up when you are not looking, or when you have disabled logging to the console. I often disable logging to the console when I run into problems, and then fail to turn it back on. Typing "sh logg" on the appropriate router will allow you to see if the demand circuit has been quiet for a long time. If the last time the ISDN circuit came up was a few minutes after the routers were last booted, that is usually a good sign.

If you can keep things straight, you might save time by configuring ISDN on R1, then cutting and pasting the config from R1 to Wordpad or Notepad, changing certain values, then pasting the modified config to R2. If you do this, just remember that it is easy to forget to change all the values that need to be changed or include all the necessary statements, especially the dialer-list, because it appears so far down in the configuration

script, and the username and password, since these appear so far down in the configuration script.

### **9. Configure R4 s0 and e0, R3 e0, and R3 s0.2 in EIGRP AS 20.**

As a general rule, you should use "no auto-summary" whenever you configure EIGRP. An exception would be if you were asked to summarize routes in EIGRP without using a command to perform the summarization. The same general rule applies to RIP version 2 and BGP.

### **10. R6 e0, s0 and s1, R1 s1, and R7 s0 all participate in RIP routing.**

Use RIP version 2 whenever you are not required to use RIP version 1. Remember to disable auto-summarization in RIP version 2 just as you would with EIGRP. When I was preparing for the CCIE, the area that most concerned me was the interaction between VLSM and FLSM routing domains. You are lucky that IGRP is no longer on the CCIE Lab exam, and anytime you can avoid using RIP version 1, you should.

### **11. Configure route redistribution throughout the network, using route tagging and route maps to prevent route feedback. Hypothetical users on all LANs in the network should be able to ping all other LANs shown here, even if the frame relay link between R3 and R4 fails. Make any adjustments to administrative distance necessary to make the VPN in Task 13 work as specified.**

Until recently, most configuration scripts in CCIE preparation books and practice scenarios used distribute-lists to prevent route feedback. This approach had the advantage of nailing down exactly which routes will or will not be redistributed from one protocol into another. Depending on how many routes are involved, this can be a very tedious task with many opportunities for mistakes.

In Jeff Doyle's *Routing TCP/IP*, the author briefly discusses the use of route maps and route tagging during redistribution in his chapter on route maps. You tag routes as they are redistributed, and then use route-maps to prevent routes carrying that tag from being redistributed back into their original routing domain. In this scenario, I used route-maps that permitted routes to retain tags that show their original routing protocol, and doing so provides much redundancy. For example, when redistributing between OSPF and EIGRP, we let through the RIP routes carrying the tag 100, and let them keep their tag value.

Let us look at a specific RIP route, 10.1.6.0/24, to see the benefits and dangers of this approach. This route will be redistributed into OSPF, and then travel throughout OSPF and then will be redistributed into EIGRP on R4 and R3. That means that a link could fail between R5 and R2, or between R3 and R1, and there will be a path from R4 and R3 to R6 lo0. The disadvantage is that if you try to carry redundancy too far, and are not

vigilant about keeping the RIP routes from being redistributed back into OSPF from EIGRP, there will be disaster. R1 will receive OSPF external routes to R6 lo0 from the EIGRP routers R4 and R3. On R1, these OSPF routes will have a lower administrative distance than the real RIP route to R6 lo0, and will therefore be preferred. You can adjust administrative distance on R1, but that will treat only the symptom, not the disease.

The lesson here is that you cannot simply configure your tagging and route-maps and then sit back and trust everything to work just fine. Not only are there the substantive issues such as the one shown here. There is the constant danger of typos in route-maps, which is similar to the danger of typos in authentication passwords. The mistakes can be very subtle.

Then you have the subtle, unexpected glitches. I found that on R1, the route-map for letting RIP routes into OSPF had to be a little different from the others. RIP routes carrying no tag yet would match whatever the first match statement in the route map was. If it was a deny statement, any RIP route would match it (no matter what match value was used) and be blocked. The solution was to match the tag value of zero, and then set the tag value for these routes to be 100 as they were redistributed into OSPF.

Another potential issue is that EIGRP routes would lose their tag when being redistributed into RIP from OSPF on R1, so I had the route-map match tag 20 and then set tag 20, to maintain the tag value of 20. This may not have caused any problems in this scenario, but if the scenario had involved any redistribution out of RIP on another router in the RIP domain, the tag value could have been very important.

It is important to remember the obvious—that routes will not have the tag values you give them until you redistribute the routes from one routing protocol into another and set tag values. Sometimes I would forget this fact and expect routes coming from a particular routing protocol to have the tag I associate with that protocol before they have been redistributed into another routing protocol. You really have to keep your logic straight with tagging and route maps, but I still prefer them to distribute-lists.

Even route tagging and route-maps do not necessarily permit you to make use of all redundant paths. Blocking routes tagged as OSPF routes from being redistributed into the OSPF domain means that you will not be able to use a redundant path through the EIGRP routing domain to get from a source in the OSPF domain to a destination in the OSPF domain. Since we also prevented RIP routes from being redistributed from EIGRP to OSPF, the same limitation applies to RIP routes.

Then there is the administrative distance issue, which is conveniently flagged for you in this scenario. R4 will prefer the longer route to R1 through the OSPF domain to the shorter route through the EIGRP domain, since OSPF routes are trusted more than EIGRP external routes. To force R4 to take the path specified in the VPN task, you have to adjust the EIGRP distance by typing "distance eigrp 70 105" under "router eigrp 20" on R4. To set the distance back to default, simply type "default distance eigrp 20." You do not have the same problem on R1 at the other end of the VPN because both paths to R4's VPN peer address have the same administrative distance because they are both OSPF routes.

Be at all times aware of the impact of administrative distance on a router's decision-making. Any individual router will prefer a route from a source it trusts more over any

route learned from a source it trusts less, even if we humans can see that by any method of calculating path cost, the latter route would have a lower cost. This behavior can lead to sub-optimal routing, which may not be a fatal problem on its own, but could cause certain applications not to work.

Once you have all the redistribution configured, with the administrative distance adjustment, go to the routing table on R4 and look at the route to 10.1.1.0/24.

r4#**sho ip rou**

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 not set

```

      172.16.0.0/16 is variably subnetted, 9 subnets, 4 masks
D EX   172.16.30.32/28 [105/2223616] via 172.16.30.21, 01:00:57,
Serial0
D EX   172.16.30.2/32 [105/2223616] via 172.16.30.21, 01:00:57,
Serial0
O IA   172.16.30.3/32 [110/256] via 172.16.30.26, 01:01:05, Serial1
D EX   172.16.30.0/29 [105/2223616] via 172.16.30.21, 01:01:01,
Serial0
D EX   172.16.30.1/32 [105/2223616] via 172.16.30.21, 01:00:57,
Serial0
D EX   172.16.30.8/29 [105/2223616] via 172.16.30.21, 01:00:57,
Serial0
D EX   172.16.30.16/30 [105/2223616] via 172.16.30.21, 01:00:58,
Serial0
C      172.16.30.20/30 is directly connected, Serial0
C      172.16.30.24/30 is directly connected, Serial1
      10.0.0.0/24 is subnetted, 6 subnets
D EX   10.1.3.0 [105/2223616] via 172.16.30.21, 01:00:58, Serial0
D EX   10.1.2.0 [105/2223616] via 172.16.30.21, 01:00:58, Serial0
D EX   10.1.1.0 [105/2223616] via 172.16.30.21, 01:00:58, Serial0
D EX   10.1.6.0 [105/2223616] via 172.16.30.21, 01:00:58, Serial0
D EX   10.1.5.0 [105/2223616] via 172.16.30.21, 01:00:58, Serial0
C      10.1.4.0 is directly connected, Ethernet0
D EX 192.168.1.0/24 [105/2223616] via 172.16.30.21, 01:00:58, Serial0
r4#

```

The link to R3 will be used. Now shut down R3 s0.2, go back to R4, wait a few seconds for convergence to complete, and examine R4 s routing table:

r4#**sho ip rou**

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 not set

```

    172.16.0.0/16 is variably subnetted, 8 subnets, 4 masks
O E1   172.16.30.32/28 [110/212] via 172.16.30.26, 00:00:10, Serial1
O IA   172.16.30.2/32 [110/128] via 172.16.30.26, 00:00:10, Serial1
O IA   172.16.30.3/32 [110/256] via 172.16.30.26, 01:03:04, Serial1
O IA   172.16.30.1/32 [110/192] via 172.16.30.26, 00:00:10, Serial1
O IA   172.16.30.8/29 [110/1690] via 172.16.30.26, 00:00:10, Serial1
O      172.16.30.16/30 [110/128] via 172.16.30.26, 00:00:10, Serial1
C      172.16.30.20/30 is directly connected, Serial0
C      172.16.30.24/30 is directly connected, Serial1
    10.0.0.0/24 is subnetted, 6 subnets
O IA   10.1.3.0 [110/266] via 172.16.30.26, 00:00:11, Serial1
O IA   10.1.2.0 [110/138] via 172.16.30.26, 00:00:11, Serial1
O IA   10.1.1.0 [110/202] via 172.16.30.26, 00:00:11, Serial1
O E1   10.1.6.0 [110/212] via 172.16.30.26, 00:00:12, Serial1
O      10.1.5.0 [110/74] via 172.16.30.26, 00:00:12, Serial1
C      10.1.4.0 is directly connected, Ethernet0
O E1 192.168.1.0/24 [110/212] via 172.16.30.26, 00:00:12, Serial1
r4#

```

R4 will now take the link to R5. Now bring R3 s0.2 back up, and shut down R3 s0.1, and then look at R3's routing table.

r3#**sh ip rou**

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 not set

```

    172.16.0.0/16 is variably subnetted, 7 subnets, 4 masks
D EX   172.16.30.32/28 [170/2223616] via 172.16.30.22, 00:00:06,
Serial0.2
D EX   172.16.30.2/32 [170/2223616] via 172.16.30.22, 00:00:06,
Serial0.2

```



```

D EX 172.16.30.1/32 [170/2223616] via 172.16.30.22, 00:00:06,
Serial0.2
D EX 172.16.30.8/29 [170/2223616] via 172.16.30.22, 00:00:06,
Serial0.2
D EX 172.16.30.16/30 [170/2223616] via 172.16.30.22, 00:00:06,
Serial0.2
C 172.16.30.20/30 is directly connected, Serial0.2
D EX 172.16.30.24/30 [170/2223616] via 172.16.30.22, 00:00:08,
Serial0.2
  10.0.0.0/24 is subnetted, 6 subnets
C 10.1.3.0 is directly connected, Ethernet0
D EX 10.1.2.0 [170/2223616] via 172.16.30.22, 00:00:07, Serial0.2
D EX 10.1.1.0 [170/2223616] via 172.16.30.22, 00:00:07, Serial0.2
D EX 10.1.6.0 [170/2223616] via 172.16.30.22, 00:00:07, Serial0.2
D EX 10.1.5.0 [170/2223616] via 172.16.30.22, 00:00:08, Serial0.2
D 10.1.4.0 [90/2195456] via 172.16.30.22, 00:00:18, Serial0.2
D EX 192.168.1.0/24 [170/2223616] via 172.16.30.22, 00:00:08, Serial0.2
r3#

```

We now take the long detour through R4. The OSPF route for 10.1.1.0/24 was redistributed into EIGRP.

Prevention of route feedback means that there is a limit to all this redundancy. If the link between R2 and R5 were to fail, R5 would not be able to reach R1 by means of the redundant path through the EIGRP domain, because the OSPF route carrying a tag of 64 will not be redistributed back into the OSPF domain on R3. Let us verify this by bringing up R3 s0.1, shutting down R2 s1, and then looking for a route to 10.1.1.0/24 in R5 s routing table.

```
r5#sh ip rou
```

```

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 not set
```

```

      172.16.0.0/30 is subnetted, 2 subnets
O E1 172.16.30.20 [110/84] via 172.16.30.25, 00:00:03, Serial1
C    172.16.30.24 is directly connected, Serial1
      10.0.0.0/24 is subnetted, 2 subnets
C    10.1.5.0 is directly connected, Ethernet0
O E1 10.1.4.0 [110/84] via 172.16.30.25, 00:00:03, Serial1
r5#

```

There is no route to 10.1.5.0/24, because it is an OSPF tagged route, which our route-maps keep from passing back into the OSPF domain from EIGRP. I experimented with

letting RIP tagged routes pass through the EIGRP domain and be redistributed into OSPF, but this caused routing loops.

Some CCIE candidates have noticed that route tagging and route-maps to prevent route feedback are sometimes used where they are not needed. If you have only one router in the network on which redistribution is occurring, and everything in the network is configured correctly, then you don't need any measures to prevent route feedback, because the default behavior is that a router will not redistribute any routes from one routing domain that were redistributed into that routing domain on that router. Because of this loop-prevention behavior, if you want to redistribute routes from EIGRP into OSPF and also into RIP on the same router, you will need to do it separately. Similarly, to redistribute "connected" routes into a routing domain, you must do it manually for each routing domain. You cannot redistribute a "connected" route into OSPF, and expect it to be redistributed into EIGRP with other OSPF routes when OSPF-to-EIGRP redistribution takes place on the same router.

**12. Configure R5 as a DHCP server for devices on the e0 LAN interface, handing out IP addresses from the range 10.1.5.10/24-10.1.5.62/24. Set R5 e0 as the default gateway and 10.1.4.7 and 10.1.1.7 as the DNS servers.**

This is fairly straightforward. The part that may seem strange at first is that you define the pool's characteristics, and then define the excluded addresses outside of the pool.

**13. Configure a VPN for traffic between R4 e0 and R1 e0. Use the pre-shared key "cisco" and DES encryption with DES authentication using an MD5 hash. Configure IPSEC using s0 as the outgoing interface on R1 and s0 as the outgoing interface on R4. The VPN does not have to be able to make use of redundant network paths.**

The key to configuring IPSEC is to write down all of your IPSEC settings to prevent misconfiguration. So many things need to match on each side of the tunnel that there is a real potential for problems if you shoot from the hip with VPNs.

Remember that the access lists used for VPNs only specify what IP traffic should use the VPN, and they should be mirror images of each other. You could just as easily specify that only telnet traffic between these two LANs should be encrypted.

There is one way in which these access-lists block. When R4 sees a packet coming in from 10.1.1.0/24 that is headed to 10.1.4.0/24 and did not use the VPN (does not come in encrypted), R4 will drop that packet. Thus, the router understands this access-list to mean that any traffic from this source to that destination will use the VPN, and any traffic from that destination to this source had better be using the VPN, or will be dropped.

Because we use one outgoing WAN interface on R1 and one outgoing interface on R4 as VPN termination points, we do not get all the benefits we might otherwise get from redundant paths between R1 and R4. As we discussed above, we had to lower the administrative distance of EIGRP external routes on R4 to a value below the

administrative distance of OSPF routes, so that R4 would not prefer the OSPF route to R1 to the EIGRP route to R1. If it preferred the OSPF route, it would take a long path through R5 and R2, while the return path would go through R3. On R1, the administrative distances of the paths from R1 to R4 through the EIGRP domain and the internal OSPF path would be the same, while on R4 the paths from R4 to R1 would have different administrative distances.

To test the VPN, I did an extended ping from R1 s e0 to R4 s e0 interface:

Sending 500, 100-byte ICMP Echos to 10.1.4.1, timeout is 2 seconds:

```
04:29:04: IPSEC(sa_request): ,
  (key eng. msg.) src= 172.16.30.1, dest= 172.16.30.22,
  src_proxy= 10.1.1.0/255.255.255.0/0/0 (type=4),
  dest_proxy= 10.1.4.0/255.255.255.0/0/0 (type=4),
  protocol= ESP, transform= esp-des esp-md5-hmac ,
  lifedur= 3600s and 4608000kb,
  spi= 0x74B220CB(1957830859), conn_id= 0, keysize= 0, flags=
0x4004....
04:29:12: IPSEC(validate_proposal_request): proposal part #1,
  (key eng. msg.) dest= 172.16.30.22, src= 172.16.30.1,
  dest_proxy= 10.1.4.0/255.255.255.0/0/0 (type=4),
  src_proxy= 10.1.1.0/255.255.255.0/0/0 (type=4),
  protocol= ESP, transform= esp-des esp-md5-hmac ,
  lifedur= 0s and 0kb,
  spi= 0x0(0), conn_id= 0, keysize= 0, flags= 0x4
04:29:12: IPSEC(key_engine): got a queue event...
04:29:12: IPSEC(initialize_sas): ,
  (key eng. msg.) dest= 172.16.30.1, src= 172.16.30.22,
  dest_proxy= 10.1.1.0/255.255.255.0/0/0 (type=4),
  src_proxy= 10.1.4.0/255.255.255.0/0/0 (type=4),
  protocol= ESP, transform= esp-des esp-md5-hmac ,
  lifedur= 3600s and 4608000kb,
  spi= 0x74B220CB(1957830859), conn_id= 2000, keysize= 0, flags= 0x4
04:29:12: IPSEC(initialize_sas): ,
  (key eng. msg.) src= 172.16.30.1, dest= 172.16.30.22,
  src_proxy= 10.1.1.0/255.255.255.0/0/0 (type=4),
  dest_proxy= 10.1.4.0/255.255.255.0/0/0 (type=4),
  protocol= ESP!!!!!!!!!!!!!!!!!!!!!!!!!!!!, transform= esp-des esp-md5-hmac
,
  lifedur= 3600s and 4608000kb,
  spi= 0x179DFOAE(396226734), conn_id= 2001, keysize= 0, flags= 0x4
04:29:12: IPSEC(create_sa): sa created,
  (sa) sa_dest= 172.16.30.1, sa_prot= 50,
  sa_spi= 0x74B220CB(1957830859),
  sa_trans= esp-des esp-md5-hmac , sa_conn_id= 2000
04:29:12: IPSEC(create_sa): sa created,
  (sa) sa_dest= 172.16.30.22, sa_prot= 50,
  sa_spi= 0x179DFOAE(396226734),
  sa_trans= esp-des esp-md5-hmac , sa_conn_id=
2001!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```



```
DLSw Remote NetBIOS Name reachability cache list
NetBIOS Name      status      Loc.      peer
```

```
r5#
```

There is a good chance that you would have DLSW+ connectivity when you see this, but it is more reassuring to see host names. After hooking up hosts and doing some Windows browsing, I went over to R1 and saw this:

```
r1#sh dls w reach
```

```
DLSw Local MAC address reachability cache list
Mac Addr      status      Loc.      port      rif
0040.5407.18b1 FOUND      LOCAL    TBridge-001 --no rif--
00d0.7d67.7c81 FOUND      LOCAL    TBridge-001 --no rif--
```

```
DLSw Remote MAC address reachability cache list
Mac Addr      status      Loc.      peer
00d0.7d67.7ca1 FOUND      REMOTE 10.1.5.1(2065)
00d0.db80.44fd FOUND      REMOTE 10.1.5.1(2065)
```

```
DLSw Local NetBIOS Name reachability cache list
NetBIOS Name  status      Loc.      port      rif
DELL800      FOUND      LOCAL    TBridge-001 --no rif--
```

```
DLSw Remote NetBIOS Name reachability cache list
NetBIOS Name  status      Loc.      peer
DELLNOTE    FOUND      REMOTE 10.1.5.1(2065) max-lf(17800)
```

```
r1#
```

I was able to reach resources on DELLNOTE from DELL800, so I knew DLSW+ was configured correctly.

Rental racks seldom have PCs connected to them for testing of DLSW+, so this is something you might want to test on a very small lab of your own. Two routers connected back-to-back with a PC (running NetBEUI only) hanging off each Ethernet LAN will do.

**15. Configure Low Latency Queueing (Class-Based Weighted Fair Queueing) on the serial link between R6 and R1, giving 100 kbps reserved bandwidth to all IP traffic with IP precedence of 5, and giving that traffic top priority. On the same link, configure multilink PPP fragmentation and interleaving so that packets will not have greater than 10 ms serialization delay. Configure RTP header compression and TCP header compression on the same link.**

This was quite a lot to ask for in one task, but in a real world network you may well want to configure all of these QOS features on one link.

I used a multilink interface, rather than the more traditional virtual-template interface. One could just as easily use a virtual-template interface, which is the more common and

traditional way to configure PPP multilink fragmentation and interleaving. The multilink interface requires a relatively new IOS like a general release 12.2.

You should explore the options available with LLQ now. Here we gave the strict priority queue to IP traffic that has been set to IP precedence 5. Various vendors have their IP phones set the IP precedence of voice payload packets to 5.

However, let us say you cannot trust IP precedence values in your network. You can simply give the strict priority queue to Real Time Protocol (RTP) packets in a range of ports often used by voice traffic, that is, 16384 to 32767:

```
r6(config)#class-map test
r6(config-cmap)#match ?
  access-group      Access group
  any                Any packets
  class-map         Class map
  cos               IEEE 802.1Q/ISL class of service/user priority
  values
  destination-address Destination address
  input-interface   Select an input interface to match
  ip                IP specific values
  mpls              Multi Protocol Label Switching specific values
  not               Negate this match result
  protocol          Protocol
  qos-group         Qos-group
  source-address    Source address
```

```
r6(config-cmap)#match ip ?
  dscp             Match IP DSCP (DiffServ CodePoints)
  precedence       Match IP precedence
  rtp              Match RTP port nos
```

```
r6(config-cmap)#match ip rtp ?
  <2000-65535>    Lower bound of UDP destination port
```

```
r6(config-cmap)#match ip rtp ?
  <2000-65535>    Lower bound of UDP destination port
```

```
r6(config-cmap)#match ip rtp 16384 ?
  <0-16383>      Range of UDP ports
```

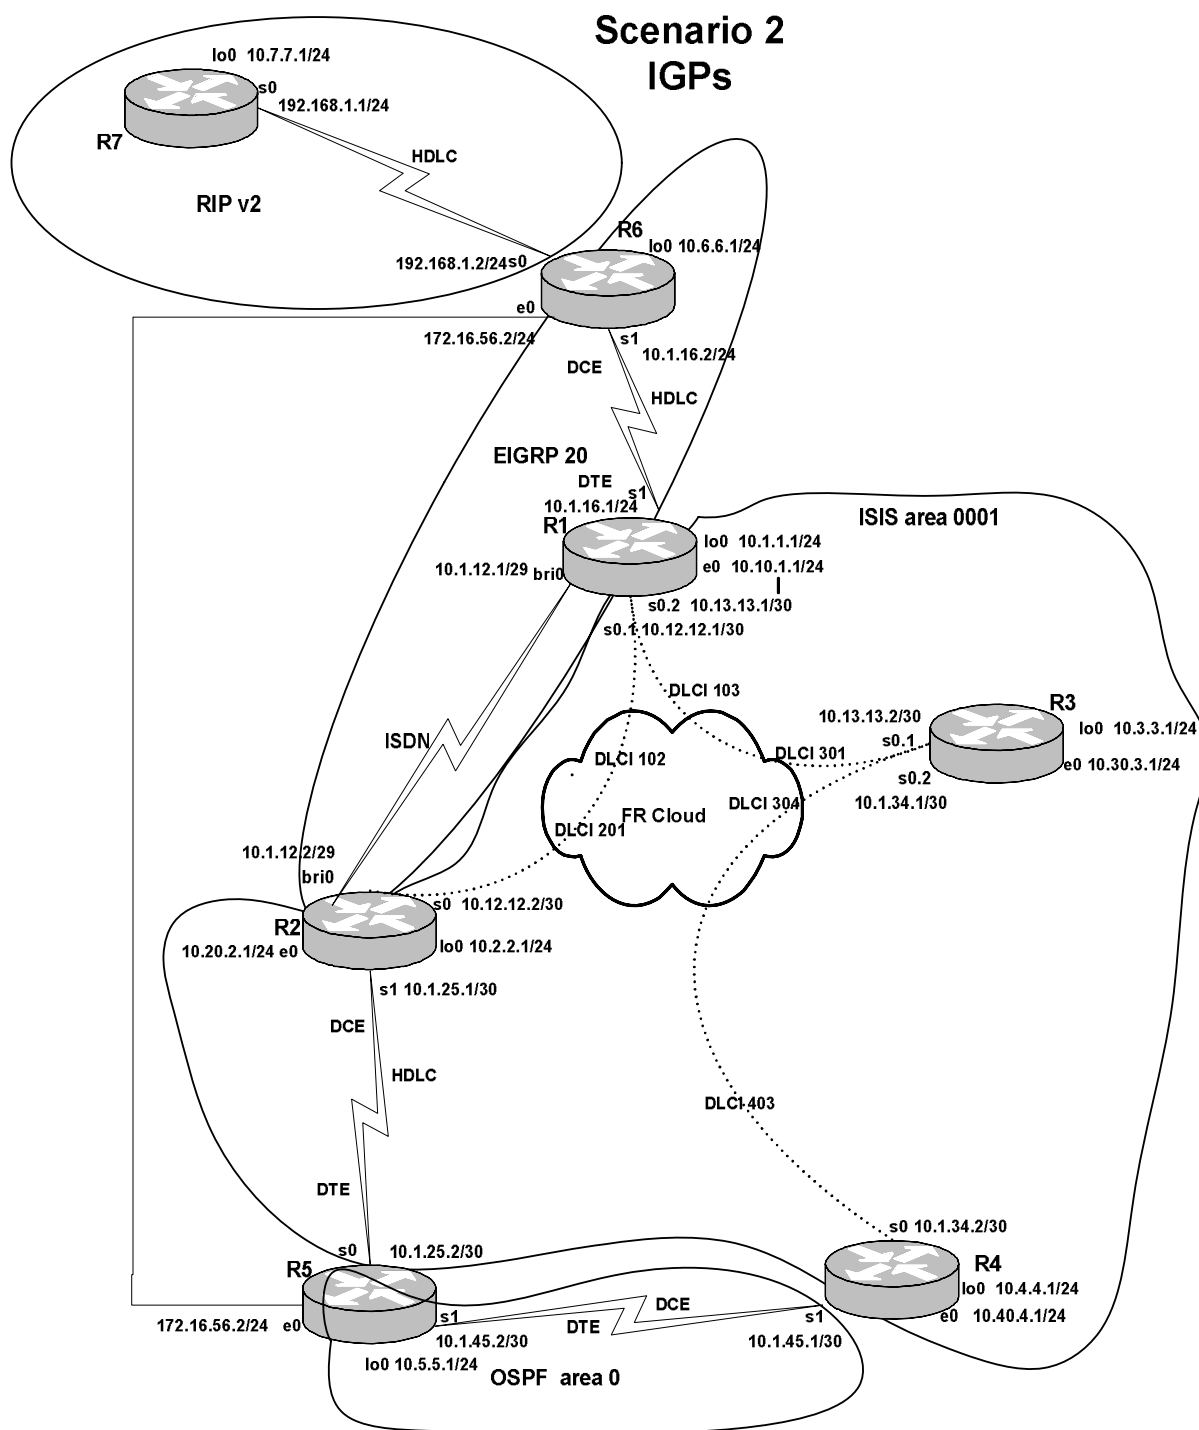
```
r6(config-cmap)#match ip rtp 16384 16383 ?
  <cr>
```

```
r6(config-cmap)#match ip rtp 16384 16383
```

You will have noticed above that you can also match based on DSCP values. One of the best features of Cisco IOS is the interactive help feature. Use it often, even when you may not need it to get the job done, simply to find out whether there are new options.

# Chapter Four

## Scenario Two





Scenario Two focuses heavily on BGP. There are several IGPs running, including a small OSPF domain that breaks up the IS-IS domain in a way that would be unforgivable in a real-world network. The VLANs are as follows:

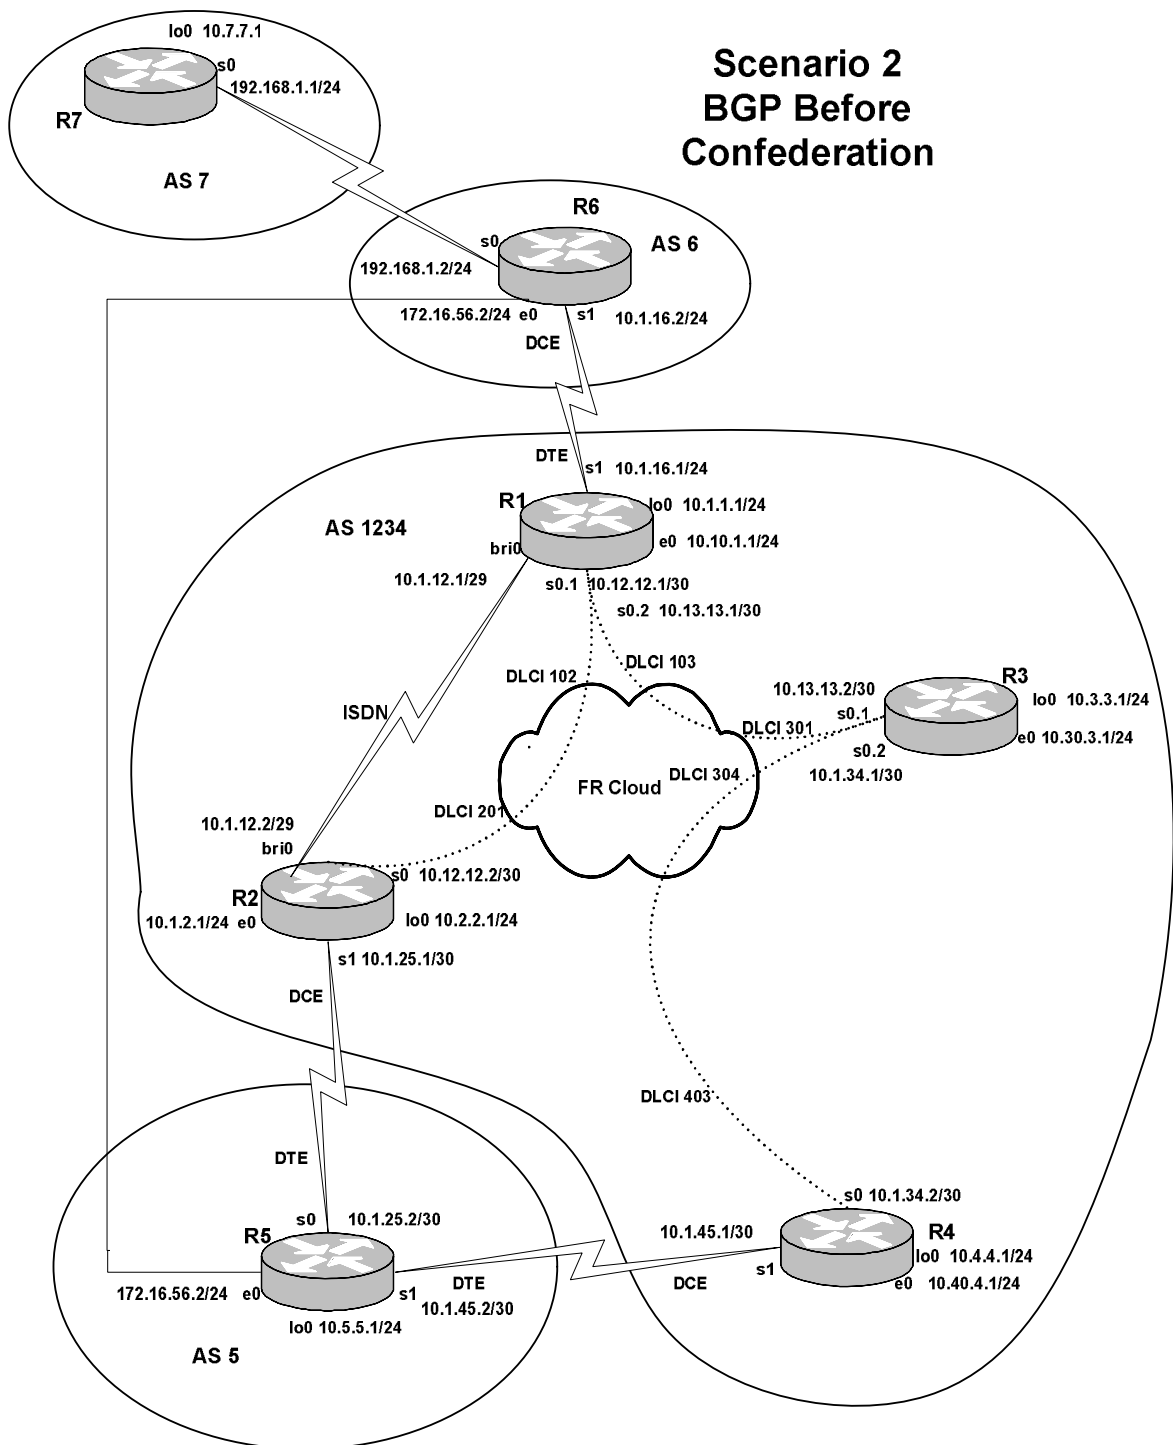
Router	Interface	VLAN #	Switch port # (fastether 0/x)
R1	e0	10	1
R2	e0	20	2
R3	e0	30	3
R4	e0	40	4
R5	e0	56	5
R6	e0	56	6

## Tasks

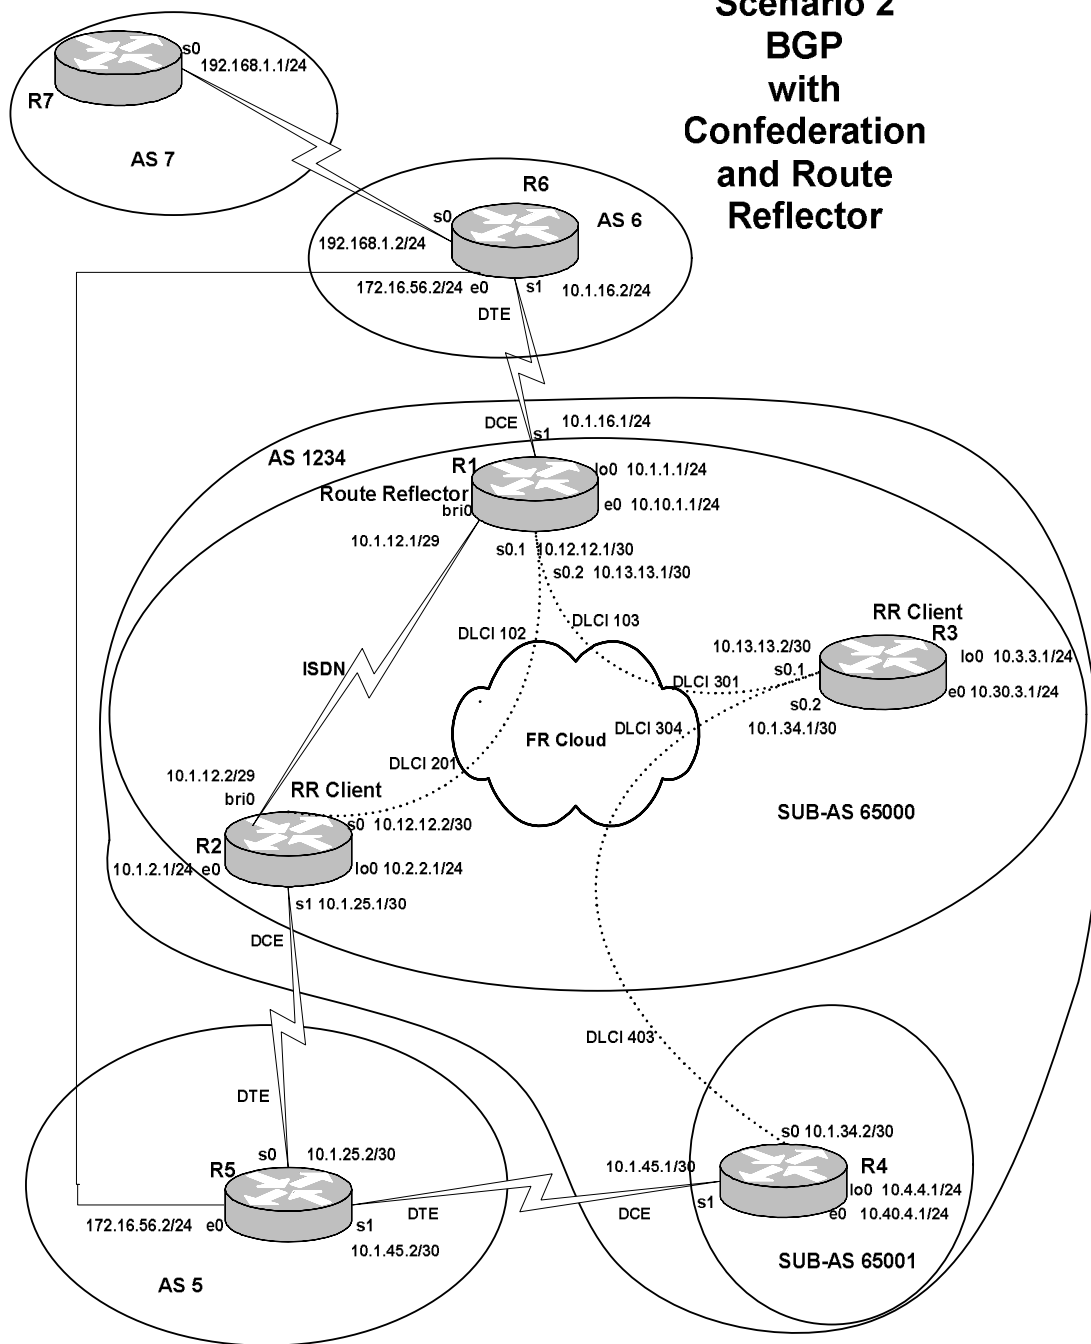
1. Configure Cat interfaces and VLANs as in Scenario One, except that Fa0/5 and Fa0/6 will be in VLAN 56. Enable spantree portfast on all interfaces.
2. Configure all ethernet and loopback addresses as shown in the diagram.
3. Configure frame relay and other serial links as shown in the diagram.
4. Configure the RIP domain as shown in the appropriate diagram. R7 s0 and lo0, and R6 s0 will run RIP.
5. R6 s1 and lo0 will participate in EIGRP 20. R6 e0 will not run any internal routing protocol.
6. R1 bri0 and R2 bri0 will participate in EIGRP 20, and R1 bri0 will be configured to come up when the watched route 10.2.2.0/24 disappears from R1's routing table.
7. The only route redistribution of any kind in this scenario should be redistribution of R1 lo0 and R2 lo0 into EIGRP 20.
8. All interfaces on R1, R2, R3 and R4, except the ISDN interfaces, R1 s1, and R4 s1, will be in ISIS area 0001. Use the router numbers in the NET so that Router 1 in area 0001 will have the NET 49.0001.1111.1111.1111.00. R5 s1 will run OSPF. Since there is only one ISIS area, all routers will be level-1 only.
9. R4 s1 and R5 s1 and lo0 will be in OSPF area 0. R5 e0 will not run any internal routing protocol.

10. For all BGP peering in this scenario, use loopbacks when IGP routes make it possible to peer from loopback to loopback. Disable BGP synchronization on all routers.
11. R7 in BGP AS 7 will peer with R6 in AS 6. Configure peering between R5 and R6 using the ethernet link.
12. R1, R2, R3 and R4 are in AS 1234. There will not be a BGP full-mesh for the routers in AS 1234, so you will need to use a confederation and a route-reflector. R1, R2, and R3 will be in sub-AS 65000, and will use a route reflector. R4 will be in sub-AS 65001. Advertise the loopbacks on all routers in BGP (not just these four routers). Advertise R4's ethernet in BGP throughout AS 1234, but manipulate a BGP attribute so that the route advertisement does not reach R6.
13. Configure BGP so that traffic between AS 5 and AS 6 will normally use the Ethernet link between R5 and R6. In the event that the Ethernet link fails, traffic between AS 5 and AS 6 should transit AS 1234.
14. Configure BGP on R5 so that traffic destined for R1 e0 or R2 e0 will normally take the link to R2, and traffic destined for R3 will normally take the link to R4.
15. On R4 s0, configure TCP header compression and shape traffic to 128kbps with no bursting above that rate on R4 s0.

## Scenario 2 BGP Before Confederation



## Scenario 2 BGP with Confederation and Route Reflector



## Solution Configuration Scripts

### R1

```
r1#sh runn
Building configuration...

Current configuration : 2626 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r1
!
logging rate-limit console 10 except errors
!
username r2 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.0
 ip router isis
!
interface Ethernet0
 ip address 10.10.1.1 255.255.255.0
 ip router isis
!
interface Serial0
 no ip address
 encapsulation frame-relay
!
interface Serial0.1 multipoint
 ip address 10.12.12.1 255.255.255.252
 ip router isis
 frame-relay map clns 102 broadcast
 frame-relay map ip 10.12.12.2 102 broadcast
 no frame-relay inverse-arp
```

```

!
interface Serial0.2 point-to-point
 ip address 10.13.13.1 255.255.255.252
 ip router isis
 frame-relay interface-dlci 103
!
interface Serial1
 ip address 10.1.16.1 255.255.255.0
!
interface BRI0
 ip address 10.1.12.1 255.255.255.248
 encapsulation ppp
 dialer idle-timeout 90
 dialer watch-disable 1
 dialer map ip 10.2.2.0 name r2 broadcast 4082222222
 dialer map ip 10.1.12.2 name r2 broadcast 4082222222
 dialer watch-group 1
 dialer-group 1
 isdn switch-type basic-ni
 isdn spid1 4081111111 4081111111
 isdn spid2 4081111112 4081111111
 isdn T310 40000
 cdapi buffers regular 0
 cdapi buffers raw 0
 cdapi buffers large 0
 ppp authentication chap
!
router eigrp 20
 redistribute connected route-map onlyloop
 network 10.1.12.0 0.0.0.7
 network 10.1.16.0 0.0.0.255
 default-metric 1500 20 255 1 1500
 no auto-summary
 eigrp log-neighbor-changes
!
router isis
 net 49.0001.1111.1111.1111.00
 is-type level-1
!
router bgp 65000
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 1234
 bgp confederation peers 65001
 network 10.1.1.0 mask 255.255.255.0
 neighbor 10.1.16.2 remote-as 6
 neighbor 10.1.16.2 send-community
 neighbor 10.2.2.1 remote-as 65000
 neighbor 10.2.2.1 update-source Loopback0
 neighbor 10.2.2.1 route-reflector-client
 neighbor 10.2.2.1 next-hop-self
 neighbor 10.2.2.1 send-community
 neighbor 10.3.3.1 remote-as 65000
 neighbor 10.3.3.1 update-source Loopback0

```

```

neighbor 10.3.3.1 route-reflector-client
neighbor 10.3.3.1 next-hop-self
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 150 deny eigrp any any
access-list 150 permit ip any any
dialer watch-list 1 ip 10.2.2.0 255.255.255.0
dialer-list 1 protocol ip list 150
route-map onlyloop permit 10
  match interface Loopback0
!
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r1#

```

## R2

```

r2#sh runn
Building configuration...

Current configuration : 2093 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r2
!
logging rate-limit console 10 except errors
!
username r1 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!

```

```

!
!
interface Loopback0
 ip address 10.2.2.1 255.255.255.0
 ip router isis
!
interface Ethernet0
 ip address 10.20.2.1 255.255.255.0
 ip router isis
!
interface Serial0
 ip address 10.12.12.2 255.255.255.252
 ip router isis
 encapsulation frame-relay
 frame-relay map clns 201 broadcast
 frame-relay map ip 10.12.12.1 201 broadcast
 frame-relay map ip 10.12.12.2 201 broadcast
 no frame-relay inverse-arp
!
interface Serial1
 ip address 10.1.25.1 255.255.255.252
 ip router isis
 clockrate 1300000
!
interface BRI0
 ip address 10.1.12.2 255.255.255.248
 encapsulation ppp
 dialer map ip 10.1.12.1 name r1 broadcast 4081111111
 dialer-group 1
 isdn switch-type basic-ni
 isdn spid1 40822222221 4082222222
 isdn spid2 40822222222 4082222222
 isdn T310 40000
 cdapi buffers regular 0
 cdapi buffers raw 0
 cdapi buffers large 0
 ppp authentication chap
!
router eigrp 20
 redistribute connected route-map onlyloop
 network 10.1.12.0 0.0.0.7
 default-metric 1500 20 255 1 1500
 no auto-summary
 no eigrp log-neighbor-changes
!
router isis
 net 49.0001.2222.2222.2222.00
 is-type level-1
!
router bgp 65000
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 1234
 bgp confederation peers 65001

```



```

network 10.2.2.0 mask 255.255.255.0
neighbor 10.1.1.1 remote-as 65000
neighbor 10.1.1.1 update-source Loopback0
neighbor 10.1.25.2 remote-as 5
neighbor 10.1.25.2 send-community
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 150 deny eigrp any any
access-list 150 permit ip any any
dialer-list 1 protocol ip list 150
route-map onlyloop permit 10
  match interface Loopback0
!
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r2#

```

### R3

```

r3#sh runn
Building configuration...

Current configuration : 1512 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r3
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!

```

```

!
!
!
interface Loopback0
 ip address 10.3.3.1 255.255.255.0
 ip router isis
!
interface Ethernet0
 ip address 10.30.3.1 255.255.255.0
 ip router isis
!
interface Serial0
 no ip address
 encapsulation frame-relay
!
interface Serial0.1 point-to-point
 ip address 10.13.13.2 255.255.255.252
 ip router isis
 frame-relay interface-dlci 301
!
interface Serial0.2 multipoint
 ip address 10.1.34.1 255.255.255.252
 ip router isis
 frame-relay map clns 304 broadcast
 frame-relay interface-dlci 304
 frame-relay ip tcp header-compression
!
interface Serial1
 no ip address
 shutdown
!
router isis
 net 49.0001.3333.3333.00
 is-type level-1
!
router bgp 65000
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 1234
 bgp confederation peers 65001
 network 10.3.3.0 mask 255.255.255.0
 neighbor 10.1.1.1 remote-as 65000
 neighbor 10.1.1.1 update-source Loopback0
 neighbor 10.1.1.1 send-community
 neighbor 10.4.4.1 remote-as 65001
 neighbor 10.4.4.1 ebgp-multi-hop 255
 neighbor 10.4.4.1 update-source Loopback0
 no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!

```

```

!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r3#

```

## R4

```

r4#sh runn
Building configuration...

Current configuration : 2243 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r4
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
 ip address 10.4.4.1 255.255.255.0
 ip router isis
!
interface Ethernet0
 ip address 10.40.4.1 255.255.255.0
 ip router isis
!
interface Serial0
 ip address 10.1.34.2 255.255.255.252
 ip router isis
 encapsulation frame-relay
 no fair-queue
 frame-relay class frts34

```

```

frame-relay traffic-shaping
frame-relay map clns 403 broadcast
frame-relay map ip 10.1.34.1 403 broadcast
frame-relay map ip 10.1.34.2 403 broadcast
no frame-relay inverse-arp
frame-relay lmi-type cisco
frame-relay ip tcp header-compression
!
interface Serial1
ip address 10.1.45.1 255.255.255.252
clockrate 1300000
!
interface TokenRing0
no ip address
shutdown
!
router ospf 64
log-adjacency-changes
network 10.1.45.1 0.0.0.0 area 0
!
router isis
net 49.0001.4444.4444.00
is-type level-1
!
router bgp 65001
no synchronization
bgp log-neighbor-changes
bgp confederation identifier 1234
bgp confederation peers 65000
network 10.4.4.0 mask 255.255.255.0
network 10.40.4.0 mask 255.255.255.0
neighbor 10.1.45.2 remote-as 5
neighbor 10.1.45.2 send-community
neighbor 10.1.45.2 route-map setcomm out
neighbor 10.3.3.1 remote-as 65000
neighbor 10.3.3.1 ebgp-multi-hop 255
neighbor 10.3.3.1 update-source Loopback0
neighbor 10.3.3.1 next-hop-self
neighbor 10.3.3.1 send-community
neighbor 10.3.3.1 route-map setcomm out
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
ip access-list standard ether0
permit 10.40.4.0 0.0.0.255
!
map-class frame-relay frts34
frame-relay traffic-rate 128000 128000
no frame-relay adaptive-shaping
route-map setcomm permit 10

```

```

    match ip address ether0
    set community no-export
    !
route-map setcomm permit 20
    !
    !
    !
line con 0
    exec-timeout 0 0
    transport input none
line aux 0
line vty 0 4
    login
    !
end

r4#

```

## R5

```

r5#
Building configuration...

02:44:56: %SYS-5-CONFIG_I: Configured from console by console
Current configuration : 2119 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r5
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
    ip address 10.5.5.1 255.255.255.0
!
interface Ethernet0
    ip address 172.16.56.1 255.255.255.0
!
interface Serial0
    ip address 10.1.25.2 255.255.255.252

```

```

ip router isis
no fair-queue
!
interface Serial1
ip address 10.1.45.2 255.255.255.252
!
router ospf 64
log-adjacency-changes
network 10.1.45.2 0.0.0.0 area 0
network 10.5.5.1 0.0.0.0 area 0
network 10.50.5.1 0.0.0.0 area 0
!
router isis
net 49.0001.5555.5555.00
is-type level-1
!
router bgp 5
no synchronization
bgp log-neighbor-changes
network 10.5.5.0 mask 255.255.255.0
neighbor 10.1.25.1 remote-as 1234
neighbor 10.1.25.1 route-map setweight251 in
neighbor 10.1.45.1 remote-as 1234
neighbor 10.1.45.1 route-map setweight451 in
neighbor 172.16.56.2 remote-as 6
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
ip access-list standard loop1and2
permit 10.1.1.0 0.0.0.255
permit 10.2.2.0 0.0.0.255
ip access-list standard loop3
permit 10.3.3.0 0.0.0.255
route-map setweight451 permit 10
match ip address loop3
set weight 32767
!
route-map setweight451 permit 20
!
route-map setweight251 permit 10
match ip address loop1and2
set weight 32767
!
route-map setweight251 permit 20
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0

```

```
line vty 0 4
 login
 !
end
```

```
r5#
```

## R6

```
r6#sh runn
Building configuration...

Current configuration : 1174 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r6
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
 ip address 10.6.6.1 255.255.255.0
!
interface Ethernet0
 ip address 172.16.56.2 255.255.255.0
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.2 255.255.255.0
 no fair-queue
 clockrate 1300000
!
interface Serial1
 ip address 10.1.16.2 255.255.255.0
 clockrate 1300000
!
router eigrp 20
```

```

network 10.1.16.0 0.0.0.255
network 10.6.6.0 0.0.0.255
no auto-summary
no eigrp log-neighbor-changes
!
router rip
version 2
network 192.168.0.0
no auto-summary
!
router bgp 6
no synchronization
bgp log-neighbor-changes
network 10.6.6.0 mask 255.255.255.0
neighbor 10.1.16.1 remote-as 1234
neighbor 172.16.56.1 remote-as 5
neighbor 192.168.1.1 remote-as 7
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end

r6#

```

## R7

```

r7#sh runn
Building configuration...

Current configuration : 920 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r7
!
logging rate-limit console 10 except errors

```



```

!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
 ip address 10.7.7.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.1 255.255.255.0
 no fair-queue
!
interface Serial1
 no ip address
 shutdown
!
router rip
 version 2
 network 192.168.0.0
 no auto-summary
!
router bgp 7
 no synchronization
 bgp log-neighbor-changes
 network 10.7.7.0 mask 255.255.255.0
 neighbor 192.168.1.2 remote-as 6
 no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
!
line con 0
 exec-timeout 0 0
 transport input none
line aux 0
line vty 0 4
 login
!

```

end

## Cat

```
cat#sh runn
Building configuration...

Current configuration : 2217 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname cat
!
!
ip subnet-zero
no ip domain-lookup
!
!
spanning-tree portfast default
spanning-tree extend system-id
!
!
!
interface FastEthernet0/1
  switchport access vlan 10
  no ip address
!
interface FastEthernet0/2
  switchport access vlan 20
  no ip address
!
interface FastEthernet0/3
  switchport access vlan 30
  no ip address
!
interface FastEthernet0/4
  switchport access vlan 40
  no ip address
!
interface FastEthernet0/5
  switchport access vlan 56
  no ip address
!
interface FastEthernet0/6
  switchport access vlan 56
  no ip address
!
[output truncated]
```

```
!  
interface Vlan1  
  no ip address  
  shutdown  
!  
ip classless  
ip http server  
!  
!  
line con 0  
  exec-timeout 0 0  
line vty 0 4  
  login  
line vty 5 15  
  login  
!  
end  
  
cat#
```

## Explanation

### **1. Configure Cat interfaces and VLANs as in Scenario One, except that Fa0/5 and Fa0/6 will be in VLAN 56. Enable spanning-tree portfast on all interfaces.**

Placing the e0 interfaces of R5 and R6 in the same VLAN creates the ethernet link between the routers shown in the diagram. As in Scenario One, use "spanning-tree portfast default" globally.

### **2. Configure all ethernet and loopback addresses as shown in the diagram.**

No explanation needed.

### **3. Configure frame relay and other serial links as shown in the diagram.**

If you read through the scenario before beginning to configure, you will see that ISIS will be configured on the frame relay links, and you can save time by configuring frame-relay map statements for CLNS at the same time that you configure "frame-relay map" statements for IP. The same principle would apply if you saw a bridging task down the road that would require "frame-relay map bridge" statements.

### **4. Configure the RIP domain as shown in the appropriate diagram. R7 s0 and lo0, and R6 s0 will run RIP.**

Unless you have to run RIP version 1, always run RIP version 2, and disable auto summary.

### **5. R6 s1 and lo0 will participate in EIGRP 20. R6 e0 will not run any internal routing protocol.**

The fact that r6 e0 will not run any internal routing protocol means that any BGP peering later between R6 and R5 will be from one ethernet interface to the other, rather than loopback-to-loopback.

### **6. R1 bri0 and R2 bri0 will participate in EIGRP 20, and R1 bri0 will be configured to come up when the watched route 10.2.2.0/24 disappears from R1 s routing table.**

Dialer watch is a tricky feature. First, the dialer watch-list uses a normal address mask (255.255.255.0) rather than a wildcard mask (0.0.0.255). Second, you need a dialer map statement to the network being watched. Third, even if it is configured correctly, the line could drop and come back up at the end of every idle-timeout

interval. In a production network, there will probably be traffic to keep the link up, but in your practice lab, the link will probably drop.

You should experiment with the "dialer idle-timeout" values and the "dialer watch-disable" values until you reach a state where things work satisfactorily for you. There is also a feature intended to delay the operation of dialer-watch just after a reboot while convergence is going on. You can try typing "dialer watch-list 1 delay route-check initial 120" in global config mode to cause a two minute delay, but I was not happy with the results. Be careful when removing this command because it is easy to end up without a dialer watch-list.

**7. The only route redistribution of any kind in this scenario should be the redistribution of R1 lo0 and R2 lo0 into EIGRP 20.**

The redistribution of these two connected networks into EIGRP should permit basic connectivity between R1 and R2 and BGP peering between them after a frame relay outage. However, that is about it. If the whole frame relay cloud fails, some BGP peerings will be impossible. You might want to explore the limits of reachability throughout the network after a failure of the frame relay link between R1 and R2, including the usefulness of the ethernet link between R5 and R6.

The end result of this scenario will not be any-to-any connectivity, even when all links are up. As you work through CCIE practice scenarios, remember to configure only what is required. Some of us are tempted to glance at a practice scenario and then merrily go about building a network in which we can ping every interface from every other interface. As you work through CCIE practice scenarios, think of the game "Simon says," because an important part of the game is following instructions carefully.

**8. All interfaces on R1, R2, R3 and R4, except the ISDN interfaces, R1 s1, and R4 s1, will be in ISIS area 0001. Use the router numbers in the NET so that Router 1 in area 0001 will have the NET 49.0001.1111.1111.1111.00. R5 s1 will run OSPF. Since there is only one ISIS area, all routers will be level-1 only.**

There is a potential for IS-IS interface-type mismatch between the physical frame-relay interfaces and the subinterfaces. The simple solution is to use point-to-multipoint subinterfaces, which will form adjacencies with frame-relay physical interfaces.

Another way to cure an IS-IS interface-type mismatch is to use tunnels, so IS-IS does not need to be concerned about the underlying physical interfaces or subinterfaces.

Making the loopback and ethernet interfaces passive prevents them from sending out useless hellos all the time on those interfaces, and will make debugging ISIS adjacency formation a lot easier, because "debug isis adj" will produce much less output to wade through.

**9. R4 s1 and R5 s1 and lo0 will be in OSPF area 0. R5 e0 will not run any internal routing protocol.**

This OSPF area would be an abominable design flaw in a production network. While it does not break the ISIS network, it certainly breaks it up and deprives us of some redundancy. Watch out for the little detail that R5 e0 does not run any internal routing protocol. Be conscious of exactly which interfaces are participating in which routing protocols.

**10. For all BGP peering in this scenario, use loopbacks when IGP routes make it possible to peer from loopback to loopback. Disable BGP synchronization on all routers.**

IS-IS routes make peering loopback-to-loopback possible between R1 and R2, R1 and R3, and R3 and R4. All other peering should be done from one end of a link to the other end of a link.

Remember that when you peer between loopbacks you need to specify your loopback as the update-source, and if it is an EBGP peering, you need to enable EBGP-multihop.

**11. R7 in BGP AS 7 will peer with R6 in AS 6. Configure peering between R5 and R6 using the ethernet link.**

Obviously, if you are peering from one end of a link to another, you can do so with or without running an internal routing protocol on that link.

**12. R1, R2, R3 and R4 are in AS 1234. There will not be a BGP full-mesh for the routers in AS 1234, so you will need to use a confederation and a route-reflector. R1, R2, and R3 will be in sub-AS 65000, and will use a route reflector. R4 will be in sub-AS 65001. Advertise the loopbacks on all routers in BGP (not just these four routers). Advertise R4 s ethernet in BGP throughout AS 1234, but manipulate a BGP attribute so that the route advertisement does not reach R6.**

You will make R1 the route-reflector because connects to both R2 and R3. Go ahead and use network statements to advertise the loopbacks. An alternate way to get loopbacks into BGP would be to "redistribute connected route-map looponly" and then create a route-map named "looponly" to permit only the loopback into BGP. You would probably only use this alternate method if the simpler "network" statement was prohibited.

Remember the word "mask" in BGP network statements, and remember to advertise the network with the precise mask. With the BGP network statement, you are advertising a network, not specifying an interface or interfaces to participate in a routing protocol.

To keep the BGP route for R4's ethernet from reaching R6, you set its attribute to "no-export" and make sure you enable sending of community attributes to all BGP peers. No-export allows a route to go to the next AS, but not beyond. The sub-AS 65000 is treated as a separate AS from sub-AS 65001, so the route will not reach R6. While it may not be immediately apparent, you need to propagate the community attribute to R5 in AS 5 to keep R6 from learning about R4's ethernet through its BGP peering with R5.

You should also note that "send-community" is required to propagate the community attribute, even to internal BGP peers.

As you configure BGP, you will find a few next-hop reachability problems that need to be overcome by using "next-hop-self." On R1 the next-hop address of the BGP route to R6's loopback is the other end of the serial link that connects R1 and R6, so R1 has no problem with reaching the next-hop. However, R1's serial link to R6 is in the EIGRP routing domain, and is never redistributed into the IS-IS routing domain, so R2 and R3 will have no idea how to reach that next-hop address. Therefore, R1 needs to be configured to tell R2 and R3 to use its own address as the next-hop address.

**13. Configure BGP so that traffic between AS 5 and AS 6 will normally use the Ethernet link between R5 and R6. In the event that the Ethernet link fails, traffic between AS 5 lo0 and AS 6 lo0 should transit AS 1234.**

You do not need a single line of script to fulfill this task. This task describes the default behavior. This is how R6's BGP table should look:

```
r6#sh ip bgp
BGP table version is 13, local router ID is 10.6.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i
- internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*  10.1.1.0/24      172.16.56.1              0         0 5 1234
i
*>
*> 10.2.2.0/24      10.1.16.1                0         0 1234 i
*  10.2.2.0/24      172.16.56.1              0         0 5 1234
i
*> 10.3.3.0/24      10.1.16.1                0         0 1234 i
*  10.3.3.0/24      172.16.56.1              0         0 5 1234
i
*> 10.4.4.0/24      10.1.16.1                0         0 1234 i
*  10.4.4.0/24      172.16.56.1              0         0 5 1234
i
*  10.5.5.0/24     10.1.16.1                0        0 1234 5
i
*>
*> 10.6.6.0/24      172.16.56.1             0         0 5 i
*  10.6.6.0/24      0.0.0.0                  0         32768 i
```

```
*> 10. 7. 7. 0/24      192. 168. 1. 1      0      0 7 i
r6#
```

By default, R6 prefers the routes with shorter AS-paths, so it will prefer the ethernet link to R5. If that link fails, it will go through AS 1234 to reach AS 5. The lesson here is to avoid rushing to configure something elaborate where little or no configuration is necessary.

#### 14. Configure BGP on R5 so that traffic destined for R1 e0 or R2 e0 will normally take the link to R2, and traffic destined for R3 will normally take the link to R4.

Here you will manipulate an attribute of certain routes as they come in from AS 1234. I manipulated weight because I only needed to influence the routing decision on one Cisco router. If I had wanted to have this preference shared among routers in an AS, I would have used local preference.

If you are using a route-map to manipulate attributes, remember to include the final "catch-all" permit statement. If you are using a route-map to permit certain routes and deny all others, then you will not need a "catch-all" permit statement at the end.

We could have omitted the route-map statements matching the non-preferred routes, as they would come in under the catch-all statements with a default weight of zero, anyway. I left them in only for logical clarity.

Here is how R5's BGP table will look:

```
r5#sh ip bgp
BGP table version is 11, local router ID is 10.5.5.1
Status codes: s suppressed, d damped, h history, * valid, > best, i
- internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
* 10.1.1.0/24       172.16.56.2              0 6 1234
i
*>                  10.1.25.1             32767 1234 i
*                   10.1.45.1              0 1234 i
* 10.2.2.0/24       172.16.56.2              0 6 1234
i
*>                  10.1.25.1              0             32767 1234 i
*                   10.1.45.1              0 1234 i
* 10.3.3.0/24       172.16.56.2              0 6 1234
i
*                   10.1.25.1              0 1234 i
*>                  10.1.45.1             32767 1234 i
* 10.4.4.0/24       172.16.56.2              0 6 1234
i
*                   10.1.25.1              0 1234 i
*>                  10.1.45.1              0             0 1234 i
```



```

*> 10.5.5.0/24      0.0.0.0          0          32768 i
* 10.6.6.0/24      10.1.45.1        0          0 1234 6
i
*                10.1.25.1        0          0 1234 6
i
*>                172.16.56.2      0          0 6 i
* 10.7.7.0/24      10.1.45.1        0          0 1234 6
7 i
*                10.1.25.1        0          0 1234 6
7 i
  Network          Next Hop          Metric LocPrf Weight Path
*>                172.16.56.2      0          0 6 7 i
*> 10.40.4.0/24     10.1.45.1        0          0 1234 i
r5#

```

As you can see, the router has chosen the best route (designated by the ">" mark) to the pertinent networks based on the weight.

You can also have problems due to typos in route-map names. The route-map names need to match exactly, and are case-sensitive. It is common to make route-map names all upper case, probably for readability as one reviews the logic of configuration scripts, but for the purposes of CCIE lab preparation I prefer using only lower case in configuration scripts to facilitate typing and to avoid mismatches due to case sensitivity.

#### 15. On R4 s0, configure TCP header compression and shape traffic to 128kbps with no bursting above that rate.

The trick here is that if you configure TCP header compression on one side of a link, you need to configure it at the other side as well. If you do not do this, you will find that your BGP peering session between R3 and R4 will fail, because it uses TCP. You will still be able to ping between R3 lo0 and R4 lo0, so if you simply configure TCP header compression but do not stay long enough on R3 or R4 to see the error message when the BGP peering fails, you could be in trouble and not know it. You can also test whether there is a break in TCP connectivity by telneting from R4 to R3's loopback address:

```

r4#telnet 10.3.3.1
Trying 10.3.3.1 ... Open

```

Password required, but none set

```

[Connection to 10.3.3.1 closed by foreign host]
r4#

```

The "Password required, but none set" indicates successful TCP connectivity.

This task should serve as a warning that one configuration task can break your network and could conceivably cost you more points than failing to configure it at all would have cost you.


This task also demonstrates the limits of testing using ICMP (ping). This is often the best testing tool available, but do not sit back and assume that you have configured everything correctly just because you can ping every interface from every other interface.

Unlike header compression, traffic shaping can be configured on one side of a link.

## Chapter Five

### Scenario Three





The Catalyst 3550 is the center of this scenario that focuses heavily on route redistribution. Since the Catalyst 3550 is the star, perhaps you will forgive my creating a scenario so unrealistic as to have only two BGP autonomous systems and only one connection between them. You can set up each interface as a router interface, but I chose to use VLAN interfaces as the routing interfaces so that no recabling needs to be done between scenarios.

The Catalyst 3550 will run three routing protocols in this scenario: RIP, EIGRP, and OSPF. Since we are using the Ethernet interfaces for connectivity between routers, we will use loopback addresses to represent LANs.

The VLANs are as follows:

Router	Interface	VLAN #	Switch port # (fastether 0/x)
R1	e0	10	1
R3	e0	30	3
R5	e0	50	5
R6	e0	60	6

## Tasks

1. Configure the frame relay links as shown in the diagram. Do not use any subinterfaces.
2. Configure the other serial interfaces with the default encapsulation.
3. Configure the appropriate 10/100 interfaces, VLANs, and VLAN interfaces, on the Cat 3550. Configure the Ethernet interfaces on R1, R3, R5, and R6.
4. Configure OSPF area 0 as shown in the diagram.
5. Configure OSPF area 34 as a totally stubby area.
6. Configure EIGRP AS 20 as shown in the diagram. For the ISDN link, configure dialer watch so that R2 dials R1 when the route to R1 lo0 disappears from its routing table. Use dialer profiles.
7. Configure RIP as shown in the diagram. R6 s0 will not run RIP, but its network will be redistributed into RIP.
8. On Cat, configure redistribution among the three internal gateway protocols.

9. Configure BGP peering between R6 and R7.
10. Advertise R7's loopback addresses in BGP. Use a prefix-list on R7 to make sure that only the 172.16.30.0/24 route is advertised to AS 6.
11. Redistribute BGP routes into RIP on R6.
12. R6 will be a firewall router. Configure it to permit inside users to initiate outgoing TCP and UDP sessions, and to let in return traffic belonging to those sessions, as well as any necessary routing protocol traffic. Permit ping testing from networks behind the firewall router to networks outside the firewall router.
13. Configure NAT on R6 so that all outgoing packets that pass through R6 e0 will have the IP address of R6 s0 as their source IP address.
14. Insure that R4 can telnet to R7's 172.16.30.1 loopback address.
15. Configure IP multicast using a shared tree. R2 lo0 will serve as the source of multicast traffic for testing purposes, and will also serve as the rendezvous point. Have the lo0 interfaces of R1, R3, R4, and R5 receive multicast traffic addressed to 235.0.0.1. Test your configuration by pinging 235.0.0.1 from R2's lo0. The ISDN links will not participate in multicast routing.

## Solution Configuration Scripts

### R1

```

r1#s
Building configuration...

Current configuration : 2160 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r1
!
logging rate-limit console 10 except errors
!
username r2 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
ip multicast-routing
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.0
 ip igmp join-group 235.0.0.1
!
interface Ethernet0
 ip address 10.10.1.2 255.255.255.252
 ip pim sparse-mode
!
interface Serial0
 ip address 10.1.12.1 255.255.255.252
 ip pim sparse-mode
 encapsulation frame-relay
 frame-relay map ip 10.1.12.1 102 broadcast
 frame-relay map ip 10.1.12.2 102 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 no ip address
 shutdown
!

```

```

interface BRI0
  no ip address
  encapsulation ppp
  dialer pool-member 1
  isdn switch-type basic-ni
  isdn spid1 4081111111 4081111111
  isdn spid2 4081111112 4081111111
  cdapi buffers regular 0
  cdapi buffers raw 0
  cdapi buffers large 0
  no cdp enable
  ppp authentication chap
!
interface Dialer1
  ip address 10.1.21.1 255.255.255.0
  encapsulation ppp
  dialer pool 1
  dialer remote-name r2
  dialer idle-timeout 30
  dialer-group 1
!
router eigrp 20
  network 10.1.1.0 0.0.0.255
  network 10.1.12.0 0.0.0.3
  network 10.1.21.0 0.0.0.255
  network 10.10.1.0 0.0.0.3
  no auto-summary
  eigrp log-neighbor-changes
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 10.1.2.1
!
dialer-list 1 protocol ip permit
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r1#

```

## R2

```

r2#sh runn
Building configuration...

```



```

Current configuration : 2305 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r2
!
logging rate-limit console 10 except errors
!
username r1 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
ip multicast-routing
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Loopback0
 ip address 10.1.2.1 255.255.255.0
 ip pim sparse-mode
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 10.1.12.2 255.255.255.252
 ip pim sparse-mode
 encapsulation frame-relay
 frame-relay map ip 10.1.12.1 201 broadcast
 frame-relay map ip 10.1.12.2 201 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 ip address 10.1.25.1 255.255.255.252
 ip pim sparse-mode
 clockrate 1300000
!
interface BRI0
 no ip address
 encapsulation ppp
 dialer pool-member 1
 isdn switch-type basic-ni
 isdn spid1 40822222221 4082222222
 isdn spid2 40822222222 4082222222
 cdapi buffers regular 0

```

```

cdapi buffers raw 0
cdapi buffers large 0
no cdp enable
ppp authentication chap
!
interface Dialer1
ip address 10.1.21.2 255.255.255.0
encapsulation ppp
dialer pool 1
dialer remote-name r1
dialer string 4081111111
dialer watch-group 1
dialer-group 1
no cdp enable
!
router eigrp 20
network 10.1.2.0 0.0.0.255
network 10.1.12.0 0.0.0.3
network 10.1.21.0 0.0.0.255
network 10.1.25.0 0.0.0.3
no auto-summary
no eigrp log-neighbor-changes
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 10.1.2.1
!
access-list 150 deny eigrp any any
access-list 150 permit ip any any
dialer watch-list 1 ip 10.1.1.0 255.255.255.0
dialer-list 1 protocol ip list 150
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end

r2#

```

### R3

```

r3#sh runn
Building configuration...

Current configuration : 1700 bytes
!

```

```

version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r3
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
ip multicast-routing
no ip dhcp-client network-discovery
!
!
!
interface Loopback0
 ip address 10.1.3.1 255.255.255.0
 ip pim sparse-mode
 ip ospf network point-to-point
 ip igmp join-group 235.0.0.1
!
interface Ethernet0
 ip address 10.10.3.2 255.255.255.252
 ip pim sparse-mode
!
interface Serial0
 ip address 10.1.34.1 255.255.255.252
 ip pim sparse-mode
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 10.1.34.1 304 broadcast
 frame-relay map ip 10.1.34.2 304 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 no ip address
 shutdown
!
router ospf 64
 log-adjacency-changes
 area 34 stub no-summary
 network 10.1.3.1 0.0.0.0 area 34
 network 10.1.34.1 0.0.0.0 area 34
 network 10.10.3.2 0.0.0.0 area 0
!
ip kerberos source-interface any
ip classless
ip http server

```

```

ip pim rp-address 10.1.2.1
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r3#

```

## R4

```

r4#sh runn
Building configuration...

Current configuration : 1786 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r4
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
ip multicast-routing
no ip dhcp-client network-discovery
!
!
!
interface Loopback0
 ip address 10.1.4.1 255.255.255.0
 ip pim sparse-mode
 ip ospf network point-to-point
 ip igmp join-group 235.0.0.1
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 10.1.34.2 255.255.255.252

```

```

ip pim sparse-mode
encapsulation frame-relay
ip ospf network point-to-point
frame-relay map ip 10.1.34.1 403 broadcast
frame-relay map ip 10.1.34.2 403 broadcast
no frame-relay inverse-arp
frame-relay lmi-type cisco
!
interface Serial1
ip address 10.1.45.1 255.255.255.252
ip pim sparse-mode
clockrate 1300000
!
interface TokenRing0
no ip address
shutdown
!
router ospf 64
log-adjacency-changes
area 34 stub no-summary
network 10.1.4.1 0.0.0.0 area 34
network 10.1.34.2 0.0.0.0 area 34
network 10.1.45.1 0.0.0.0 area 34
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 10.1.2.1
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end

r4#

```

## R5

```

r5#sh runn
Building configuration...

Current configuration : 1795 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime

```

```

no service password-encryption
!
hostname r5
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
ip multicast-routing
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
 ip address 10.1.5.1 255.255.255.0
 ip pim sparse-mode
 ip ospf network point-to-point
 ip igmp join-group 235.0.0.1
!
interface Ethernet0
 ip address 10.10.5.2 255.255.255.252
 ip pim sparse-mode
!
interface Serial0
 ip address 10.1.25.2 255.255.255.252
 ip pim sparse-mode
 no fair-queue
!
interface Serial1
 ip address 10.1.45.2 255.255.255.252
 ip pim sparse-mode
!
router eigrp 20
 redistribute connected metric 1500 20 255 1 1500 route-map looponly
 network 10.1.25.0 0.0.0.3
 no auto-summary
 no eigrp log-neighbor-changes
!
router ospf 64
 log-adjacency-changes
 area 34 stub no-summary
 network 10.1.5.1 0.0.0.0 area 0
 network 10.1.25.2 0.0.0.0 area 0
 network 10.1.45.2 0.0.0.0 area 34
 network 10.10.5.2 0.0.0.0 area 0
!
 ip kerberos source-interface any
 ip classless
 ip http server
 ip pim rp-address 10.1.2.1
!

```

```

route-map looponly permit 10
  match interface Loopback0
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r5#

```

## R6

```

r6#sh runn
Building configuration...

Current configuration : 2087 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r6
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
interface Loopback0
 ip address 10.1.6.1 255.255.255.0
!
interface Ethernet0
 ip address 10.10.6.2 255.255.255.252
 ip nat inside
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0

```

```

ip address 192.168.1.2 255.255.255.0
ip access-group firewallin in
ip access-group firewallout out
ip nat outside
no fair-queue
clockrate 1300000
!
interface Serial1
no ip address
shutdown
!
router rip
version 2
redistribute connected route-map s0only
redistribute bgp 6 metric 3
network 10.0.0.0
no auto-summary
!
router bgp 6
no synchronization
bgp log-neighbor-changes
neighbor 192.168.1.1 remote-as 7
no auto-summary
!
ip kerberos source-interface any
ip nat inside source list incogniti interface Serial0 overload
ip classless
ip http server
!
!
ip access-list standard incogniti
deny 192.168.1.0 0.0.0.255
permit any
!
ip access-list extended firewallin
permit tcp any any eq bgp
permit tcp any eq bgp any
evaluate outboundssessions
ip access-list extended firewallout
permit tcp any any reflect outboundssessions
permit udp any any reflect outboundssessions
permit icmp any any reflect outboundssessions
route-map s0only permit 10
match interface Serial0
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end

```



r6#

## R7

```
r7#sh runn
Building configuration...

Current configuration : 1705 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r7
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
interface Loopback0
 ip address 172.16.29.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.30.1 255.255.255.0
!
interface Loopback2
 ip address 172.16.31.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.1 255.255.255.0
 no fair-queue
!
interface Serial1
 no ip address
 shutdown
```

```
!  
router bgp 7  
  no synchronization  
  bgp log-neighbor-changes  
  network 172.16.29.0 mask 255.255.255.0  
  network 172.16.30.0 mask 255.255.255.0  
  network 172.16.31.0 mask 255.255.255.0  
  neighbor 192.168.1.2 remote-as 6  
  neighbor 192.168.1.2 prefix-list evenonly out  
  no auto-summary  
!  
ip kerberos source-interface any  
ip classless  
ip http server  
!  
!  
ip prefix-list evenonly seq 5 permit 172.16.30.0/24  
!  
line con 0  
  exec-timeout 0 0  
  transport input none  
line aux 0  
line vty 0 4  
  password cisco  
  login  
!  
end  
  
r7#
```

## Cat

```
cat#sh runn  
Building configuration...  
  
Current configuration : 3054 bytes  
!  
version 12.1  
no service pad  
service timestamps debug uptime  
service timestamps log uptime  
no service password-encryption  
!  
hostname cat  
!  
!  
ip subnet-zero  
ip routing  
no ip domain-lookup  
!  
ip multicast-routing  
!
```

```
spanning-tree portfast default
spanning-tree extend system-id
!
!
!
interface FastEthernet0/1
  switchport access vlan 10
  no ip address
!
interface FastEthernet0/2
  no ip address
!
interface FastEthernet0/3
  switchport access vlan 30
  no ip address
!
interface FastEthernet0/4
  no ip address
!
interface FastEthernet0/5
  switchport access vlan 50
  no ip address
!
interface FastEthernet0/6
  switchport access vlan 60
  no ip address
!
[output truncated for brevity]
!
interface Vlan1
  no ip address
  shutdown
!
interface Vlan10
  ip address 10.10.1.1 255.255.255.252
  ip pim sparse-mode
!
interface Vlan30
  ip address 10.10.3.1 255.255.255.252
  ip pim sparse-mode
!
interface Vlan50
  ip address 10.10.5.1 255.255.255.252
  ip pim sparse-mode
!
interface Vlan60
  ip address 10.10.6.1 255.255.255.252
  ip pim sparse-mode
!
router eigrp 20
  redistribute ospf 64 metric 1500 20 255 1 1500
  redistribute rip metric 1500 20 255 1 1500
  network 10.10.1.0 0.0.0.3
```

```
no auto-summary
no eigrp log-neighbor-changes
!
router ospf 64
log-adjacency-changes
redistribute eigrp 20 metric-type 1 subnets
redistribute rip metric-type 1 subnets
network 10.10.3.1 0.0.0.0 area 0
network 10.10.5.1 0.0.0.0 area 0
!
router rip
version 2
redistribute eigrp 20 metric 2
redistribute ospf 64 metric 2
passive-interface default
no passive-interface Vlan60
network 10.0.0.0
no auto-summary
!
ip classless
ip http server
ip pim rp-address 10.1.2.1
!
line con 0
exec-timeout 0 0
line vty 0 4
login
line vty 5 15
login
!
end

cat#
```

## Explanation

### 1. Configure the frame relay links as shown in the diagram. Do not use any subinterfaces.

Obviously, you do not have to configure the tasks in order. You might well go from router to router, configuring all of the basic connectivity tasks. Get a lot of practice, and you will find a routine that works well for you.

The fact that we cannot use any subinterfaces means that all the links should have inverse arp disabled and have frame relay map statements. You can use inverse arp, but it is a good habit to avoid relying on it. You also need to consider such things DR/BDR elections, split horizon, and network type, depending on the routing protocol.

Here there are no frame-relay hub-and-spoke networks, so things are pretty simple. You should enable split horizon on the frame relay link between R1 and R2, to be on the safe side.

### 2. Configure the other serial interfaces with the default encapsulation.

Default encapsulation on Cisco serial interfaces is HDLC.

### 3. Configure the appropriate 10/100 interfaces, VLANs, and VLAN interfaces, on the Cat 3550. Configure the Ethernet interfaces on R1, R3, R5, and R6.

Here we put each pertinent router's ethernet interface in the same VLAN as the appropriate Cat fastethernet interface.

### 4. Configure OSPF area 0 as shown in the diagram.

No explanation needed.

### 5. Configure OSPF area 34 as a totally stubby area.

The totally stubby area is worth studying closely. Take a look at the routing table and OSPF database on R4, before we make a cost adjustment to make the multicast task work correctly:

```
r4#sh ip rou
```

```
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 10.1.34.1 to network 0.0.0.0

```

10.0.0.0/8 is variably subnetted, 4 subnets, 3 masks
O    10.1.3.1/32 [110/65] via 10.1.34.1, 01:00:38, Serial0
C    10.1.4.0/24 is directly connected, Loopback0
C    10.1.45.0/30 is directly connected, Serial1
C    10.1.34.0/30 is directly connected, Serial0
O*IA 0.0.0.0/0 [110/65] via 10.1.34.1, 01:00:38, Serial0
      [110/65] via 10.1.45.2, 01:00:38, Serial1
r4#

```

r4#**sh ip ospf database**

OSPF Router with ID (10.1.4.1) (Process ID 64)

Router Link States (Area 34)

Link ID count	ADV Router	Age	Seq#	Checksum	Link
10.1.3.1	10.1.3.1	1899	0x80000003	0xD824	3
10.1.4.1	10.1.4.1	1856	0x80000004	0x3BA0	5
10.1.5.1	10.1.5.1	43	0x80000004	0x9D63	2

Summary Net Link States (Area 34)

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	10.1.3.1	1899	0x80000002	0x32FB
0.0.0.0	10.1.5.1	43	0x80000003	0x2209

r4#

All you get on R4 is connected routes, a route to R3's loopback address, because it is also within the totally stubby area, and two equal-cost default routes leading you to area 0.

In the OSPF database, you see that only the default routes show up under "Summary Net Link States (Area 34)."

Examine the routing tables and OSPF databases on the OSPF routers, and get a feel for the differences. Think about what the stubbiness of area 34 would mean for redundancy in the event of a link failure between R3 and Cat. Because no inter-area routes are carried in a totally stubby area, Cat will not know how to reach R3 through area 34 if the ethernet between R3 and Cat fails.

**6. Configure EIGRP AS 20 as shown in the diagram. For the ISDN link, configure dialer watch so that R2 dials R1 when the route to R1 lo0 disappears from its routing table. Use dialer profiles.**

The first thing you should do when configuring EIGRP is to disable auto-summary. It is so easy to forget to do this when you are concentrating on getting your network masks right or considering subtle issues raised by a scenario.

To make the dialer watch feature work with EIGRP, it is important to use an access list with your dialer list to prevent EIGRP hellos from bringing up the line. In addition, I configured a dialer string on only one side of the link to make sure that only R2 dials R1. You can use "dialer map" on a dialer profile, but "dialer map" and "dialer remote-name" are mutually exclusive.

One of the "gotchas" of using dialer watch and legacy DDR is that you need a dialer map statement for each watched route. This configuration using dialer profiles did not require a "dialer map" statement for any routes.

Note that "ppp authentication chap" is on the physical interface, rather than on the dialer interface. This is a command that is applied before authentication is complete, so it needs to be configured on the physical interface.

Here is the debug output on R2 when we type "**debug dialer**" and then shut down the frame relay interface on R2:

```
r2(config)#int s0
r2(config-if)#shu
r2(config-if)#
04:25:19: DDR: Dialer Watch: watch-group = 1
04:25:19: DDR: network 10.1.1.0/255.255.255.0 DOWN,
04:25:19: DDR: primary DOWN
04:25:19: DDR: Dialer Watch: Dial Reason: Primary of group 1 DOWN
04:25:19: DDR: Dialer Watch: watch-group = 1,
04:25:19: BRO DDR: rotor dialout [priority]
04:25:19: DDR: dialing secondary by dialer string 4081111111 on Di 1
04:25:19: BRO DDR: Attempting to dial 4081111111
04:25:19: DDR: Dialer Watch: watch-group = 1
04:25:19: DDR: network 10.1.1.0/255.255.255.0 DOWN,
04:25:19: DDR: primary DOWN
04:25:19: DDR: Dialer Watch: Dial Reason: Primary of group 1 DOWN
04:25:19: DDR: Dialer Watch: watch-group = 1,
04:25:19: %LINK-3-UPDOWN: Interface BRI0:1, changed state to up
04:25:19: BRO:1 DDR: Dialer Watch: resetting call in progress
04:25:19: BRO:1: interface must be fifo queue, force fifoexi
04:25:19: %DIALER-6-BIND: Interface BRO:1 bound to profile Di 1
r2(config)#
04:25:21: %LINK-5-CHANGED: Interface Serial0, changed state to
admini strati vely
downexi
r2#
04:25:22: BRO:1 DDR: di aler protocol up
```

```

04:25:22: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0,
changed state
to down
04:25:22: %LINEPROTO-5-UPDOWN: Line protocol on Interface BRI0:1,
changed state
to up
04:25:23: %SYS-5-CONFIG_I: Configured from console by console
r2#sh ip rou
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
* - candi date default, U - per-user static route, o - ODR
P - periodic downloaded static route

```

Gateway of last resort is not set

```

172.16.0.0/24 is subnetted, 1 subnets
D EX 172.16.30.0 [170/46256896] via 10.1.21.1, 00:00:06, Dialer1
10.0.0.0/8 is variably subnetted, 16 subnets, 3 masks
D 10.10.1.0/30 [90/46251776] via 10.1.21.1, 00:00:06, Dialer1
D EX 10.10.3.0/30 [170/46256896] via 10.1.21.1, 00:00:06, Dialer1
D EX 10.10.5.0/30 [170/46256896] via 10.1.21.1, 00:00:06, Dialer1
D EX 10.10.6.0/30 [170/46256896] via 10.1.21.1, 00:00:06, Dialer1
D 10.1.12.0/30 [90/46738176] via 10.1.21.1, 00:00:06, Dialer1
D EX 10.1.3.0/24 [170/46256896] via 10.1.21.1, 00:00:07, Dialer1
C 10.1.2.0/24 is directly connected, Loopback0
D 10.1.1.0/24 [90/46354176] via 10.1.21.1, 00:00:07, Dialer1
D EX 10.1.6.0/24 [170/46256896] via 10.1.21.1, 00:00:07, Dialer1
D EX 10.1.5.0/24 [170/2223616] via 10.1.25.2, 00:00:07, Serial1
D EX 10.1.4.0/24 [170/46256896] via 10.1.21.1, 00:00:10, Dialer1
C 10.1.25.0/30 is directly connected, Serial1
C 10.1.21.0/24 is directly connected, Dialer1
C 10.1.21.1/32 is directly connected, Dialer1
D EX 10.1.45.0/30 [170/46256896] via 10.1.21.1, 00:00:10, Dialer1
D EX 10.1.34.0/30 [170/46256896] via 10.1.21.1, 00:00:10, Dialer1
D EX 192.168.1.0/24 [170/46256896] via 10.1.21.1, 00:00:10, Dialer1
r2#

```

Dialer watch does not trigger the ISDN link as quickly when you shut down only the frame relay interface on R1. You have to wait while the route completely disappears from the routing table. EIGRP convergence is very fast, but in this case R2 has lost its EIGRP adjacency with R1, so it cannot get the news from R1, and its own frame relay link is still up. If it were another link (running EIGRP) on R1 that had gone down, R2 would know about it immediately.

If you have trouble getting dialer watch to make the call that brings up the link, try rebooting the router on which dialer watch is configured. I found that it helped to reboot R2 after doing "no shut" on R2 s0 and saving the config. This way, the route to R1 lo0 via R2 s0 was in R2's routing table when R2 is booted up.



**7. Configure RIP as shown in the diagram. R6 s0 will not run RIP, but its network will be redistributed into RIP.**

In order for the routers behind R6 to be able to reach AS 7, they need an internal route to the network that will be the next-hop address for any BGP routes advertised from AS 7. If you were not explicitly told what to do with R6 s0 network, reading through the entire scenario to the BGP tasks should make you consider next-hop realities. If you miss the issue in the early stages and configure BGP and then see problems, then remember to think about potential next-hop reachability problems.

**8. On Cat, configure redistribution among the three internal gateway protocols.**

The solution config for Cat shows very simple redistribution among the three internal routing protocols. You will note that there are no route-maps or distribute-lists to prevent route feedback. In a scenario like this, they should not be necessary, if all else is configured correctly, because there is only one redistribution point for the IGPs, and we are only redistributing BGP one-way into RIP on R6.

With redistribution, there is the possibility for sub-optimal routing. In this case, all traffic between routing domains is going to go through Cat. Because Cat's ethernet links to the routers are relatively fast, going through Cat does not add much delay to most trips, even though it may add hops. Take a look at the diagram and the various routing tables to see if you can find any suboptimal paths.

Just looking at the diagram, you should be able to see that there would not be serious problems involving the RIP domain, because it connects to the other IGP domains only through Cat. The sub-optimal routing problems, if any, would involve traffic between EIGRP and OSPF. EIGRP and OSPF intersect at R5 as well as at Cat, and there is no redistribution specified on R5. Because R5 participates in both EIGRP and OSPF, it should be able to reach all networks in both domains even if there were no redistribution on Cat. It does need to go through Cat to reach R3 in the same OSPF domain, because of the totally stubby area between R5 and R3.

R2 is in EIGRP only, and this is the source of a sub-optimal routing problem. While R5 knows to take the direct serial link to reach R2 s loopback address, R2 will only know the redistributed route through Cat to reach R5 s loopback. R5 s routing table looks like this (before improvements discussed below):

```
r5#sh ip rou
```

```
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 not set

```

172.16.0.0/24 is subnetted, 1 subnets
O E1   172.16.30.0 [110/30] via 10.10.5.1, 01:44:55, Ethernet0
10.0.0.0/8 is variably subnetted, 17 subnets, 3 masks
D      10.10.1.0/30 [90/2707456] via 10.1.25.1, 01:44:49, Serial0
O      10.10.3.0/30 [110/11] via 10.10.5.1, 01:45:22, Ethernet0
C      10.10.5.0/30 is directly connected, Ethernet0
O E1   10.10.6.0/30 [110/30] via 10.10.5.1, 01:45:22, Ethernet0
D      10.1.12.0/30 [90/2681856] via 10.1.25.1, 01:45:36, Serial0
O      10.1.3.1/32 [110/129] via 10.1.45.1, 01:45:23, Serial1
D      10.1.2.0/24 [90/2297856] via 10.1.25.1, 01:45:37, Serial0
D      10.1.1.0/24 [90/2809856] via 10.1.25.1, 01:44:50, Serial0
O E1   10.1.6.0/24 [110/30] via 10.10.5.1, 01:45:23, Ethernet0
O      10.1.4.1/32 [110/65] via 10.1.45.1, 01:45:23, Serial1
C      10.1.5.0/24 is directly connected, Loopback0
C      10.1.25.0/30 is directly connected, Serial0
O E1   10.1.21.2/32 [110/30] via 10.10.5.1, 01:44:39, Ethernet0
D      10.1.21.0/24 [90/46738176] via 10.1.25.1, 01:45:38, Serial0
D      10.1.21.1/32 [90/46738176] via 10.1.25.1, 01:44:39, Serial0
C      10.1.45.0/30 is directly connected, Serial1
O      10.1.34.0/30 [110/128] via 10.1.45.1, 01:45:23, Serial1
O E1   192.168.1.0/24 [110/30] via 10.10.5.1, 01:45:23, Ethernet0
r5#

```

The route to 10.1.2.0/24 is an internal EIGRP route. In R2's routing table, the route to R5's loopback is an external EIGRP route redistributed from OSPF, and a host route (/32 mask) at that:

r2#sh ip rou

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 not set

```

172.16.0.0/24 is subnetted, 1 subnets
D EX   172.16.30.0 [170/2249216] via 10.1.12.1, 03:05:34, Serial0
10.0.0.0/8 is variably subnetted, 16 subnets, 3 masks
D      10.10.1.0/30 [90/2195456] via 10.1.12.1, 03:05:43, Serial0
D EX   10.10.3.0/30 [170/2249216] via 10.1.12.1, 03:05:37, Serial0
D EX   10.10.5.0/30 [170/2249216] via 10.1.12.1, 03:05:37, Serial0
D EX   10.10.6.0/30 [170/2249216] via 10.1.12.1, 03:05:37, Serial0
C      10.1.12.0/30 is directly connected, Serial0

```

```

D EX 10.1.3.1/32 [170/2249216] via 10.1.12.1, 03:04:46, Serial0
C 10.1.2.0/24 is directly connected, Loopback0
D 10.1.1.0/24 [90/2297856] via 10.1.12.1, 03:06:04, Serial0
D EX 10.1.6.0/24 [170/2249216] via 10.1.12.1, 03:05:35, Serial0
D EX 10.1.4.1/32 [170/2249216] via 10.1.12.1, 03:04:46, Serial0
D EX 10.1.5.1/32 [170/2249216] via 10.1.12.1, 03:04:47, Serial0
C 10.1.25.0/30 is directly connected, Serial1
C 10.1.21.0/24 is directly connected, Dialer1
C 10.1.21.1/32 is directly connected, Dialer1
D EX 10.1.45.0/30 [170/2249216] via 10.1.12.1, 03:04:47, Serial0
D EX 10.1.34.0/30 [170/2249216] via 10.1.12.1, 03:04:47, Serial0
D EX 192.168.1.0/24 [170/2249216] via 10.1.12.1, 03:05:36, Serial0
r2#

```

The detour through Cat will not add much delay, but it is a suboptimal path, nevertheless. There will be asymmetric routing between R2 and R5's loopback. You could still ping or traceroute from R2 to R5 if you used R5's closest serial interface as the destination, because it is directly connected to R2.

```
r2#trace 10.1.25.2
```

```
Type escape sequence to abort.
Tracing the route to 10.1.25.2
```

```
 1 10.1.25.2 4 msec * 4 msec
r2#
```

Even with asymmetric routing, I could still open a telnet session between R2 and R5's loopback, so connectivity was adequate to support an application layer protocol.

An imperfect network may be okay for some applications. However, in this scenario, we are asked to implement multicast routing with R2's loopback0 as the multicast source and rendezvous point. Reverse path forwarding (RPF) means that any multicast traffic that does not come in on the link that a given router would use to reach the source will be rejected by that router (or will not be forwarded by that router). If R5 expects traffic from R2's loopback to take the direct serial link between R2 and R5, then it will reject packets from R2 that come in on any other link.

It would appear that you could fix the problem by redistributing R5's loopback into EIGRP as a connected route, like so:

```

router eigrp 20
 redistribute connected metric 1500 20 255 1 1500 route-map looponly

[in global config mode]
route-map looponly permit 10
 match interface Loopback0

```

Unfortunately, that is not enough. If you redistribute the connected route into EIGRP, it will go in as a /24 route, while OSPF will produce a /32 route. Because both routes are redistributed into EIGRP, they will have equal administrative distance. It does not matter which route has the better EIGRP metric, the /32 route from OSPF will be chosen by R2 as the best route because /32 is a longer match than /24. **Longest match trumps metric and administrative distance.** You need to go back and put "ip ospf network point-to-point" on your OSPF loopback interfaces to eliminate these host routes. Since several loopback addresses in the OSPF routing domain will be multicast receivers, you may as well add this statement to all of them.

### **9. Configure BGP peering between R6 and R7.**

Because there is no internal routing protocol that would permit a TCP session between R6's loopback and R7's loopback, peering will need to be from one end of the HDLC link to the other. Besides, since there is only one path between R6 and R7, peering loopback-to-loopback would not add any redundancy, anyway.

### **10. Advertise R7's loopback addresses in BGP. Use a prefix-list on R7 to make sure that only the 172.16.30.0/24 route is advertised to AS 6.**

This is pretty simple. You use network statements to put the networks into BGP, and then use a simple one-line prefix-list to permit only the one route to be advertised to the peer in AS 6.

By the way, it is better to filter BGP routes before sending them to a peer than to waste bandwidth by sending all the routes and then filtering them as they arrive at the peer. Here, the difference between three routes and one route does not save a lot of bandwidth, but it is the principle that counts.

### **11. Redistribute BGP routes into RIP on R6.**

We all know that advertising all BGP routes into an IGP is not usually advisable in the real world. Still, this is a practice scenario, not real life, and we have filtered the BGP routes before they are advertised to R6. We have filtered them down to one BGP route, in fact.

### **12. R6 will be a firewall router. Configure it to permit inside users to initiate outgoing TCP and UDP sessions, and to let in return traffic belonging to those sessions, as well as any necessary routing protocol traffic. Permit ping testing from networks behind the firewall router to networks outside the firewall router.**

Reflexive access-lists work the way a Pix firewall is often configured to work, in that they permit inside users to initiate outgoing sessions and then permit the return traffic back in. I did a ping from R4 to 172.16.30.1 on R7, and look at the access-list that R6 generated.

```
r6#sh access-list
Standard IP access list incogniti
  deny 192.168.1.0, wildcard bits 0.0.0.255 (188 matches) check=18
  permit any (18 matches)
Extended IP access list firewallin
  permit tcp any any eq bgp (189 matches)
  permit tcp any eq bgp any
  evaluate outboundssessions
Extended IP access list firewallout
  permit tcp any any reflect outboundssessions
  permit udp any any reflect outboundssessions
  permit icmp any any reflect outboundssessions
Reflexive IP access list outboundssessions
  permit icmp host 172.16.30.1 host 192.168.1.2 (10 matches) (time
left 283)
```

Look at the temporary hole created in the firewall for return ICMP traffic from 172.16.30.1 to the s0 interface on R6. This interface's address is the destination address because of NAT, which will be discussed below.

**13. Configure NAT on R6 so that all outgoing packets that pass through R6 e0 will have the IP address of R6 s0 as their source IP address.**

The nice thing about using NAT on R6 is that AS 7 will not need routes to addresses to subnets behind R6 in order to be able to send return traffic to those subnets.

When all outgoing traffic uses one IP address, we may loosely refer to it as NAT, but it is really Port Address Translation (PAT). If you do a "sh ip nat trans" when there are many outgoing sessions, you will see one IP address with a lot of ports.

I found that configuring NAT using the serial interface as the outside address broke the BGP peering. The router insisted on NATing the BGP session traffic, even though it was ostensibly sourced from the serial interface that serves as the NAT outside address. The NAT process would change the source port of the BGP packets as they left the serial interface, which would break the BGP peering, because the remote peer would reject the source port number. I fixed the problem by excluding R6 s0 s IP address from the NAT inside source list. This is one of those examples of how configuring one simple task near the end of a scenario can break the network in a way that you might not even spot unless you do careful testing.

**14. Insure that R4 can telnet to R7 s 172.16.30.1 loopback address.**

You need to enable login and provide a telnet password under some VTY line(s) for telnet to work. For speed, I have chosen not to use enable passwords or enable secret passwords in these scenarios. A blank password will not work for telnet sessions. Therefore, if you want to do any serious configuration using a Telnet session, you will need to configure real passwords.

**15. Configure IP multicast using a shared tree. R2 lo0 will serve as the source of multicast traffic for testing purposes, and will also serve as the rendezvous point. Have the lo0 interfaces of R1, R3, R4, and R5 receive multicast traffic addressed to 235.0.0.1. Test your configuration by pinging 235.0.0.1 from R2 s lo0. The ISDN links will not participate in multicast routing.**

This task created issues that we discussed in the mutual redistribution task.

At one time I thought there would be a problem for multicast traffic destined for R4 lo0 in the totally stubby OSPF area. R4 had two equal cost default routes that lead out of the totally stubby area. I thought that to eliminate reverse-path forwarding issues, it would help to raise the OSPF cost of one of the links on R4 to prevent load balancing between the two default routes. I configured on R4 s0 "ip ospf cost 130." It turned out that this was not necessary, because the multicast source on R2 will always use the path to R4 lo0 going through R4 s1. If you ever do run into a situation where you need to manually set the reverse-path manually, you can use a static mroute or you can adjust the cost of a link.

The IOS syntax aspect of configuring multicast is pretty simple, and PIM with a static rendezvous point is extremely simple to configure. What makes multicast challenging is spotting the issues and working through them.

In this case, I configured pim on the multicast source interface itself, because it was also the rendezvous point. It is not always necessary for the multicast source interface to have multicast configured on it. For instance, one router could have no multicast routing configured on it, and still have an interface serve as a multicast source. Interfaces on router(s) connected to this interface would need to have multicast enables on them in order to forward multicast packets.

Here is the output from a ping test sourced from the multicast source interface.

```
r2#ping
Protocol [ip]:
Target IP address: 235.0.0.1
Repeat count [1]: 3
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Interface [All]: loopback0
Time to live [255]:
Source address: 10.1.2.1
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 235.0.0.1, timeout is 2 seconds:
```

```
Reply to request 0 from 10.1.25.2, 20 ms [R5]
Reply to request 0 from 10.10.3.2, 160 ms [R3]
Reply to request 0 from 10.1.12.1, 156 ms [R1]
Reply to request 0 from 10.10.3.2, 68 ms [R3]
Reply to request 0 from 10.1.12.1, 60 ms [R1]
Reply to request 0 from 10.1.45.1, 28 ms [R4]
Reply to request 1 from 10.1.25.2, 8 ms
Reply to request 1 from 10.10.3.2, 44 ms
Reply to request 1 from 10.10.3.2, 40 ms
Reply to request 1 from 10.1.12.1, 40 ms
Reply to request 1 from 10.1.12.1, 36 ms
Reply to request 1 from 10.1.45.1, 12 ms
Reply to request 2 from 10.1.25.2, 8 ms
Reply to request 2 from 10.10.3.2, 28 ms
Reply to request 2 from 10.10.3.2, 24 ms
Reply to request 2 from 10.1.12.1, 20 ms
Reply to request 2 from 10.1.12.1, 16 ms
Reply to request 2 from 10.1.45.1, 12 ms
r2#
```

**(Router numbers added in brackets.)**

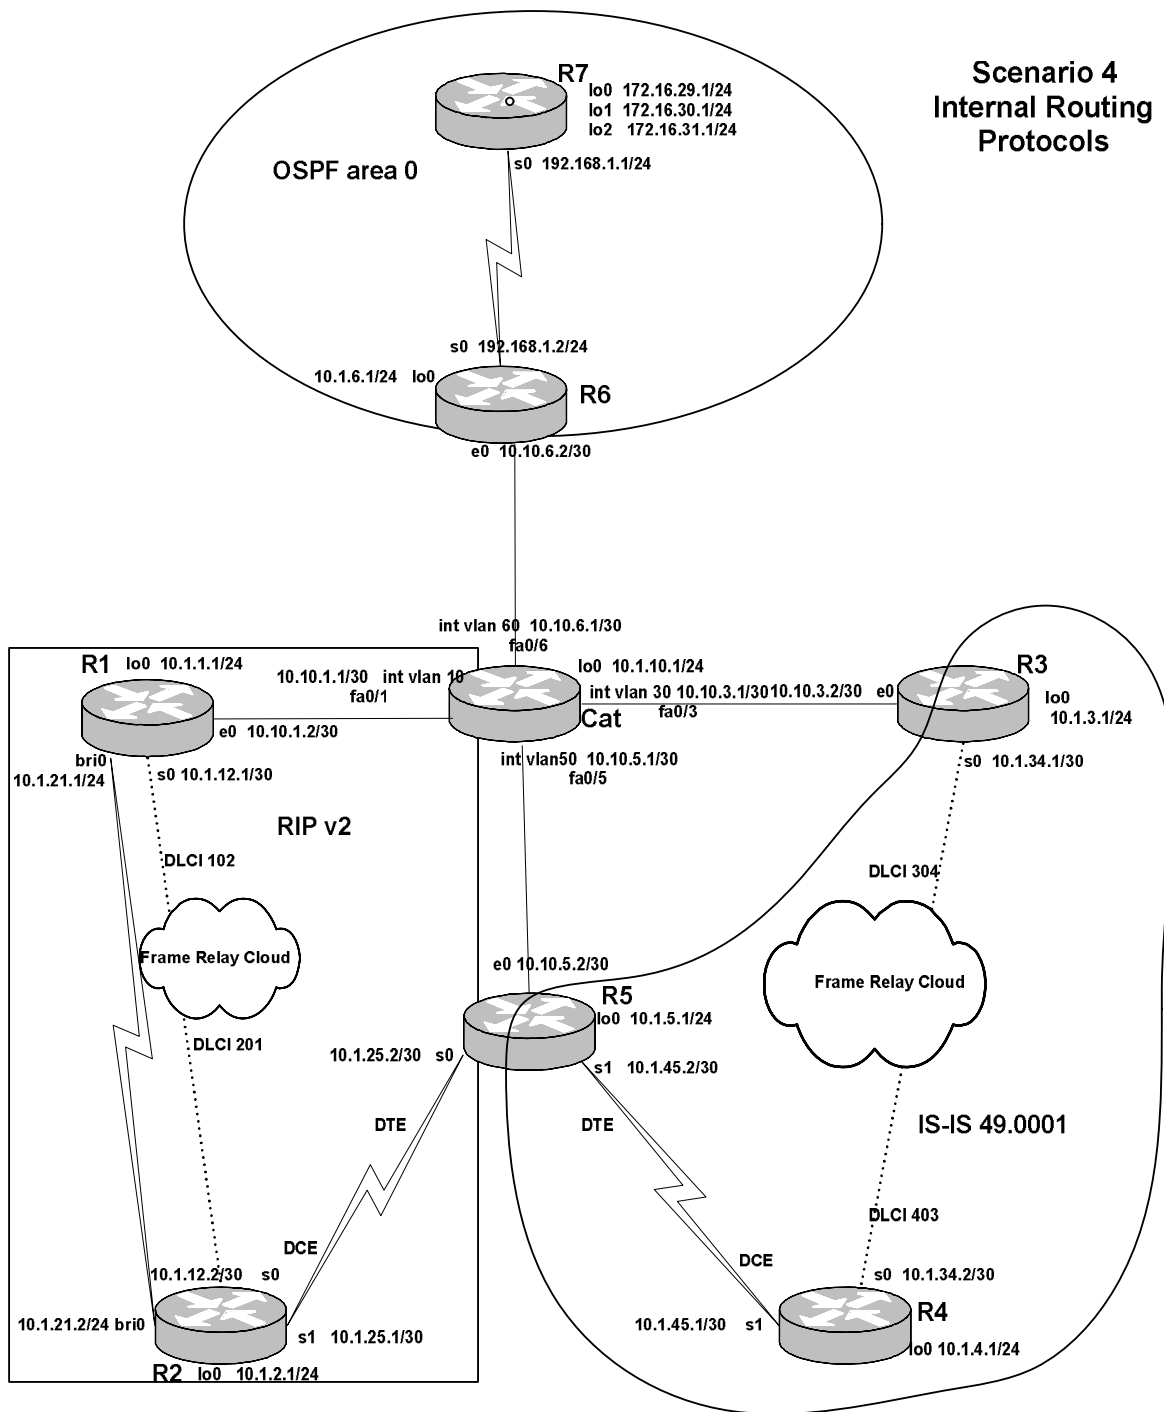
Note that the responses come from the closest interface on the various routers, rather than from the loopback addresses of those routers. Often, you will see more than one response from a given router. This is not necessarily a sign of any problem.



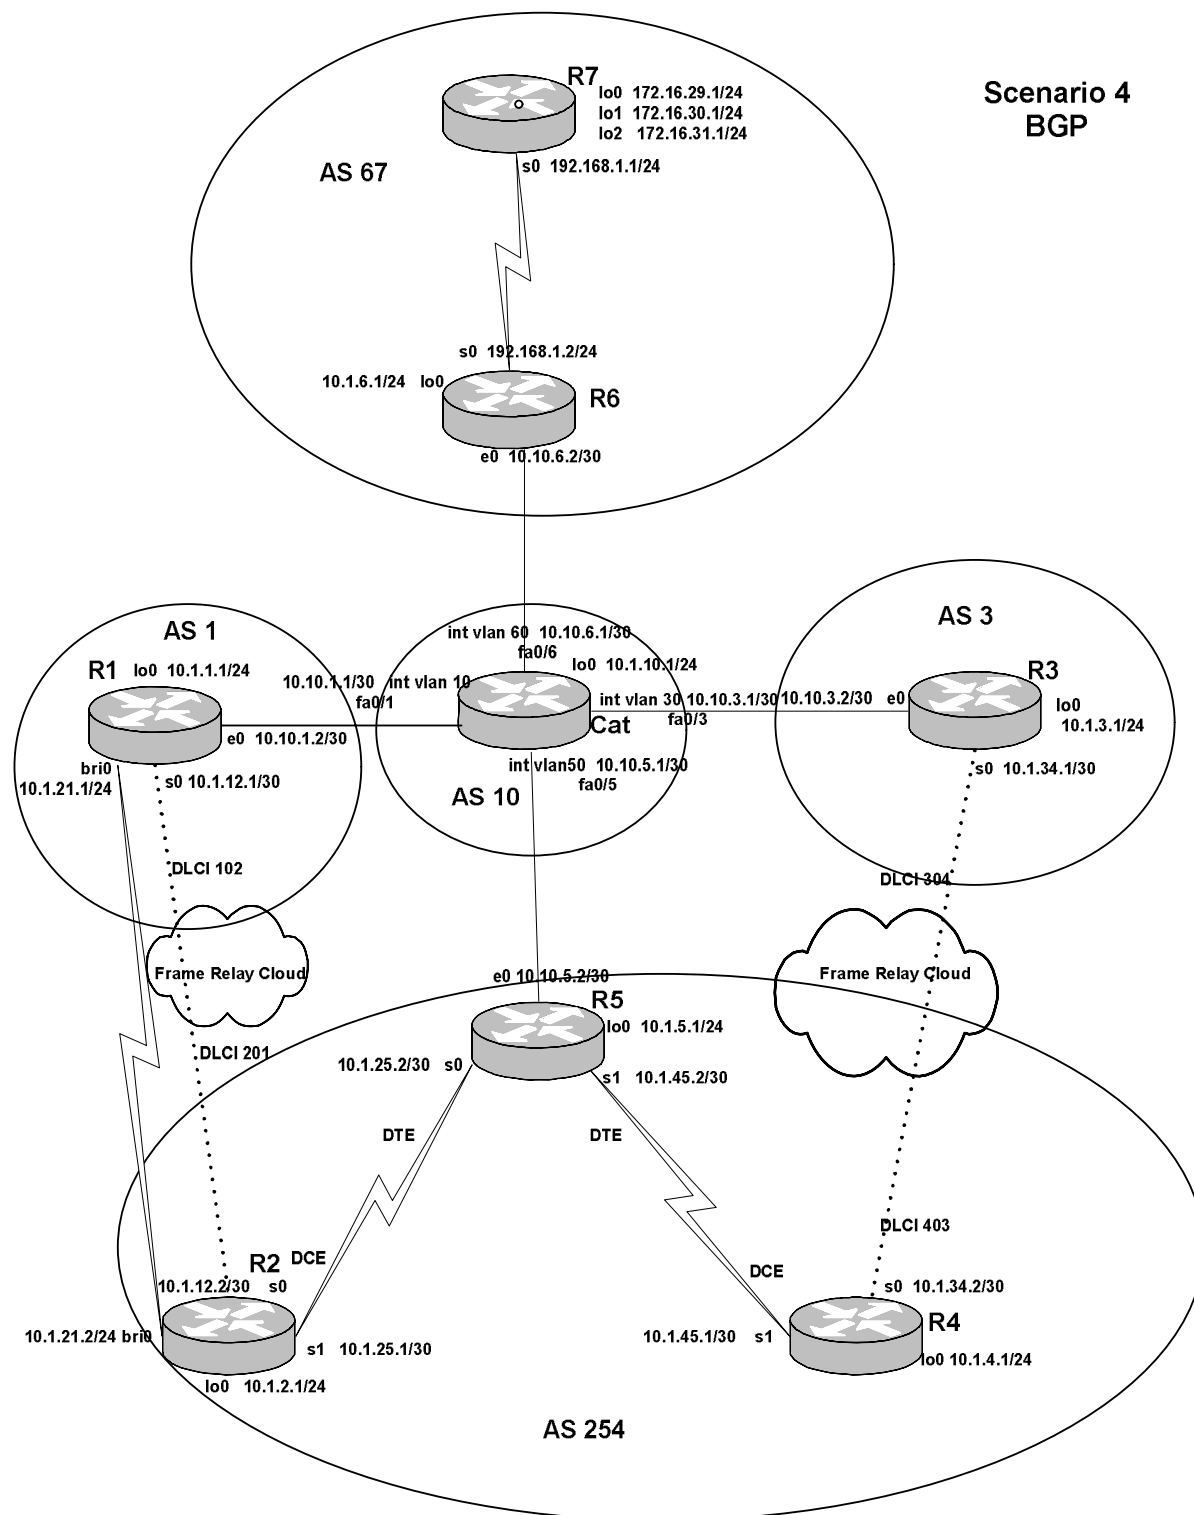
# Chapter Six

## Scenario Four

**Scenario 4  
Internal Routing  
Protocols**



**Scenario 4  
BGP**



Scenario Four requires careful configuration, but be prepared for the result to be less than any-to-any connectivity. The BGP autonomous systems do not correspond neatly with the internal routing domains, which makes it an interesting learning exercise. If, after completing the tasks, you would like to add greater connectivity and a layer of difficulty to the scenario, use network statements to advertise more and more networks in BGP until you reach any-to-any connectivity.

This scenario uses the same topology, including IP addresses, as Scenario 3, with the addition of a loopback address on Cat of 10.1.10.1/24.

The VLANs are as follows:

Router	Interface	VLAN #	Switch port # (fastether 0/x)
R1	e0	10	1
R3	e0	30	3
R5	e0	50	5
R6	e0	60	6

## Tasks

1. Configure the router ethernet links and switch VLANs and VLAN interfaces. Configure the loopback addresses. Manually configure a MAC address of 0001.0001.0001 on R1 e0, and configure the switch to shut down fa0/1 if an ethernet interface with any other MAC address is connected to that switch port.
2. Configure the frame-relay links.
3. Configure the remaining serial links with HDLC encapsulation.
4. Configure OSPF as shown in the Internal Routing Protocols diagram. Use one OSPF network statement for all loopback addresses on R7. Insure that OSPF does not advertise any host (/32) routes.
5. Configure IS-IS area 0001. The NET will be 49.0001.nnnn.nnnn.nnnn.00, where n=router number. Do not run IS-IS on interfaces where there is no other IS-IS router with which to form an adjacency. Redistribute appropriate connected networks into IS-IS.
6. Configure RIP version 2. Use floating static routes so that R1 can use the ISDN link to reach R2 s loopback0 interface when the R1 s0 interface fails, and thereby maintain the BGP peering between R1 and R2.
7. No internal routing protocol will run on the Cat-R6 link, the Cat-R3 link, or the Cat-R5 link. Redistribute between RIP and IS-IS on R5, and redistribute connected

routes into internal routing protocols wherever appropriate. BGP peering should be from loopback to loopback where possible.

8. Configure BGP on R6 and R7 in AS 67, and configure R6 to peer with Cat.
9. Configure BGP peering between R1 in AS 1 and Cat in AS 10.
10. Configure BGP on R2, R5, and R4 in AS 254. Configure peering between R1 and R2, R5 and Cat, and R4 and R3. R3 will be in AS 3.
11. Configure BGP peering between R3 in AS 3 and Cat in AS 10.
12. Advertise all of R7's loopback interfaces in BGP. Advertise R4 s0 in BGP.
13. Configure a BGP attribute so that any BGP speaker in AS 254 will prefer the R3-R4 serial link for traffic destined for R7 lo0. It should be possible to ping from R4 s0 to any loopback address on R7 and vice versa.
14. Configure queueing on the R4 s0 HDLC link so that Telnet traffic always goes first, web traffic goes when the queue for Telnet is empty, and all other traffic waits until both of these two queues are empty.

## Solution Configuration Scripts

### R1

```
r1#sh runn
Building configuration...

Current configuration : 2446 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r1
!
logging rate-limit console 10 except errors
!
username r2 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.0
!
interface Ethernet0
 mac-address 0001.0001.0001
 ip address 10.10.1.2 255.255.255.252
!
interface Serial0
 ip address 10.1.12.1 255.255.255.252
 encapsulation frame-relay
 ip split-horizon
 frame-relay map ip 10.1.12.1 102 broadcast
 frame-relay map ip 10.1.12.2 102 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 no ip address
 shutdown
!
interface BRI0
```

```

ip address 10.1.21.1 255.255.255.0
encapsulation ppp
dialer map ip 10.1.21.2 name r2 broadcast 4082222222
dialer-group 1
isdn switch-type basic-ni
isdn spid1 4081111111 4081111111
isdn spid2 4081111112 4081111111
cdapi buffers regular 0
cdapi buffers raw 0
cdapi buffers large 0
ppp authentication chap
!
router rip
version 2
network 10.0.0.0
distribute-list 12 in BRI0
no auto-summary
!
router bgp 1
no synchronization
bgp log-neighbor-changes
neighbor 10.1.2.1 remote-as 254
neighbor 10.1.2.1 ebgp-multi-hop 255
neighbor 10.1.2.1 update-source Loopback0
neighbor 10.1.10.1 remote-as 10
neighbor 10.1.10.1 ebgp-multi-hop 255
neighbor 10.1.10.1 update-source Loopback0
no auto-summary
!
ip kerberos source-interface any
ip classless
ip route 10.1.2.0 255.255.255.0 10.1.21.2 240
ip http server
!
access-list 12 deny 10.1.2.0 0.0.0.255
access-list 12 deny 10.1.12.0 0.0.0.3
access-list 12 permit any
access-list 101 deny udp any any eq rip
access-list 101 permit ip any any
dialer-list 1 protocol ip list 101
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end r1#

```

## R2

```
r2#sh runn
Building configuration...

Current configuration : 2127 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r2
!
logging rate-limit console 10 except errors
!
username r1 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Loopback0
 ip address 10.1.2.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 10.1.12.2 255.255.255.252
 encapsulation frame-relay
 ip split-horizon
 frame-relay map ip 10.1.12.1 201 broadcast
 frame-relay map ip 10.1.12.2 201 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 ip address 10.1.25.1 255.255.255.252
 clockrate 1300000
!
interface BRI0
 ip address 10.1.21.2 255.255.255.252
 encapsulation ppp
 dialer map ip 10.1.21.1 name r1 broadcast
 isdn switch-type basic-ni
 isdn spid1 40822222221 4082222222
```



```

isdn spid2 4082222222 4082222222
cdapi buffers regular 0
cdapi buffers raw 0
cdapi buffers large 0
ppp authentication chap
!
router rip
  version 2
  network 10.0.0.0
  no auto-summary
!
router bgp 254
  no synchronization
  bgp log-neighbor-changes
  neighbor 10.1.1.1 remote-as 1
  neighbor 10.1.1.1 ebgp-multi-hop 255
  neighbor 10.1.1.1 update-source Loopback0
  neighbor 10.1.5.1 remote-as 254
  neighbor 10.1.5.1 update-source Loopback0
  no auto-summary
!
ip kerberos source-interface any
ip classless
ip route 10.1.1.0 255.255.255.0 10.1.12.1 250
ip http server
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r2#
r2#

```

### R3

```

r3#sh runn
Building configuration...

Current configuration : 1712 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r3

```

```
!  
logging rate-limit console 10 except errors  
!  
ip subnet-zero  
no ip finger  
no ip domain-lookup  
!  
no ip dhcp-client network-discovery  
!  
!  
!  
interface Loopback0  
 ip address 10.1.3.1 255.255.255.0  
!  
interface Ethernet0  
 ip address 10.10.3.2 255.255.255.252  
!  
interface Serial0  
 ip address 10.1.34.1 255.255.255.252  
 ip router isis  
 encapsulation frame-relay  
 frame-relay map clns 304 broadcast  
 frame-relay map ip 10.1.34.1 304 broadcast  
 frame-relay map ip 10.1.34.2 304 broadcast  
 no frame-relay inverse-arp  
 frame-relay lmi-type cisco  
!  
interface Serial1  
 no ip address  
 shutdown  
!  
router isis  
 redistribute connected metric 10  
 net 49.0001.3333.3333.00  
!  
router bgp 3  
 no synchronization  
 bgp log-neighbor-changes  
 neighbor 10.1.4.1 remote-as 254  
 neighbor 10.1.4.1 ebgp-multi-hop 255  
 neighbor 10.1.4.1 update-source Loopback0  
 neighbor 10.10.3.1 remote-as 10  
 no auto-summary  
!  
ip kerberos source-interface any  
ip classless  
ip http server  
!  
!  
line con 0  
 exec-timeout 0 0  
 transport input none  
line aux 0
```

```
line vty 0 4
  login
!
end
r3#
```

## R4

```
r4#sh runn
Building configuration...

Current configuration : 2261 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r4
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
 ip address 10.1.4.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Serial0
 ip address 10.1.34.2 255.255.255.252
 ip router isis
 encapsulation frame-relay
 priority-group 1
 frame-relay map clns 403 broadcast
 frame-relay map ip 10.1.34.1 403 broadcast
 frame-relay map ip 10.1.34.2 403 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 ip address 10.1.45.1 255.255.255.252
 ip router isis
```

```

clockrate 1300000
!
interface TokenRing0
no ip address
shutdown
!
router isis
redistribute connected metric 10 route-map looponly
net 49.0001.4444.4444.4444.00
!
router bgp 254
no synchronization
bgp log-neighbor-changes
network 10.1.34.0 mask 255.255.255.252
neighbor 10.1.3.1 remote-as 3
neighbor 10.1.3.1 ebgp-multi-hop 255
neighbor 10.1.3.1 update-source Loopback0
neighbor 10.1.3.1 route-map setpreflo0 in
neighbor 10.1.5.1 remote-as 254
neighbor 10.1.5.1 update-source Loopback0
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 70 permit 172.16.29.0 0.0.0.255
priority-list 1 protocol ip high tcp telnet
priority-list 1 protocol ip medium tcp www
route-map looponly permit 10
match interface Loopback0
!
route-map setpreflo0 permit 10
match ip address 70
set local-preference 300
!
route-map setpreflo0 permit 20
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end
r4#

```

## R5

```
r5#s
```

Building configuration...

Current configuration : 1906 bytes

```
!  
version 12.2  
no service single-slot-reload-enable  
service timestamps debug uptime  
service timestamps log uptime  
no service password-encryption  
!  
hostname r5  
!  
logging rate-limit console 10 except errors  
!  
ip subnet-zero  
no ip finger  
no ip domain-lookup  
!  
no ip dhcp-client network-discovery  
!  
!  
!  
interface Loopback0  
 ip address 10.1.5.1 255.255.255.0  
!  
interface Ethernet0  
 ip address 10.10.5.2 255.255.255.252  
!  
interface Serial0  
 ip address 10.1.25.2 255.255.255.252  
 no fair-queue  
!  
interface Serial1  
 ip address 10.1.45.2 255.255.255.252  
 ip router isis  
!  
router isis  
 redistribute connected metric 10  
 redistribute rip metric 30  
 net 49.0001.5555.5555.00  
!  
router rip  
 version 2  
 redistribute connected metric 1 route-map looponly  
 redistribute isis level-2 metric 2  
 passive-interface default  
 no passive-interface Serial0  
 network 10.0.0.0  
 no auto-summary  
!  
router bgp 254  
 no synchronization  
 bgp log-neighbor-changes
```

```

neighbor 10.1.2.1 remote-as 254
neighbor 10.1.2.1 update-source Loopback0
neighbor 10.1.4.1 remote-as 254
neighbor 10.1.4.1 update-source Loopback0
neighbor 10.1.4.1 route-reflector-client
neighbor 10.10.5.1 remote-as 10
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
route-map looonly permit 10
  match interface Loopback0
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end
r5#

```

## R6

```

r6#sh runn
Building configuration...

Current configuration : 1702 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r6
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0

```

```

ip address 10.1.6.1 255.255.255.0
ip ospf network point-to-point
!
interface Ethernet0
ip address 10.10.6.2 255.255.255.252
!
interface Ethernet1
no ip address
shutdown
!
interface Serial0
ip address 192.168.1.2 255.255.255.0
no fair-queue
clockrate 1300000
!
interface Serial1
no ip address
shutdown
!
router ospf 64
log-adjacency-changes
redistribute connected subnets route-map ether
network 10.1.6.1 0.0.0.0 area 0
network 192.168.1.2 0.0.0.0 area 0
!
router bgp 67
no synchronization
bgp log-neighbor-changes
neighbor 10.10.6.1 remote-as 10
neighbor 172.16.29.1 remote-as 67
neighbor 172.16.29.1 update-source Loopback0
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
route-map ether permit 10
match interface Ethernet0
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end
r6#

```

## R7

```
r7#sh runn
Building configuration...

Current configuration : 1841 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r7
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
 ip address 172.16.29.1 255.255.255.0
 ip ospf network point-to-point
!
interface Loopback1
 ip address 172.16.30.1 255.255.255.0
 ip ospf network point-to-point
!
interface Loopback2
 ip address 172.16.31.1 255.255.255.0
 ip ospf network point-to-point
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.1 255.255.255.0
 no fair-queue
!
interface Serial1
 no ip address
 shutdown
!
router ospf 64
 log-adjacency-changes
```



```

network 172.16.28.0 0.0.3.255 area 0
network 192.168.1.1 0.0.0.0 area 0
!
router bgp 67
no synchronization
bgp log-neighbor-changes
network 172.16.29.0 mask 255.255.255.0
network 172.16.30.0 mask 255.255.255.0
network 172.16.31.0 mask 255.255.255.0
neighbor 10.1.6.1 remote-as 67
neighbor 10.1.6.1 update-source Loopback0
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end
r7#

```

## Cat

```

cat#sh runn
Building configuration...

Current configuration : 3073 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname cat
!
!
ip subnet-zero
ip routing
no ip domain-lookup
!
!
spanning-tree portfast default
spanning-tree extend system-id
!

```

```
!  
!  
interface Loopback0  
 ip address 10.1.10.1 255.255.255.0  
!  
interface FastEthernet0/1  
 switchport access vlan 10  
 switchport mode access  
 switchport port-security  
 switchport port-security mac-address 0001.0001.0001  
 no ip address  
!  
interface FastEthernet0/2  
 no ip address  
!  
interface FastEthernet0/3  
 switchport access vlan 30  
 no ip address  
!  
interface FastEthernet0/4  
 no ip address  
!  
interface FastEthernet0/5  
 switchport access vlan 50  
 no ip address  
!  
interface FastEthernet0/6  
 switchport access vlan 60  
 no ip address  
!  
[output truncated for brevity]  
interface Vlan1  
 no ip address  
 shutdown  
!  
interface Vlan10  
 ip address 10.10.1.1 255.255.255.252  
!  
interface Vlan30  
 ip address 10.10.3.1 255.255.255.252  
!  
interface Vlan50  
 ip address 10.10.5.1 255.255.255.252  
!  
interface Vlan60  
 ip address 10.10.6.1 255.255.255.252  
!  
router rip  
 version 2  
 redistribute connected metric 1 route-map looponly  
 passive-interface default  
 no passive-interface Vlan10  
 network 10.0.0.0  
 no auto-summary
```

```
!  
router bgp 10  
  no synchronization  
  bgp log-neighbor-changes  
  neighbor 10.1.1.1 remote-as 1  
  neighbor 10.1.1.1 ebgp-multi-hop 255  
  neighbor 10.1.1.1 update-source Loopback0  
  neighbor 10.10.3.2 remote-as 3  
  neighbor 10.10.5.2 remote-as 254  
  neighbor 10.10.6.2 remote-as 67  
  no auto-summary  
!  
ip classless  
ip http server  
!  
!  
!  
route-map looponly permit 10  
  match interface Loopback0  
!  
!  
line con 0  
  exec-timeout 0 0  
line vty 0 4  
  login  
line vty 5 15  
  login  
!  
end  
cat#
```

## Explanation

**1. Configure the router ethernet links and switch VLANs and VLAN interfaces. Configure the loopback addresses. Manually configure a MAC address of 0001.0001.0001 on R1 e0, and configure the switch to shut down fa0/1 if an ethernet interface with any other MAC address is connected to that switch port.**

The tricky part of setting up port-security on a Cat 3550 switch is that you need to type "switchport mode access" on the interface, even if you have already manually set the VLAN for the interface. If there is a violation, and you use the default violation policy of shutting down the interface, the interface will remain shutdown until you manually reset it. The easiest way to reset it, and the simplest to remember, is to type "shut" and then "no shut" on the interface. Simply typing "no shut" will not bring the interface back up. You can also manually shut down the interface to remove a MAC address from the dynamic list of permitted MAC addresses. You may need to do this if you have enabled port security on a port before specifying a static MAC address. A dynamically learned MAC address may take up the whole table (default value is one permitted MAC address), so you need to clear the table using "shutdown" before the switch will let you statically specify the same address.

Remember to include the Cat loopback address that has been added in this scenario. You should enable "spanning-tree portfast default." Be careful to use the correct masks for the ethernet addresses, as they are unlikely masks for ethernet interfaces.

**2. Configure the frame-relay links.**

Remember to use static frame-relay mapping of IP addresses and to disable inverse arp.

**3. Configure the remaining serial links with HDLC encapsulation.**

HDLC is the default encapsulation for synchronous serial interfaces, so HDLC may not always be specified on the network diagrams in practice scenarios.

**4. Configure OSPF as shown in the Internal Routing Protocols diagram. Use one OSPF network statement for all loopback addresses on R7. Insure that OSPF does not advertise any host (/32) routes.**

Setting the OSPF network type for the loopback interfaces to point-to-point eliminates the advertisement by OSPF of host routes for these interfaces.

**5. Configure IS-IS area 0001. The NET will be 49.0001.nnnn.nnnn.nnnn.00, where n=router number. Do not run IS-IS on interfaces where there is no other IS-IS router with which to form an adjacency. Redistribute appropriate connected networks into IS-IS.**

The main trick here is that you will need to provide a frame relay map statement for CLNS on the frame relay link before ISIS adjacencies will form. When you "debug isis adj" to watch your adjacencies trying to form, and you see "encapsulation failed" on a frame relay interface, think of "frame map clns."

Instead of running IS-IS on the ethernet and loopback interfaces, you will simply redistribute connected networks into IS-IS. If you use a route-map to limit the redistribution of connected networks into IS-IS, you may find this limiting route-map affects redistribution from other routing protocols. Redistribution from RIP into IS-IS on R5 worked when I did not use a limiting route-map for redistributing connected networks into IS-IS. If you are permitted to do so, a simpler way to get these networks into IS-IS is to use "passive-interface e0" under router isis. This will automatically remove any IS-IS configuration from the interface, but its route will be in IS-IS

#### **6. Configure RIP version 2. Use floating static routes so that R1 can use the ISDN link to reach R2's loopback0 interface when the R1 s0 interface fails, and thereby maintain the BGP peering between R1 and R2.**

The floating static route with a high administrative distance of 240 will only trigger the ISDN link when the other route to the loopback address in question disappears from the routing table. We use an access-list with the dialer list to prevent RIP from bringing up the link, but once it is up, the BGP-related TCP traffic will bring up the link.

Go to R1, type "debug ip bgp" and "debug dialer." Then shutdown R1 s0, and watch the BGP traffic bring up the ISDN link within a minute. The link may drop from time to time, but the BGP peering should be maintained.

A distribute-list on R1 to prevent certain RIP routes from entering R1's routing table helped the floating static route configuration to work properly. I did not need a distribute-list on R2 as well, because I configured ISDN so that only R1 dials R2. There is no dial string in the dialer map statement on R2, nor is there any dialer-list on R2 to trigger any dialing.

You can watch RIP's slow convergence after the ISDN link finally drops after the idle-timeout. If you send a packet to one of the networks for which there is both an ISDN route and a frame-relay route, and the ISDN route gets used, that will bring up the ISDN link again. In a real-world network with plenty of traffic, you might want to do something to add to the cost of all routes using the ISDN link. Consider, perhaps, an offset-list or an adjustment to administrative distance for all those routes. This was not required here, however.

You should enable split horizon on the frame-relay links that are running RIP.

#### **7. No internal routing protocol will run on the Cat-R6 link, the Cat-R3 link, or the Cat-R5 link. Redistribute between RIP and IS-IS on R5, and redistribute connected routes into internal routing protocols wherever appropriate. BGP peering should be from loopback to loopback where possible.**

I did redistribution between RIP and IS-IS without using any mechanisms to prevent route feedback. This was possible because there was only one point of redistribution. If I had not enabled split horizon on the frame relay link between R1 and R2, there could have been problems. BGP peering will be from one end of the link to the other on links where no internal routing protocol is running. For all other links, it should be possible to do peering loopback-to-loopback, as long as you redistribute any necessary loopback addresses into the appropriate routing protocol.

Remember to use "no synch" and "no auto" under BGP wherever possible.

### **8. Configure BGP on R6 and R7 in AS 67, and configure R6 to peer with Cat.**

Since OSPF includes the loopback addresses, BGP peering can be between loopbacks. Peering between R6 and Cat will use the ethernet IP addresses.

### **9. Configure BGP peering between R1 in AS 1 and Cat in AS 10.**

Just redistribute the Cat loopback into RIP, and you can peer loopback to loopback. Remember that you do not need "ebgp-multihop" for IGP peering between loopbacks, but you do need to specify the loopback as the source addresses for BGP updates.

### **10. Configure BGP on R2, R5, and R4 in AS 254. Configure peering between R1 and R2, R5 and Cat, and R4 and R3. R3 will be in AS 3.**

R1 and R2 participate in RIP, so peering will be loopback to loopback. Due to IGP redistribution on R5, R2, R5, and R4 all have routes to each other's loopback addresses, so peering will be loopback to loopback. R5 will peer with Cat using the ethernet link. R4 will peer with R3 using loopback addresses.

### **11. Configure BGP peering between R3 in AS 3 and Cat in AS 10.**

Peering cannot use loopbacks here because there is no routing protocol running between Cat and R3 to carry the routes to loopback interfaces.

### **12. Advertise all of R7's loopback interfaces in BGP. Advertise R4 s0 in BGP.**

The simplest way to advertise routes in BGP is to use the network statement. If using the network statement is prohibited, you can redistribute connected routes with a route-map to limit which connected networks are redistributed into BGP.

The main trick with BGP network statements is that the mask needs to match the route precisely for the network to be advertised. The mask for R4 s0 is a /30 mask, which is odd for an ethernet interface. It would be easy to make the mistake of using a /24 mask, and then waste a minute or two wondering why the route does not appear in your BGP tables.

**13. Configure a BGP attribute so that any BGP speaker in AS 254 will prefer the R3-R4 serial link for traffic destined for R7 lo0. It should be possible to ping from R4 s0 to any loopback address on R7 and vice versa.**

On R4, we set a high local preference value of 300 for the BGP route to R7 lo0 as it comes in from R3. Here is R5's BGP table:

```
r5#sh ip bgp
BGP table version is 6, local router ID is 10.1.5.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*>i 10.1.34.0/30    10.1.4.1           0     100     0 i
*>i 172.16.29.0/24 10.1.3.1           300    0 3 10 67 i
*                  10.10.5.1         0     100     0 i 10 67 i
*> 172.16.30.0/24  10.10.5.1         0     100     0 i 10 67 i
*> 172.16.31.0/24  10.10.5.1         0     100     0 i 10 67 i
r5#
```

As you can see, the "best" BGP route to 172.16.29.0/24 is the route with higher local preference, rather than the one with the shorter AS-Path.

To test, you will need to use extended ping on R4 and R7 to set the source address for each ping. The one exception is that you can use a simple (non-extended) ping from R4 to R7 lo0, because the source interface will be R4 s0. The source address for pings to the other two loopbacks on R7 would be s1, since the AS-Path of the route through AS 10 will be shorter than the AS-Path of the BGP route through AS 3.

When adjusting BGP attributes, do not expect the changes to propagate quickly through the network. Typing "clear ip bgp \* soft in" or the like will not always be enough. Do not hesitate to break all of your BGP sessions to make the changes appear in your BGP tables quickly. This is not a production network, so you can and should be bold. For example, when you adjust local-preference on R4 you may need to "clear ip bgp \*" on R4 and then do the same on R2.

By the same token, you should not hesitate to clear your internal routes when a problem appears. Be on the lookout for "possibly down" in your routing tables. This could be a symptom of a real problem, or it could just be due to your having made a lot of configuration changes in the network in the last few minutes. You should clear your IP routing table using "clear ip rout \*" and see if the problem persists.

**14. Configure queueing on the R4 s0 HDLC link so that Telnet traffic always goes first, web traffic goes when the queue for Telnet is empty, and all other traffic waits until both of these two queues are empty.**

To test priority queueing, go to R4 and type "debug priority" and then telnet to R7 lo0. Remember when configuring priority queueing that you can use the built-in www and telnet options with priority lists, and the router will catch traffic to or from ports 80 and 23, respectively. The debugging output below shows **Telnet** packets using the priority queue:

Trying 172.16.29.1 ... Open

Password required, but none set

```
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 48/0)
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 44/0)
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 53/0)
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 44/0)
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 47/0)
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 47/0)
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 53/0)
02:34:09: PQ: Serial0: ip (tcp 23) -> high
02:34:09: PQ: Serial0 output (Pk size/Q 44/0)
02:34:10: PQ: Serial0 output (Pk size/Q 1505/0)
[Connection to 172.16.29.1 closed by foreign host]
r4#
02:34:11: PQ: Serial0 output (Pk size/Q 1505/0)
02:34:11: PQ: Serial0 output (Pk size/Q 104/0)
02:34:11: PQ: Serial0: ip (tcp 23) -> high
02:34:11: PQ: Serial0 output (Pk size/Q 44/0)
02:34:11: PQ: Serial0: ip (tcp 23) -> high
02:34:11: PQ: Serial0 output (Pk size/Q 44/0)
02:34:11: PQ: Serial0: ip (defaulting) -> normal
```

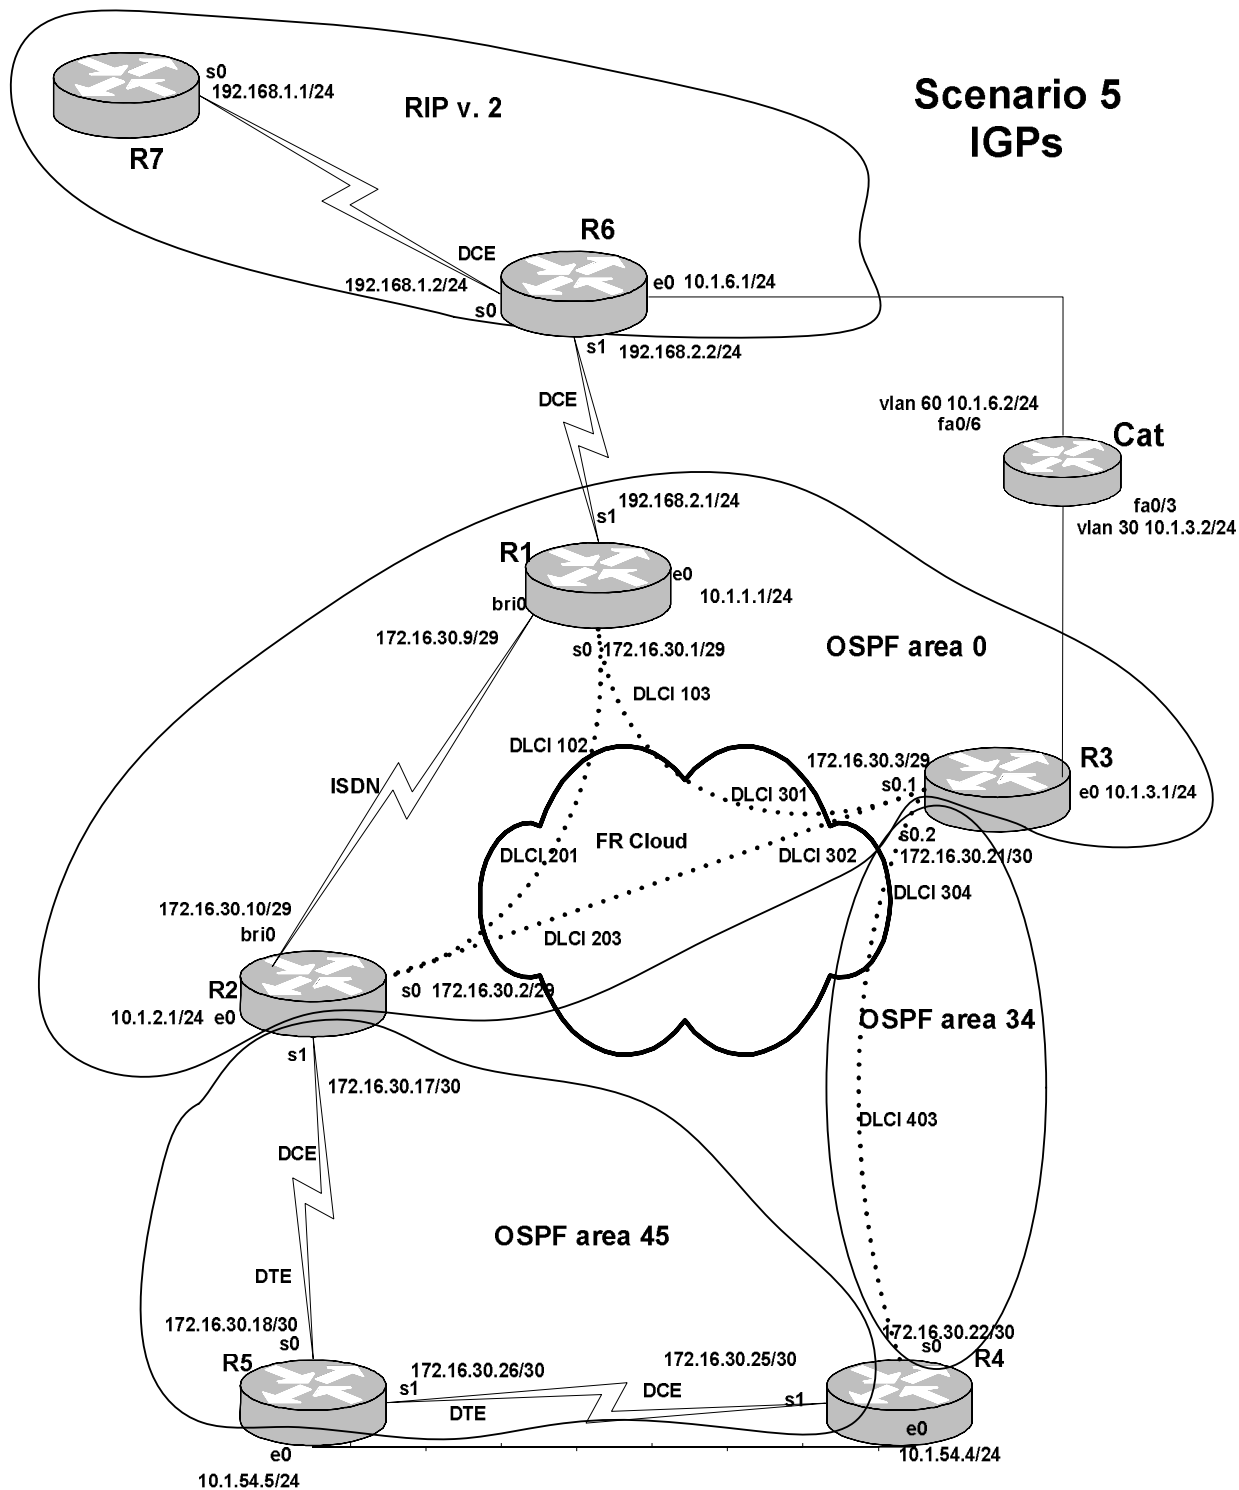
Note that the Telnet packets tend to be pretty small.



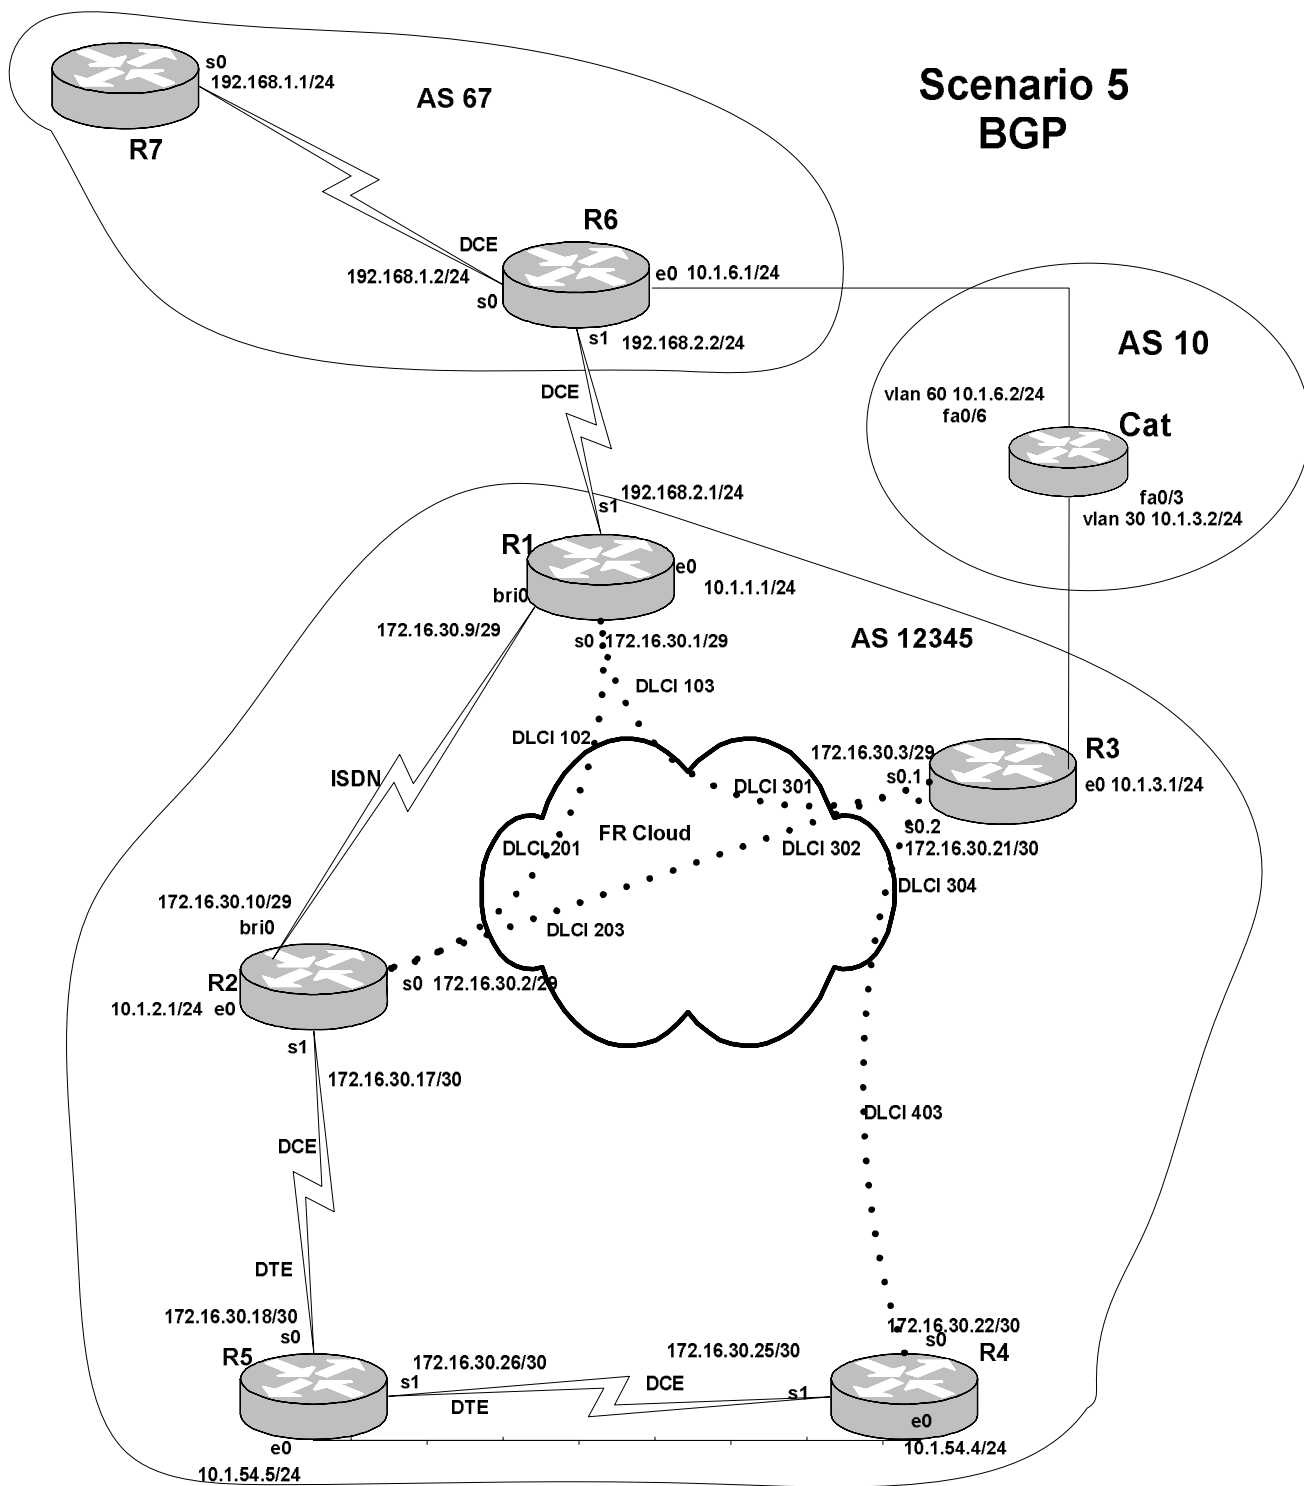
# Chapter Seven

## Scenario Five

### Scenario 5 IGPs



### Scenario 5 BGP



In Scenario Five, we will experiment with mutual redistribution between OSPF and BGP on two routers. This is something you will probably never do in a production network, so enjoy the chance to do it here. We have progressed from the comfortable, realistic-looking network of Scenario One to the unrealistic alternate universe of Scenario Five.

Ironically, Scenario Five uses a similar topology to that used in Scenario One, but with a few alterations. There is an added frame relay link between R2 and R3 to create a full-mesh frame relay cloud. The frame relay switch always offered this link, but we did not use it until now. R4 e0 and R5 e0 are now on the same ethernet, in VLAN 54, so as to permit a Hot Standby Routing Protocol (HSRP) task. In addition, we enable routing on Cat and put Cat (as a BGP-speaking router) between R3 and R6.

Try not to get distracted too much by the odd fact that there is a serial link between two routers that share an ethernet. The ethernet will not participate actively in OSPF. I wanted to provide an HSRP task without collapsing the topology too much. The VLANs are as follows:

Router	Interface	VLAN #	Switch port # (fastether 0/x)
R1	e0	10	1
R2	e0	20	2
R3	e0	30	3
R4	e0	54	4
R5	e0	54	5
R6	e0	60	6

## Tasks

1. Configure a full-mesh of frame relay links among R1, R2, and R3, so that they share the same subnet. R3 will connect to R1 and R2 using its s0.1 interface. Configure a separate frame relay link between R3 and R4. R3 s0.2 will connect to R4 s0. R1 and R2 will not use any frame relay subinterfaces.
2. Configure all other serial links.
3. Configure LAN links, VLANs and VLAN interfaces.
4. Configure RIP version 2 on R7 and R6. Use RIP MD5 authentication on this link. The link between R1 and R6 will not actively participate in any routing protocol.
5. Configure R1, R2 and R3 in OSPF area 0, according to the IGP diagram. The frame relay links among these three routers will use the broadcast network type.

6. Configure OSPF area 34 and area 45. Insure that even if the R2-R5 serial link fails, the end users in OSPF area 45 will still be able to reach R1 e0, R2 e0, and R3 e0. R4 and R5 e0 will not try to form an OSPF adjacency.
7. Configure OSPF demand circuit on the ISDN link between R1 and R2. Use dialer profiles.
8. Configure BGP on R1 and R3 in AS 12345, on R6 and R7 in AS 67, and on Cat in AS 10. Configure peering between R1 and R3, R1 and R6, R3 and Cat, and R6 and Cat.
9. On R6, advertise R6 s0 and R6 e0 in BGP using network statements.
10. Redistribute OSPF into BGP on R1 and R3. Redistribute BGP into OSPF on R1 and R3. Insure that traffic from AS 12345 to destinations in AS 67 prefers the R1-R6 link.
11. On R1 and R3, summarize all the 172.16.30.x routes in AS 12345 to 172.16.30.0/24 so that AS 67 will see the 172.16.30.0/24 route, but will not see the more specific routes in that range. Insure that no other routers in AS 12345 see any route to 172.16.30.0/24. Do not summarize any of the 10.1.x.x routes on R1 or R3.
12. R6 will advertise a 10.1.0.0/16 BGP aggregate, but will suppress only the specific routes 10.1.54.0/24 and 10.1.3.0/24.
13. Take steps to keep the 10.1.0.0/16 and 172.16.30.0/24 summary routes from being advertised from R6 to R1 or Cat.
14. Configure Hot Standby Routing Protocol on R4 e0 and R5 e0 so that R5 will normally serve the hosts on the 10.1.54.0/24 ethernet. If R5 s0 fails, R4 will serve the hosts on the shared ethernet. All end users on the shared ethernet will be using 10.1.54.1/24 as their default gateway.
15. We need to severely limit the bandwidth of traffic coming into R3 e0 from one server. The server's MAC address is 2121.2020.2121. Limit all incoming traffic from that server to a maximum of 160,000 kbps.
16. Configure custom queueing on R5 s0 so that each of four queues will get roughly one-quarter of the bandwidth. Queue 1 will be for Telnet. Queue 2 will be for FTP. Queue 3 will be for Web traffic. Queue 4 will be for everything else. The size of each queue should be 1500 bytes.

## Solution Configuration Scripts

### R1

```

r1#r1#s
Building configuration...

Current configuration : 2649 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r1
!
logging rate-limit console 10 except errors
!
username r2 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Ethernet0
 ip address 10.1.1.1 255.255.255.0
!
interface Serial0
 ip address 172.16.30.1 255.255.255.248
 encapsulation frame-relay
 ip ospf network broadcast
 frame-relay map ip 172.16.30.1 103 broadcast
 frame-relay map ip 172.16.30.2 102 broadcast
 frame-relay map ip 172.16.30.3 103 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 ip address 192.168.2.1 255.255.255.0
!
interface BRI0
 no ip address
 encapsulation ppp
 dialer pool-member 1

```

```

isdn switch-type basic-ni
isdn spid1 4081111111 4081111111
isdn spid2 4081111112 4081111111
cdapi buffers regular 0
cdapi buffers raw 0
cdapi buffers large 0
ppp authentication chap
!
interface Dialer1
ip address 172.16.30.9 255.255.255.248
encapsulation ppp
ip ospf demand-circuit
dialer pool 1
dialer remote-name r2
dialer string 4082222222
dialer-group 1
ppp authentication chap
!
router ospf 64
log-adjacency-changes
redistribute bgp 12345 metric-type 1 subnets route-map blocksummary
passive-interface Serial1
network 10.1.1.1 0.0.0.0 area 0
network 172.16.30.0 0.0.0.255 area 0
network 192.168.2.1 0.0.0.0 area 0
!
router bgp 12345
no synchronization
bgp log-neighbor-changes
aggregate-address 172.16.30.0 255.255.255.0 summary-only
redistribute ospf 64 match internal
neighbor 172.16.30.3 remote-as 12345
neighbor 172.16.30.3 route-map blocksummary out
neighbor 192.168.2.2 remote-as 67
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
ip prefix-list summary seq 10 permit 172.16.30.0/24
dialer-list 1 protocol ip permit
route-map blocksummary deny 10
match ip address prefix-list summary
!
route-map blocksummary permit 20
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login

```

```
!
end
```

```
r1#
```

## R2

```
r2#s
```

```
Building configuration...
```

```
Current configuration : 2134 bytes
```

```
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r2
!
logging rate-limit console 10 except errors
!
username r1 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
isdn switch-type basic-ni
!
!
!
!
interface Ethernet0
 ip address 10.1.2.1 255.255.255.0
!
interface Serial0
 ip address 172.16.30.2 255.255.255.248
 encapsulation frame-relay
 ip ospf network broadcast
 frame-relay map ip 172.16.30.1 201 broadcast
 frame-relay map ip 172.16.30.2 203 broadcast
 frame-relay map ip 172.16.30.3 203 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 ip address 172.16.30.17 255.255.255.252
 clockrate 1300000
!
interface BRI0
 no ip address
 encapsulation ppp
```



```

dialer pool-member 1
isdn switch-type basic-ni
isdn spid1 40822222221 4082222222
isdn spid2 40822222222 4082222222
cdapi buffers regular 0
cdapi buffers raw 0
cdapi buffers large 0
ppp authentication chap
!
interface Dialer1
ip address 172.16.30.10 255.255.255.248
encapsulation ppp
dialer pool 1
dialer remote-name r1
dialer string 4081111111
dialer-group 1
!
router ospf 64
log-adjacency-changes
passive-interface Ethernet0
network 10.1.2.1 0.0.0.0 area 0
network 172.16.30.2 0.0.0.0 area 0
network 172.16.30.10 0.0.0.0 area 0
network 172.16.30.17 0.0.0.0 area 45
!
ip kerberos source-interface any
ip classless
ip http server
!
dialer-list 1 protocol ip permit
!
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end

r2#

```

### R3

```

r3#s
Building configuration...

Current configuration : 2138 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime

```

```
service timestamps log uptime
no service password-encryption
!
hostname r3
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
ip cef
no ip dhcp-client network-discovery
!
!
!
!
interface Ethernet0
 ip address 10.1.3.1 255.255.255.0
 rate-limit input access-group rate-limit 120 160000 30000 60000
 conform-action transmit exceed-action drop
!
interface Serial0
 no ip address
 encapsulation frame-relay
 frame-relay lmi-type cisco
!
interface Serial0.1 multipoint
 ip address 172.16.30.3 255.255.255.248
 ip ospf network broadcast
 frame-relay map ip 172.16.30.1 301 broadcast
 frame-relay map ip 172.16.30.2 302 broadcast
 frame-relay map ip 172.16.30.3 302 broadcast
 no frame-relay inverse-arp
!
interface Serial0.2 point-to-point
 ip address 172.16.30.21 255.255.255.252
 frame-relay interface-dlci 304
!
interface Serial1
 no ip address
 shutdown
!
router ospf 64
 log-adjacency-changes
 area 34 virtual-link 172.16.30.25
 redistribute bgp 12345 metric 500 metric-type 1 subnets
 passive-interface Ethernet0
 network 10.1.3.1 0.0.0.0 area 0
 network 172.16.30.3 0.0.0.0 area 0
 network 172.16.30.21 0.0.0.0 area 34
!
router bgp 12345
 no synchronization
```

```

bgp log-neighbor-changes
redistribute ospf 64 match internal
neighbor 10.1.3.2 remote-as 10
neighbor 172.16.30.1 remote-as 12345
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list rate-limit 120 2121.2020.2121
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  login
!
end

r3#

```

## R4

```

r4#s
Building configuration...

Current configuration : 1671 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r4
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
!
no ip dhcp-client network-discovery
!
!
!
interface Ethernet0
  ip address 10.1.54.4 255.255.255.0
  standby 1 priority 100 preempt
  standby 1 ip 10.1.54.1

```

```

!
interface Serial0
 ip address 172.16.30.22 255.255.255.252
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 172.16.30.21 403 broadcast
 frame-relay map ip 172.16.30.22 403 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
!
interface Serial1
 ip address 172.16.30.25 255.255.255.252
 clockrate 1300000
!
interface TokenRing0
 no ip address
 shutdown
!
router ospf 64
 log-adjacency-changes
 area 34 virtual-link 172.16.30.21
 passive-interface Ethernet0
 network 10.1.54.4 0.0.0.0 area 45
 network 172.16.30.20 0.0.0.3 area 34
 network 172.16.30.25 0.0.0.0 area 45
!
 ip kerberos source-interface any
 ip classless
 ip http server
!
!
 line con 0
   exec-timeout 0 0
   transport input none
 line aux 0
 line vty 0 4
   login
!
end

r4#

```

## R5

```

r5#s
Building configuration...

Current configuration : 1556 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime

```

```
no service password-encryption
!
hostname r5
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
interface Ethernet0
 ip address 10.1.54.5 255.255.255.0
 standby 1 priority 105 preempt
 standby 1 ip 10.1.54.1
 standby 1 track Se0 10
!
interface Serial0
 ip address 172.16.30.18 255.255.255.252
 custom-queue-list 1
!
interface Serial1
 ip address 172.16.30.26 255.255.255.252
!
router ospf 64
 log-adjacency-changes
 passive-interface Ethernet0
 network 10.1.54.5 0.0.0.0 area 45
 network 172.16.30.18 0.0.0.0 area 45
 network 172.16.30.26 0.0.0.0 area 45
!
ip kerberos source-interface any
ip classless
ip http server
!
queue-list 1 protocol ip 1 tcp telnet
queue-list 1 protocol ip 2 tcp ftp
queue-list 1 protocol ip 3 tcp www
queue-list 1 default 4
!
line con 0
 exec-timeout 0 0
 transport input none
line aux 0
line vty 0 4
 login
!
end

r5#
```

**R6**

```
r6#sh runn
Building configuration...

Current configuration : 2264 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname r6
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
key chain ripkey
  key 1
    key-string cisco
!
!
!
!
interface Ethernet0
 ip address 10.1.6.1 255.255.255.0
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.2 255.255.255.0
 ip rip authentication mode md5
 ip rip authentication key-chain ripkey
 no fair-queue
 clockrate 1300000
!
interface Serial1
 ip address 192.168.2.2 255.255.255.0
 clockrate 1300000
!
router rip
 version 2
 passive-interface default
 no passive-interface Serial0
 network 10.0.0.0
```

```

network 192.168.1.0
network 192.168.2.0
no auto-summary
!
router bgp 67
no synchronization
bgp log-neighbor-changes
network 10.1.6.0 mask 255.255.255.0
network 192.168.1.0
aggregate-address 10.1.0.0 255.255.0.0 suppress-map suppress543
neighbor 10.1.6.2 remote-as 10
neighbor 10.1.6.2 prefix-list blocksummary out
neighbor 192.168.1.1 remote-as 67
neighbor 192.168.2.1 remote-as 12345
neighbor 192.168.2.1 prefix-list blocksummary out
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
ip prefix-list blocksummary seq 10 deny 172.16.30.0/24
ip prefix-list blocksummary seq 20 deny 10.1.0.0/16
ip prefix-list blocksummary seq 30 permit 0.0.0.0/0 le 32
!
ip prefix-list prefix543 seq 10 permit 10.1.54.0/24
ip prefix-list prefix543 seq 20 permit 10.1.3.0/24
route-map suppress543 permit 10
match ip address prefix-list prefix543
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
end

r6#

```

## R7

```

r7#s
Building configuration...

Current configuration : 1469 bytes
!
version 12.2
no service single-slot-reload-enable
service timestamps debug uptime

```

```
service timestamps log uptime
no service password-encryption
!
hostname r7
!
logging rate-limit console 10 except errors
!
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
key chain ripkey
  key 1
    key-string cisco
!
!
interface Ethernet0
  no ip address
  shutdown
!
interface Ethernet1
  no ip address
  shutdown
!
interface Serial0
  ip address 192.168.1.1 255.255.255.0
  ip rip authentication mode md5
  ip rip authentication key-chain ripkey
  no fair-queue
!
interface Serial1
  no ip address
  shutdown
!
router rip
  version 2
  network 192.168.1.0
  no auto-summary
!
router bgp 67
  no synchronization
  bgp log-neighbor-changes
  neighbor 192.168.1.2 remote-as 67
  no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
!
line con 0
  exec-timeout 0 0
```



```
transport input none
line aux 0
line vty 0 4
 login
!
end

r7#
```

## Cat

```
cat#s
Building configuration...

Current configuration : 2478 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname cat
!
!
ip subnet-zero
ip routing
no ip domain-lookup
!
!
spanning-tree portfast default
spanning-tree extend system-id
!
!
!
interface FastEthernet0/1
 switchport access vlan 10
 no ip address
!
interface FastEthernet0/2
 switchport access vlan 20
 no ip address
!
interface FastEthernet0/3
 switchport access vlan 30
 no ip address
!
interface FastEthernet0/4
 switchport access vlan 54
 no ip address
!
interface FastEthernet0/5
```

```
switchport access vlan 54
no ip address
!
interface FastEthernet0/6
switchport access vlan 60
no ip address
!
```

**[output truncated]**

```
interface Vlan1
no ip address
shutdown
!
interface Vlan30
ip address 10.1.3.2 255.255.255.0
!
interface Vlan60
ip address 10.1.6.2 255.255.255.0
!
router bgp 10
no synchronization
bgp log-neighbor-changes
neighbor 10.1.3.1 remote-as 12345
neighbor 10.1.6.1 remote-as 67
no auto-summary
!
ip classless
ip http server
!
!
line con 0
exec-timeout 0 0
line vty 0 4
login
line vty 5 15
login
!
end

cat#
```

## Explanation

**1. Configure a full-mesh of frame relay links among R1, R2, and R3, so that they share the same subnet. R3 will connect to R1 and R2 using its s0.1 interface. Configure a separate frame relay link between R3 and R4. R3 s0.2 will connect to R4 s0. R1 and R2 will not use any frame relay subinterfaces.**

Manually configuring frame-relay map statements should become a habit, but with all these CCIE preparation habits you need to keep your mind engaged and wary. I have found that I will sometimes merrily type the frame map statement, but forget to put the IP address on the interface.

An interesting thing in this scenario is that you will not need a frame relay map statement on the R3 s0.2 interface, because it is a point-to-point subinterface. You simply set the dlci using the "frame-relay interface-dlci" statement. Nevertheless, you will need a frame relay map statement on R4 at the other end of the link (two, if you want to be able to ping R4 s0 from R4 itself and other locations on the same side of that frame relay link).

**2. Configure all other serial links.**

The other serial links use the default HDLC encapsulation. These are very simple to bring up. Just remember to put a clock rate on the DCE interface. You can see which interface is the DCE interface by typing "sh controllers s 0", for example, with the space between the 'S' and the "0," or you can simply try setting the clock rate by typing "clock rate" under the interface, and seeing if the computer accepts it. The space between "clock" and "rate" is needed when configuring but does not show up in the configuration script.

**3. Configure LAN links, VLANs, and VLAN interfaces.**

We have put R4 e0 and R5 e0 on the same ethernet by simply putting them in the same VLAN. If you want to hook up a PC as a user device, you can simply plug it into any available switchport and then put that port into VLAN 54.

**4. Configure RIP version 2 on R7 and R6. Use RIP MD5 authentication on this link. The link between R1 and R6 will not actively participate in any routing protocol.**

A lot of people have trouble with configuring authentication for RIP. In part the problem is the multilayered configuration of the numbered key chain with the passwords on the key-chain, but the fatal error is generating an extra space at the end of the password. One of the best features of Cisco IOS is interactive help. It is a good habit to use it often to find out what options or new features are available at any point in the configuration. The only caveat is that this healthy habit sometimes results in an extra space at the end

of a password, generated when one types a space and then checks to see what help says. So, when configuring a password, if you type an extra space at the end for any reason, go back and retype the last letter of the password, and then press "Enter."

On the serial interfaces that link R1 and R6, I used an OSPF network statement to put them in OSPF but made them passive interfaces. I went ahead and did the same with several other interfaces that I wanted in OSPF but did not want to generate OSPF hellos. You could also use "redistribute connected" to get routes for these networks into OSPF, but they would all be external routes. With all the mutual route redistribution going on, I thought it would a good idea to redistribute only OSPF internal routes into BGP, so I wanted all the networks that logically look like they belong in the OSPF domain (looking at the IGP diagram) to show up as internal networks in OSPF.

The routing table for R5 shows that all the routes for local networks in the OSPF domain are internal routes. The only external routes are redistributed from BGP:

```
r5#sh ip rou
```

```
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 not set
```

```
       172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA   172.16.30.0/29 [110/128] via 172.16.30.17, 00:30:35, Serial0
O IA   172.16.30.8/29 [110/1849] via 172.16.30.17, 00:30:35, Serial0
C      172.16.30.16/30 is directly connected, Serial0
O IA   172.16.30.20/30 [110/128] via 172.16.30.25, 00:07:46, Serial1
C      172.16.30.24/30 is directly connected, Serial1
       10.0.0.0/24 is subnetted, 5 subnets
O IA   10.1.3.0 [110/138] via 172.16.30.25, 00:30:35, Serial1
       [110/138] via 172.16.30.17, 00:30:36, Serial0
O IA   10.1.2.0 [110/74] via 172.16.30.17, 00:30:36, Serial0
O IA   10.1.1.0 [110/138] via 172.16.30.17, 00:30:36, Serial0
O E1   10.1.6.0 [110/129] via 172.16.30.17, 00:30:36, Serial0
C      10.1.54.0 is directly connected, Ethernet0
O E1   192.168.1.0/24 [110/129] via 172.16.30.17, 00:30:38, Serial0
O IA   192.168.2.0/24 [110/192] via 172.16.30.17, 00:30:38, Serial0
r5#
```

**5. Configure R1, R2 and R3 in OSPF area 0, according to the IGP diagram. The frame relay links among these three routers will use the broadcast network type.**

You set the OSPF network type on the interface. If this were a hub-and-spoke network, you would then need to make sure to “fix” the winner of the DR/BDR election. Since this is a full-mesh network as between R1, R2, and R3, it is not as necessary to “fix” the DR/BDR election.

**6. Configure OSPF area 34 and area 45. Insure that even if the R2-R5 serial link fails, the end users in OSPF area 45 will still be able to reach R1 e0, R2 e0, and R3 e0. R4 and R5 e0 will not try to form an OSPF adjacency.**

Configuring a virtual-link on R4 to connect area 45 to area 0 via area 34 fulfills the first requirement. When configuring a virtual-link, remember that the IP address you use is the router ID of the area 0 router you are connecting to. In a scenario like this, where no loopback interfaces are used, the router number could change in the course of configuration.

To fulfill the second requirement, I put the R4 e0 and R5 e0 interfaces into OSPF using network statements, and made them passive interfaces to keep them from trying to form an adjacency. You could also redistribute into OSPF the “connected” route for each e0 interface into OSPF.

**7. Configure OSPF demand circuit on the ISDN link between R1 and R2. Use dialer profiles.**

I used PPP CHAP authentication even though it was not specifically required, because I have found that ISDN works more reliably with authentication. In some circumstances, I believe that the identification of the other end of the ISDN link that occurs when you use authentication can help with basic layer 2 connectivity. I generally use some form of PPP authentication when I can, unless it is prohibited to me.

The most basic formula to remember is “username remote password cisco.” You can, and probably should, change the password if it is a production network, but the key is to remember that the username is the hostname of the remote router.

Remember to create a dialer list and apply it to the interface. Remember to specify your spids carefully, if you are using a switch-type like basic-ni1 that requires them.

You can run ISDN using the default Cisco HDLC encapsulation, but you seldom see this in practice scenarios, and it is probably not a good idea to abandon the benefits of PPP on an ISDN link. I mention this because you could someday find that your ISDN is up at layer 2, but you are still having problems. Make sure that you did not set it up without setting encapsulation to PPP. Using authentication as a matter of course makes me remember to set the encapsulation to PPP.

The ISDN configuration provided here worked well for me in 12.2, but it may not work as well in every IOS version or on every platform. There have been real improvements to

ISDN since early 12.1. In the past, demand circuit often did not work reliably without some tinkering. The host route (/32) to the other end of the ISDN link sometimes caused problems as it disappeared when the ISDN link dropped. One of the traditional ways to fix the problem was to put "no ip neighbor-route" on the ISDN interface. I did not need to use that in this configuration, but you may find that you need it.

The OSPF demand circuit should come up once or twice when you first boot the routers, but should soon settle down and not come up until a link fails or until you make a configuration change. As you configure R1 and refine your configuration with route-maps and such, you will find that the ISDN link goes up. This is exactly what should happen. If you stop configuring for 45 minutes and "sh dialer" reveals that the link has been down for almost 45 minutes, that is an indicator of success.

### **8. Configure BGP on R1 and R3 in AS 12345, on R6 and R7 in AS 67, and on Cat in AS 10. Configure peering between R1 and R3, R1 and R6, R3 and Cat, and R6 and Cat.**

This is an interesting topology, because running BGP on both R6 and R7 means that RIP is almost superfluous. I say "almost" because RIP is doing one thing without which BGP would not work right in AS 67. Without RIP, R7 would not know the route to R6 s1, and without that route, R7 could not reach the next-hop address of the routes advertised from AS 12345. Of course, you could always change the next-hop address for those routes as R6 advertises them to R7. This discussion may seem tedious and unnecessary, but you need to get used to thinking about exactly what routes are known by what routers, and how they were learned.

### **9. On R6, advertise R6 s0 and R6 e0 in BGP using network statements.**

Using this method of getting routes into BGP is the most natural way to do it in BGP and avoids a lot of potential route feedback problems (see the following section). You could redistribute connected routes into BGP with a route-map to limit the redistribution to specific routes, but that it not as simple as the good old network statement.

### **10. Redistribute OSPF into BGP on R1 and R3. Redistribute BGP into OSPF on R1 and R3. Insure that traffic from AS 12345 to destinations in AS 67 prefers the R1-R6 link.**

The first set of issues to deal with involves the problem of doing mutual route redistribution on two routers without creating route feedback and routing loops. Redistributing only OSPF internal routes into BGP avoids any route feedback from BGP routes going into OSPF on one ASBR and then being redistributed back into BGP on the other ASBR.

As for OSPF routes in AS 12345 going into BGP, being sent over to AS 10 or AS 67 and then being sent back into AS 12345 by an EBGP peer, BGP's built-in loop prevention takes care of that problem, for the most part. BGP routers in AS 12345 will reject any routes with AS 12345 in them. I say "for the most part," because BGP's loop prevention using AS-paths

would not be effective with regard to an aggregate route where no AS-path information was retained.

The redistribution of OSPF into BGP on two routers in AS 12345 could potentially create some problems. Look at R1's BGP table:

**r1#sh ip bgp**

BGP table version is 34, local router ID is 192.168.2.1

Status codes: s suppressed, d damped, h history, \* valid, > best, i - internal

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
* <b>i 10.1.1.0/24</b>	<b>172.16.30.3</b>	<b>74</b>	<b>100</b>	<b>0</b>	<b>?</b>
* >	<b>0.0.0.0</b>	<b>0</b>		<b>32768</b>	<b>?</b>
* i 10.1.2.0/24	172.16.30.2	74	100	0	?
* >	172.16.30.2	74		32768	?
* i 10.1.3.0/24	172.16.30.3	0	100	0	?
* >	172.16.30.3	74		32768	?
* > 10.1.6.0/24	192.168.2.2	0		0	67 i
* i 10.1.54.0/24	172.16.30.22	74	100	0	?
* >	172.16.30.3	138		32768	?
s> 172.16.30.0/29	0.0.0.0	0		32768	?
* > 172.16.30.0/24	0.0.0.0			32768	i
s> 172.16.30.8/29	0.0.0.0	0		32768	?
s> 172.16.30.16/30	172.16.30.2	192		32768	?
s> 172.16.30.20/30	172.16.30.3	128		32768	?
s> 172.16.30.24/30	172.16.30.3	192		32768	?
* > 192.168.1.0	192.168.2.2	0		0	67 i
* i 192.168.2.0	172.16.30.3	128	100	0	?
* >	0.0.0.0	0		32768	?

R3 advertises a route to R1's own e0 interface. This phenomenon is not creating any routing loops within AS 12345 as a practical matter, since R1 prefers its own route to its e0 interface. The IBGP routes will all have a high administrative distance of 200, so they are unlikely to cause problems within the AS.

R1's routing table will not show any of these superfluous BGP-sourced routes, because it only shows the **best** routes.

**r1#sh ip rou**

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 not set

```

    172.16.0.0/16 is variably subnetted, 6 subnets, 3 masks
B       172.16.30.0/24 [200/0] via 0.0.0.0, 00:43:31, Null0
C       172.16.30.0/29 is directly connected, Serial0
C       172.16.30.8/29 is directly connected, Dialer1
O IA    172.16.30.16/30 [110/192] via 172.16.30.2, 00:42:35, Serial0
O IA    172.16.30.20/30 [110/128] via 172.16.30.3, 00:19:37, Serial0
O IA    172.16.30.24/30 [110/192] via 172.16.30.2, 00:42:35, Serial0
        [110/192] via 172.16.30.3, 00:42:35, Serial0
    10.0.0.0/24 is subnetted, 5 subnets
O       10.1.3.0 [110/74] via 172.16.30.3, 00:42:36, Serial0
O       10.1.2.0 [110/74] via 172.16.30.2, 00:42:36, Serial0
C       10.1.1.0 is directly connected, Ethernet0
B       10.1.6.0 [20/0] via 192.168.2.2, 00:42:47
O IA    10.1.54.0 [110/138] via 172.16.30.2, 00:42:37, Serial0
        [110/138] via 172.16.30.3, 00:42:37, Serial0
B       192.168.1.0/24 [20/0] via 192.168.2.2, 00:42:48
C       192.168.2.0/24 is directly connected, Serial1
r1#

```

Just to make sure that things are fine elsewhere in the OSPF domain, let us look at R5's routing table:

r5#sh ip rou

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 not set

```

    172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA    172.16.30.0/29 [110/128] via 172.16.30.17, 00:44:10, Serial0
O IA    172.16.30.8/29 [110/1849] via 172.16.30.17, 00:44:10, Serial0
C       172.16.30.16/30 is directly connected, Serial0
O IA    172.16.30.20/30 [110/128] via 172.16.30.25, 00:21:21, Serial1
C       172.16.30.24/30 is directly connected, Serial1
    10.0.0.0/24 is subnetted, 5 subnets
O IA    10.1.3.0 [110/138] via 172.16.30.25, 00:44:10, Serial1
        [110/138] via 172.16.30.17, 00:44:11, Serial0
O IA    10.1.2.0 [110/74] via 172.16.30.17, 00:44:11, Serial0
O IA    10.1.1.0 [110/138] via 172.16.30.17, 00:44:11, Serial0
O E1    10.1.6.0 [110/129] via 172.16.30.17, 00:44:11, Serial0
C       10.1.54.0 is directly connected, Ethernet0
O E1    192.168.1.0/24 [110/129] via 172.16.30.17, 00:44:11, Serial0
O IA    192.168.2.0/24 [110/192] via 172.16.30.17, 00:44:11, Serial0

```



r5#

We see none of the strange, superfluous routes there, either, so these extra BGP routes are not causing any problem within the OSPF domain. Outside the OSPF domain, there could be problems. It turns out that the AS-path length differences head off serious problems in this scenario, but if R6 were peered directly with both R1 and R2, we can speculate that there could be serious problems. In that hypothetical scenario, R1 and R3 would advertise routes to R1's ethernet interface that would look equally good to R6 (with AD of 20 and an equal AS-path distance). That cannot be a healthy thing.

The task required that traffic in AS 12345 should prefer the R1-R6 link. If all the routers in AS 12345 were running BGP, this would be easy, because they would all see that the AS-path is shorter using the R1-R6 link. However, the routers that run only OSPF will not see the AS-path. All they see is the OSPF cost. Therefore, when I redistributed BGP into OSPF on R3 I gave the redistributed routes a high default-metric of 500. If you do a traceroute from R4 to R7 s0, you will see the traffic go to R3, but then it will go to R1 before crossing over to AS 67. If the R1-R6 link were to fail, then the traffic would take the R3-Cat link.

As for redistributing BGP into OSPF, not too many people would want to do this in a corporate production network. The memory required to deal with a full set of BGP routes for the world is substantial. Even if all your OSPF routers have a lot of memory and fast CPUs, you still would not want to slow down them down with all the redistributed information from BGP.

Regarding redistribution into OSPF in general, an important thing to remember is to specify "subnets" if you want any classless networks to be included. If you forget to include "subnets" the router will warn you that only classful networks will be included, but you could miss the warning. In addition, you should consider what metric-type to use. If you are only redistributing a route into OSPF on one ASBR, then it probably will not matter whether you use the default metric-type of Type 2. If you are redistributing a route into OSPF on more than one ASBR, then you should seriously consider setting the metric-type to Type 1, so the internal OSPF metric to reach the ASBR will be included in the total cost. In addition, you need to consider whether to manually set the OSPF cost of redistributed routes, as we did here to make OSPF routers in AS 12345 choose one way out of the AS over another.

As for redistributing OSPF into BGP, you probably would not do this in production, and if you did, you or your ISP(s) would definitely need to summarize all the BGP routes to a shorter prefix length and suppress all the more specific BGP routes. An ISP in China does not want to know the precise route for the ethernet network that you are on right now, and that ISP definitely does not want the entire BGP table to change when that ethernet network goes down. It makes a lot more sense to advertise summaries of large ranges of addresses using "network" statements and "aggregate-address" statements in BGP.

**11. On R1 and R3, summarize all the 172.16.30.x routes in AS 12345 to 172.16.30.0/24 so that AS 67 will see the 172.16.30.0/24 route, but will not see the more specific routes in**

that range. Insure that no other routers in AS 12345 see any route to 172.16.30.0/24. Do not summarize any of the 10.1.x.x routes on R1 or R3.

To summarize under BGP on R1 and R3, you put "aggregate-address 172.16.30.0 255.255.255.0 summary-only." You then need to make sure that the summary is not advertised back into OSPF, and a route-map referencing a prefix list handles that nicely. I also applied the same route-map to the peering relationship between R1 and R3 so that each would not send the summary to the other. This was purely optional, but it tidies up the BGP tables a bit. Let us look at R1's BGP table:

```
r1#sh ip bgp
```

```
BGP table version is 43, local router ID is 192.168.2.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* i 10.1.1.0/24	172.16.30.3	74	100	0	?
*>	0.0.0.0	0		32768	?
* i 10.1.2.0/24	172.16.30.2	74	100	0	?
*>	172.16.30.2	74		32768	?
*> 10.1.3.0/24	172.16.30.3	74		32768	?
* i	172.16.30.3	0	100	0	?
*> 10.1.6.0/24	192.168.2.2	0		0	67 i
*> 10.1.54.0/24	172.16.30.3	138		32768	?
* i	172.16.30.22	74	100	0	?
<b>s&gt; 172.16.30.0/29</b>	<b>0.0.0.0</b>	<b>0</b>		<b>32768</b>	<b>?</b>
*> <b>172.16.30.0/24</b>	<b>0.0.0.0</b>			<b>32768</b>	<b>i</b>
<b>s&gt; 172.16.30.8/29</b>	<b>0.0.0.0</b>	<b>0</b>		<b>32768</b>	<b>?</b>
<b>s&gt; 172.16.30.16/30</b>	<b>172.16.30.2</b>	<b>192</b>		<b>32768</b>	<b>?</b>
<b>s&gt; 172.16.30.20/30</b>	<b>172.16.30.3</b>	<b>192</b>		<b>32768</b>	<b>?</b>
<b>s&gt; 172.16.30.24/30</b>	<b>172.16.30.3</b>	<b>192</b>		<b>32768</b>	<b>?</b>
*> 192.168.1.0	192.168.2.2	0		0	67 i
* i 192.168.2.0	172.16.30.3	128	100	0	?
*>	0.0.0.0	0		32768	?

The "s" to the left of several BGP routes indicates that the route will be suppressed in advertisements to BGP neighbors.

In order to keep other routers in AS 12345 from seeing this summary route, I used a route-map referencing a prefix-list to keep this route from being redistributed into OSPF. I used the same route-map to prevent the summary route from being advertised to the IBGP neighbor.

Let us look at the BGP table on R6 to see the result:

```
r6#sh ip bgp
```

```
BGP table version is 14, local router ID is 192.168.2.2
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```

Network          Next Hop          Metric LocPrf Weight Path
*> 10.1.0.0/16    0.0.0.0          32768 i
* 10.1.1.0/24    10.1.6.2         0 10 12345 ?
*> 10.1.2.0/24    192.168.2.1     0 12345 ?
* 10.1.3.0/24    10.1.6.2         0 10 12345 ?
*> 10.1.54.0/24   192.168.2.1     74 0 12345 ?
s 10.1.3.0/24    10.1.6.2         0 10 12345 ?
s> 10.1.54.0/24   192.168.2.1     0 12345 ?
*> 10.1.6.0/24    0.0.0.0          0 32768 i
s 10.1.54.0/24   10.1.6.2         0 10 12345 ?
s> 10.1.6.0/24    192.168.2.1     138 0 12345 ?
* 172.16.30.0/24 10.1.6.2 0 10 12345 i
*> 172.16.30.0/24 192.168.2.1 0 12345 i
*> 192.168.1.0    0.0.0.0          0 32768 i
* 192.168.2.0    10.1.6.2         0 10 12345 ?
*> 192.168.2.0    192.168.2.1     0 12345 ?
r6#

```

You can see the summarized 172.16.30.0/24 route, and no more specific routes in that range.

## 12. R6 will advertise a 10.1.0.0/16 BGP aggregate, but will suppress only the specific routes 10.1.54.0/24 and 10.1.3.0/24.

This task gives you a chance to use a very simple prefix-list. It is simple because it involves matching exact routes of a precise mask length. On R6, configure "aggregate-address 10.1.0.0 255.255.0.0 suppress-map suppress543" along with the route-map and prefix-list referenced by the route-map, specifically:

```

route-map suppress543 permit 10
  match ip address prefix-list prefix543

```

and

```

ip prefix-list blocksummary seq 10 deny 172.16.30.0/24
ip prefix-list blocksummary seq 20 deny 10.1.0.0/16
ip prefix-list blocksummary seq 30 permit 0.0.0.0/0 le 32

```

Then reset BGP completely using "clear ip bgp \*." With the suppress-map, you are only specifying which routes you want to suppress, so the route-map will only match on the prefix-list specifying the routes you want suppressed.

When using route-maps, it is easy for the head to begin to spin from the logic. Note that in this configuration, we do not have a "catch-all" route-map with a permit statement and no match statement. Nor do we have a catch-all "permit 0.0.0.0/0 le 32 prefix-list. Watch your routing tables and BGP tables carefully, and if you have a problem, check the logic of your route-maps and prefix-lists and access-lists. Let us take another look at the BGP table on R6, this time with emphasis on different routes:

## r6#sh ip bgp

BGP table version is 25, local router ID is 192.168.2.2

Status codes: s suppressed, d damped, h history, \* valid, &gt; best, i - internal

Origin codes: i - IGP, e - EGP, ? - incomplete

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	<b>10.1.0.0/16</b>	<b>0.0.0.0</b>			<b>32768</b>	<b>i</b>
*	10.1.1.0/24	10.1.6.2			0	10 12345 ?
*>		192.168.2.1	0		0	12345 ?
*	10.1.2.0/24	10.1.6.2			0	10 12345 ?
*>		192.168.2.1	74		0	12345 ?
<b>s</b>	<b>10.1.3.0/24</b>	<b>10.1.6.2</b>			<b>0</b>	<b>10 12345 ?</b>
<b>s&gt;</b>		<b>192.168.2.1</b>			<b>0</b>	<b>12345 ?</b>
*>	10.1.6.0/24	0.0.0.0	0		32768	i
<b>s&gt;</b>	<b>10.1.54.0/24</b>	<b>192.168.2.1</b>	<b>138</b>		<b>0</b>	<b>12345 ?</b>
<b>s</b>		<b>10.1.6.2</b>			<b>0</b>	<b>10 12345 ?</b>
*	172.16.30.0/24	10.1.6.2			0	10 12345 i
*>		192.168.2.1			0	12345 i
*>	192.168.1.0	0.0.0.0	0		32768	i
*	192.168.2.0	10.1.6.2			0	10 12345 ?
*>		192.168.2.1	0		0	12345 ?

Note the 10.1.0.0/16 route and the various suppressed routes within that range. The BGP table of R7 will not show any of the suppressed routes, of course:

## r7#sh ip bgp

BGP table version is 19, local router ID is 192.168.1.1

Status codes: s suppressed, d damped, h history, \* valid, &gt; best, i - internal

Origin codes: i - IGP, e - EGP, ? - incomplete

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>i	10.1.0.0/16	192.168.1.2		100	0	i
*>i	10.1.1.0/24	192.168.2.1	0	100	0	12345 ?
*>i	10.1.2.0/24	192.168.2.1	74	100	0	12345 ?
*>i	10.1.6.0/24	192.168.1.2	0	100	0	i
*>i	172.16.30.0/24	192.168.2.1		100	0	12345 i
*>i	192.168.1.0	192.168.1.2	0	100	0	i
*>i	192.168.2.0	192.168.2.1	0	100	0	12345 ?

In the real world, we would probably not want to see in our BGP table even this amount of detail about routes in other ASes (here we see /24 routes). An ISP might see some detailed routes with mask-lengths like /24 from smaller customers, but will summarize (aggregate) them before advertising BGP routes to other ISPs. Of course, you cannot simply summarize routes beyond the range that you own or manage, or you could end up blackholing traffic destined for someone else's network included in the range that you summarized.

**13. Take steps to keep the 10.1.0.0/16 and 172.16.30.0/24 summary routes from being advertised from R6 to R1 or Cat.**

R1 and R3 will reject any advertisement with 12345 in the AS-path, so 172.16.30.0/24 will not be a problem for AS 12345, but we need to keep it and 10.1.0.0/16 from getting to Cat and 10.1.0.0/16 from being advertised into AS 12345. Without some form of filtering, the 10.1.0.0/16 summary route generated on R6 will be advertised into AS 10 and AS 12345. Having this route floating around in AS 12345 would be bad form, to say the least.

I used a relatively simple prefix-list covering both summary routes to accomplish the task. Under "router bgp 67," we will type:

```
neighbor 10.1.6.2 prefix-list blocksummary out
neighbor 192.168.2.1 prefix-list blocksummary out
```

In addition, globally we will configure:

```
ip prefix-list blocksummary seq 10 deny 172.16.30.0/24
ip prefix-list blocksummary seq 20 deny 10.1.0.0/16
ip prefix-list blocksummary seq 30 permit 0.0.0.0/0 le 32
```

Note that since we are blocking specific routes and permitting all other routes, we need a catch-all prefix-list statement at the end to permit all routes that were not already denied. We also have other ways to block certain routes and permit others, such as a distribute-list or route-map applied at the end of a neighbor statement.

**14. Configure Hot Standby Routing Protocol on R4 e0 and R5 e0 so that R5 will normally serve the hosts on the 10.1.54.0/24 ethernet. If R5 s0 fails, R4 will serve the hosts on the shared ethernet. All end users on the shared ethernet will be using 10.1.54.1/24 as their default gateway.**

Hot Standby Routing Protocol (HSRP) sounds like a big deal, but you can see that it is just about the easiest thing in the world to configure. That does not mean that you cannot think of a difficult HSRP task, or mess up an HSRP configuration, but its syntax and operation are straightforward. What sometimes delays me a few seconds is that I forget to put the group number, and the router complains because it thinks I am trying to configure more than one standby group on an interface.

Here is the output from "show standby" on R5 before I bring down R5 s0:

```
r5#sh standby
Ethernet0 - Group 1
  Local state is Active, priority 105, may preempt
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:00.966
  Hot standby IP address is 10.1.54.1 configured
  Active router is local
  Standby router is 10.1.54.4 expires in 00:00:09
```

```

Standby virtual mac address is 0000.0c07.ac01
2 state changes, last state change 04:58:31
Tracking interface states for 1 interface, 1 up:
Up Serial0 Priority decrement: 10
r5#

```

The trick with tracking is that you need to make sure that the value that you lose when a tracked interface fails (10 by default) will bring the priority value of the primary router low enough that the standby router will take over. If you set the starting priorities to 105 and 95 respectively, for example, you could have a problem, because when the tracked serial interface on the active router fails you will have two routers with a priority of 95. This is not a problem if a router fails entirely. In that event, the priority will decrement plenty, but it is a problem for the tracking situation, where the default behavior is to decrement the priority by 10, and to decrement only once.

Let us bring down R5 s0 and watch the output from "**debug standby**":

```

04:35:10: SB1: Et0 Hello out 10.1.54.5 Active pri 105 ip 10.1.54.1
04:35:11: SB1: Et0 Hello in 10.1.54.4 Standby pri 100 ip 10.1.54.1shu
r5(config-if)#
04:35:13: SB1: Et0 Hello out 10.1.54.5 Active pri 105 ip 10.1.54.1
04:35:14: SB1: Et0 Hello in 10.1.54.4 Standby pri 100 ip 10.1.54.1
04:35:14: SB: Et0 Sbstate adv in, Passive, active 0 passive 1, from
10.1.54.4
04:35:14: SB1: Et0 Tracked interface Se0 down, 0/1 now up
04:35:14: SB1: Et0 Priority was 105 now 95, configured as 105
04:35:14: SB1: Et0 Hello out 10.1.54.5 Active pri 95 ip 10.1.54.1
04:35:14: %OSPF-5-ADJCHG: Process 64, Nbr 172.16.30.17 on Serial0 from
FULL to D
OWN, Neighbor Down: Interface down or detached
04:35:14: SB1: Et0 Coup in 10.1.54.4 Standby pri 100 ip 10.1.54.1
04:35:14: SB1: Et0 Active: j/Coup rcvd from higher pri router
(100/10.1.54.4)
04:35:14: SB1: Et0 Active router is 10.1.54.4, was local
04:35:14: SB: Et0 Remove active hash 10.1.54.5 (vIP 10.1.54.1)
04:35:14: SB: Et0 Add active hash 10.1.54.4 (vIP 10.1.54.1)
04:35:14: SB1: Et0 Active -> Speak
04:35:14: %STANDBY-6-STATECHANGE: Standby: 1: Ethernet0 state Active
-> Speak
k
04:35:14: SB: Et0 Sbstate adv start
04:35:14: SB: Et0 Sbstate adv out, Passive, active 0 passive 1
04:35:14: SB1: Et0 Hello out 10.1.54.5 Speak pri 95 ip 10.1.54.1
04:35:14: SB: Et0 Sbstate adv in, Active, active 1 passive 0, from
10.1.54.4
04:35:14: SB: Et0 Remove passive hash 10.1.54.4 (frc 0)
04:35:14: SB1: Et0 Hello in 10.1.54.4 Active pri 100 ip 10.1.54.1
04:35:16: %LINK-5-CHANGED: Interface Serial0, changed state to
administratively
down
04:35:17: SB1: Et0 Hello in 10.1.54.4 Active pri 100 ip 10.1.54.1

```

```
04:35:17: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0,
changed state
to down
04:35:17: SB1: Et0 Hello out 10.1.54.5 Speak pri 95 ip 10.1.54.1
```

When R5 s0 went down, R5 decremented its own priority value by 10, and then advertised that lower value to R4 over the shared ethernet. R4 then did a “coup” and took over from R5.

**15. We need to severely limit the bandwidth of traffic coming into R3 e0 from one server. The server's MAC address is 2121.2020.2121. Limit all incoming traffic from that server to a maximum of 160,000 kbps.**

Rate-limiting is a particularly useful feature. It makes it easy for an ISP or an internet data center operator to limit the total bandwidth of incoming traffic from one customer router or server. A company might also want to limit the network bandwidth used by certain of its own servers or applications. You can limit the bandwidth of a particular type of traffic, such as FTP. Imagine that there is a very small software company that hosts its own on-line store as well as its own dedicated FTP server for customer support downloads. The company does not want customer FTP support downloads to hog network bandwidth on certain links so that sessions involving on-line purchases fail. It can limit FTP traffic coming from the FTP server into the router, or it can limit all traffic coming into the router from the FTP server's MAC address.

When tinkering with things like rate limiting in a production network, you will need to take into account political realities and the limits of your authority with regard to network matters. The drastic rate limiting based on MAC address that we are doing in this scenario could ruffle some feathers. This also shows the importance of documenting everything in your network. If someone decides to move the NIC card from the FTP server to the on-line store server, or rebuilds the FTP server as the online store server, they will definitely want to know about the rate limiting. I would speculate that that would be one advantage of using rate-limiting based on the type of traffic.

When you set the bits per second, you should also manually set the burst bytes and excess burst bytes according to Cisco guidelines rather than simply entering a bits per second value and letting the router set the burst and excess burst values.

Cisco recommends that you divide the bits per second value by 8 to get bytes and multiply that value by 1.5 to get your normal bytes burst value. Your excess burst value will be twice your normal burst value. Thus, in this scenario:

160,000 kbps / 8 = 20,000 bytes.

20,000 bytes X 1.5=30,000 bytes burst.

30,000 bytes X 2=60,000 bytes excess burst.

You can also set traffic that conforms to the limit to be sent with a certain IP precedence value, and traffic that exceeds to be sent with a lower IP precedence value. You should explore this feature in depth, as the task in this scenario is very simple.

You can show the results of your rate-limit configuration:

```
r3#sh int e0 rate-limit
Ethernet0
  Input
    matches: access-group rate-limit 120
      params: 160000 bps, 30000 limit, 60000 extended limit
      conformed 1 packets, 60 bytes; action: transmit
      exceeded 0 packets, 0 bytes; action: drop
      last packet: 2841948ms ago, current burst: 0 bytes
      last cleared 00:48:49 ago, conformed 0 bps, exceeded 0 bps
r3#
```

Note that the special rate-limit access-lists, such as the MAC-address access-list that we used here, use the same access-list numbers as standard IP access lists and extended IP access-lists. Make sure that you specify "rate-limit" with the access list.

**16. Configure custom queueing on R5 s0 so that each of four queues will get roughly one-quarter of the bandwidth. Queue 1 will be for Telnet. Queue 2 will be for FTP. Queue 3 will be for Web traffic. Queue 4 will be for everything else. The size of each queue should be 1500 bytes.**

Custom queueing does not have any fast lane or "HOV" lane. Rather, each queue gets a turn. The default size of a queue is 1500 bytes, which is pretty easy to remember. You can view your queueing configuration with "show queueing":

```
r5#sh queueing
```

Current fair queue configuration:

Interface	Discard threshold	Dynamic queues	Reserved queues	Link queues	Priority queues
Serial 1	64	256	0	8	1

Current DLCI priority queue configuration:

Current priority queue configuration:

Current custom queue configuration:

List	Queue	Args
1	4	default
1	1	protocol ip tcp port telnet
1	2	protocol ip tcp port ftp
1	3	protocol ip tcp port www

Current random-detect configuration:

```
r5#
```



Debugging is as easy as typing "debug custom" and telneting somewhere across the link on which queueing is configured:

```
r5#telnet 10.1.2.1
```

```
Trying 10.1.2.1 ... Open
```

```
Password required, but none set
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 48/1) Q # was 4 now 1
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 44/1) Q # was 1 now 1
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 53/1) Q # was 1 now 1
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 44/1) Q # was 1 now 1
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 47/1) Q # was 1 now 1
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 47/1) Q # was 1 now 1
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 53/1) Q # was 1 now 1
```

```
05:12:42: CQ: Serial0 output (Pk size/Q: 44/1) Q # was 1 now 1
```

```
[Connection to 10.1.2.1 closed by foreign host]
```

```
r5#
```

```
05:12:44: CQ: Serial0 output (Pk size/Q: 44/1) Q # was 1 now 1
```

```
05:12:44: CQ: Serial0 output (Pk size/Q: 44/1) Q # was 1 now 1
```

```
05:12:49: CQ: Serial0 output (Pk size/Q: 279/4) Q # was 1 now 4
```

You see all the small Telnet packets that do not even use up the 1500 bytes queue-length, and then finally a larger packet using the default queue. We skipped the second and third queues because there was no FTP or HTTP traffic.

# **Appendix A**

General Information  
and Opinion about the CCIE

## On the value of the CCIE

Anyone who is reading this book probably already knows that the CCIE is a highly regarded certification, but let us talk about it a bit, because the economic value of the CCIE has become a matter of debate since the dramatic downturn in the telecom industry. I recommend that you do your own thorough research and soul-searching about the value of pursuing the CCIE for your career or personal development before you set out to achieve the CCIE. You should satisfy yourself that it is the right thing for you to do before you set out on the journey, so that when the tough times and doubts and discouraging words arrive, you will be ready to carry on until you reach your goal.

The CCIE has been a highly respected certification for several years now, but it has now become a “hot” certification, in the sense that many people in the IT field would like to have it, and many are currently pursuing it or planning to pursue it in the near future. To be honest, there used to be very few people who could realistically consider pursuing it, but now many networkers have progressed in their career or studies so that the CCIE appears to be a logical next step. It is a very big step.

A recent survey done by [www.certcities.com](http://www.certcities.com) found that the CCIE is the “hottest” certification of 2003. Normally, I like to shy away from things that are “hot” because it is often best to avoid being where the crowds are. Numerous stocks were “hot” in 1999 that you would not want to have purchased in that year and held until 2003.

The CCIE is a valuable certification to pursue because it is so difficult to achieve, and because one learns so much from preparing for the Lab Exam. In my opinion, the CCIE will remain valuable as long as Cisco maintains the rigor of the Lab Exam, even if the number of CCIEs continues to grow at the current rate of over 2,000 per year, and even if this rate were to increase slightly.

The lab exam is so difficult and wide-ranging in its coverage that it requires that even very experienced network engineers engage in specialized preparation in some area of theory or Cisco IOS before they can pass the CCIE Lab. Many of these very experienced network engineers would not need the CCIE to obtain or retain a good job, but it gives them the opportunity and incentive to broaden or deepen their knowledge and bring additional recognition to themselves or their organization.

I suspect that the more common situation is that a network engineer or network technician pursues the CCIE as a graduate student pursues the Ph.D.—that is, as a life-transforming experience. These are not gurus getting Cisco's imprimatur, but people setting out to improve themselves. Many Ph.D. candidates never finish their dissertations, and are referred to colloquially as “A.B.D.” (“all but dissertation”). Similarly, many who set out to get the CCIE do not reach the goal. This is a harsh reality for those pursuing the CCIE, but the fact that many senior network engineers and IT managers know several well-regarded senior network engineers who failed the CCIE more than once helps keep the certification respected.

When you hear people on bulletin boards speculating that the CCIE will soon become ubiquitous like the MCSE, you can rest assured that they are wrong. The number of CCIEs is growing at a higher rate than it did a few years ago, primarily because there are now

more testing seats available. Even with the increased number of lab seats to ease what used to be a horrendous scheduling backlog, there is an absolute limit on how many candidates can even take the Lab exam in a year. Furthermore, very few candidates pass the Lab on the first attempt. Of those who eventually pass the Lab, many more take two or three or four or five attempts to do it. Legend has it that the average number of tries is around three.

As of July 31, 2003, there were 10,144 people with current CCIE certification in the world. This does not include people who once had the CCIE but let it lapse. Someone on the [www.groupstudy.com](http://www.groupstudy.com) ccielab list sent a message that they had passed on July 30, 2003, with the number 12,035, so there are a couple of thousand numbers to be accounted for. CCIE numbers begin at 1,024, with the first number going to the program itself. Therefore, there have been almost 1,000 folks who achieved the CCIE at some point, and for whatever reason let it lapse. These reasons could conceivably include death, promotion to a position where the CCIE would be superfluous, retirement, loss of CCIE status through violation of Cisco's rules on non-disclosure and the like, or leaving the industry entirely. Some of these one thousand folks are probably busy network engineers who just need to get around to passing a written exam to recertify within a year after their certification lapses. CCIE numbers are not re-used, and someone who gets a second or third CCIE has only one CCIE number.

The corps of CCIEs is not nearly as small and elite a group as it was four years ago during the tech boom, when there were only a few thousand and the demand for them was so great. Nevertheless, I speculate that as the world market for networking equipment and services improves, and more and more people gain internet access, the CCIE will be a very good certification to have. How many people have you met who had internet access and then disconnected themselves from the internet, never to return? The internet and computer networking are not going away. The knowledge one gains and the skills one hones in the process (including skill in learning about new technologies) will remain valuable.

The CCIE certification is a requirement or desired qualification for many advertised full-time positions, and for many public and private contracts. It is also required as part of the Cisco Silver and Gold Partner programs for resellers. I am not suggesting that the CCIE, standing alone, will get you a great job these days, but it is very valuable when combined with one or more other valuable qualities or skills such as experience, education, writing skills, or personal skills.

## CCIE Tracks

The vast majority of the people who have achieved the CCIE have achieved it in the Routing and Switching track. Currently, there are the following tracks:

- 1) Routing and Switching
- 2) Security
- 3) Service Provider (formerly Communications and Services)

#### 4) Voice

Each track requires a written exam, which you need to pass before you can schedule a date to take the Lab Exam. You have 18 months from when you pass the written qualifier exam to take the Lab Exam for the first time. After that first attempt, you have to make an attempt within the next year to keep from having to retake the written exam, and so on until three years have passed since you passed the written exam. Once three years have passed, you have to take the written exam again to be able to take the Lab Exam.

You sign on to the on-line scheduling tool using your testing identification number, the date you passed the written qualifier exam, and your score. You should keep this information in a safe place, because you may need to sign into the CCIE scheduling tool several times over many months and even after you have passed the Lab Exam.

Make sure you verify the time periods with the CCIE website and make sure you understand the situation. You would not believe how many people get the different time periods mixed up. As with anything else in the Cisco world, read and digest the information Cisco provides. I would go so far as to say that you should not rely on what I have said here or on any advice you get in an on-line discussion group. You should read the **most current** policies yourself, and if you still have any questions, ask Cisco itself. If you have questions about the CCIE program, contact [ccie@cisco.com](mailto:ccie@cisco.com).

For each track, check the pertinent blueprint. This tells you what can be covered on the written qualifier exam. You will not get quite so specific guidance about what is on the Lab Exam, so remember not to use the written exam blueprint as your blueprint for the lab exam.

In the case of Service Provider, there are different kinds of written exams you can take, with emphasis on different specialized areas: 1) optical, 2) DSL, 3) dial, 4) cable, 5) WAN switching, and 6) IP telephony. All Service Provider candidates who pass a written exam go on to take the same form of Lab Exam.

The Voice track is brand new. It makes sense that Cisco would offer a CCIE in this specialty since it is so difficult to master and there is currently considerable demand for specialists in this area. If you are going to spend the time necessary to master this difficult area and pass the numerous computer-based IP telephony tests, you might well want to go a further and have a CCIE to show for all your efforts.

The practical reality is that a certification such as the CCIE is not solely intended to measure and identify people who already have certain skills. It is also intended to *encourage* people to *develop* the skills being tested. When IT managers are deciding what equipment to buy, the ability to support that equipment is as important a factor as any other. It is the availability of so many people trained on Cisco equipment, and so many resources on Cisco's website, in the printed documentation, and in the many Cisco-focused reference and training books, that sets Cisco apart from so much of the competition.

## CCIE Lab Testing Centers

The CCIE Lab testing centers are in the following cities, but not all exams are offered in each location:

Bangalore, India  
Beijing, P.R.C.  
Brussels, Belgium  
Johannesburg, S.A.  
Hong Kong, P.R.C.  
RTP, North Carolina, USA  
San Jose, California, USA  
Sao Paulo, Brazil  
Sydney, Australia  
Tokyo, Japan

On December 1, 2003, the South Africa testing center will close, and a new testing center in Dubai, UAE, will open.

I believe that the Routing and Switching Lab Exam is offered at all Lab exam testing centers.

The Security Lab Exam is now offered in Research Triangle Park (RTP), Beijing, San Jose, Brussels, and Sydney.

The Service Provider (formerly Communications and Services) Lab Exam is offered in Sao Paolo, Hong Kong, RTP, Brussels, and Sydney.

You should check Cisco's website often, as administrative details change regularly, and it is best to keep up with the newest testing realities. If you participate actively in an on-line forum you will probably hear about changes that Cisco has posted.

A colleague suggested that I discuss in this book the various hotels near CCIE Lab testing centers, but I think this is one of those areas where an on-line forum is so valuable. When you schedule your CCIE Lab Exam after passing the written qualifier test, Cisco will send you an email that will include a list of hotels near the testing center. You may want to seek advice from friends about these and other hotels. For the most part, there is a "well-trodden path" in the sense that most people tend to stay at the same few hotels. For example, at RTP, anecdotal evidence suggests that many CCIE candidates favor the Wingate Inn near the Raleigh-Durham Airport because it has a shuttle to Cisco and high-speed internet access. I did not stay there, because I drove to RTP both times, and therefore did not need a shuttle. The on-the-ground realities can change, so you should read the information Cisco sends you and then consult your friends on-line for further details.

## Recertification

Once you pass the CCIE, it is not over, for you need to recertify every two years. Recertifying currently consists of passing any CCIE Written Qualifier Exam. Thus, you could

pass the CCIE Routing and Switching Lab, but recertify on the CCIE Security Written Qualifier Exam.

There is no incentive to put off recertifying, because if you pass the Lab on November 1, 2003, and take the recertification test on November 10, 2003, your CCIE will be valid until November 1, 2007. If you wait until October 31, 2005 to recertify, your CCIE will be valid until November 1, 2007. There is something to be said for recertifying right after you pass the Lab exam, when much of the material is fresher than it will be in a year or more.

Cisco used to have a continuing education requirement that could be met only by attending at least one Networkers convention every two years. That requirement no longer exists. I thought it was burdensome when I first heard about it, but some CCIEs miss it, because it made their employers send them to Networkers.

# Appendix B

## Annotated Bibliography



## Annotated Bibliography

CCIE lab preparation is like graduate school in requiring a great deal of reading. Some people like to read a few "great books" cover-to-cover, but I prefer to read liberally from many books, seeking answers to specific questions raised by practice scenarios or exploring topics that interest me at the moment. This approach may not be right for everyone, but it worked for me. I did not sit down and attempt to read one book at a time straight through, except in the case of Doyle's *Routing TCP/IP (Volume I)*. I had a large library of books as well as the entire printed Cisco IOS 12.1 documentation set, and read exactly what interested me most at any given time. I absorbed what I read because I had a keen interest in what I read.

I provide here a list of books I used in the course of preparing for the CCIE lab, along with a few newer ones that I did not get to use, with some brief comments. Use this list as a starting point for your own research. There are some good books that I never got around to purchasing and new books useful for CCIE Lab preparation are published every month or so.

Caputo, Robert, *Cisco Packetized Voice & Data Integration*, New York: McGraw Hill; 1999. This was one of my favorite CCIE preparation books. It contains detailed explanations of Quality of Service features such as frame relay traffic shaping, including explanations of how different values are calculated.

Caslow, Bruce; Pavlichenko, Valeriy, *Cisco Certification: Bridges, Routers and Switches for CCIEs (Second Edition)*. Upper Saddle River, New Jersey: Prentice Hall PTR; 2001. This book is aimed directly at CCIE preparation, and teaches an analytical, issue-spotting approach to the CCIE, similar to the approach one uses in law school exams. It covers many technologies, teaching many of the crucial lessons to learn or issues to watch out for. The first edition of this book was my favorite CCIE preparation book, and the Second Edition, which I obtained late in my CCIE preparation, is even better.

*CCIE Fundamentals: Network Design and Case Studies (Second Edition)*. Indianapolis, Indiana: Cisco Press; 2000. Contains in-depth coverage of many technologies, and contains both theory and IOS syntax.

Clark, Kennedy; Hamilton, Kevin, *Cisco LAN Switching*, Indianapolis, Indiana: Cisco Press; 1999. This was a wonderful book on switching for preparing for the CCNP Switching exam as well as for the CCIE Lab exam, but its emphasis on the Catalyst 5000 means that it is due for a new edition.

Doyle, Jeff, *Routing TCP/IP (Volume I)*. Indianapolis, Indiana: Cisco Press; 1998. This is a crucial book on routing. This book is on just about anybody's "must-read" list for the CCIE Lab.

Doyle, Jeff, *Routing TCP/IP (Volume II)*. Indianapolis, Indiana: Cisco Press; 2001. Doyle Covers BGP and IP Multicast in some depth, and includes some discussion of miscellaneous router management matters such as syslog, security, SNMP, and the like. He includes a whole chapter on EGP, the predecessor to BGP, and another on IPv6.

Halabi, Sam; McPherson, Danny, *Internet Routing Architectures (Second Edition)*, Indianapolis, Indiana: Cisco Press; 2000. This is a must-have book on BGP.

Hutnik, Stephen; Saterlee, Michael, *All-in-One CCIE Lab Study Guide*, New York: McGraw-Hill, 2000. While it is not really an all-in-one guide, it makes an excellent first book for CCIE Lab preparation.

Hutnik, Stephen; Saterlee, Michael, *CCIE Lab Practice Kit*, New York: McGraw-Hill, 2001. This is a wonderful book with lab scenarios focusing on different technologies, solution configuration scripts, and task-by-task explanation.

*Internetworking Technologies Handbook (Second Edition)*, Indianapolis, Indiana: Cisco Press; 1998. This book contains theoretical discussion of many technologies, with no coverage of IOS syntax.

*Internetworking Troubleshooting Handbook*, Indianapolis, Indiana: Cisco Press; 1999. I used this book a lot in the initial stages of CCIE preparation, as well as for CCNP preparation. It is an indispensable reference for troubleshooting everything from physical layer and data link layer problems to routing protocols.

Lewis, Chris, *Cisco Switched Internetworks*. New York: McGraw-Hill; 1999. In addition to covering basic switching concepts and VLANs, Lewis also covers ATM in considerable detail, including a useful explanation of the ATM hierarchical addressing system, and also covers voice over IP using Cisco MC3810. I consulted this book often about ATM.

Lewis, Chris, *Cisco TCP/IP Routing Professional Reference (Third Edition)*, New York: McGraw-Hill; 2000. This book is worth having as a reference.

McQuery, Steve; McGrew, Kelly; Foy, Stephen, *Cisco Voice over Frame Relay, ATM, and IP*. Indianapolis, Indiana: Cisco Press; 2001. Designed to prepare one for an older version of the CVoice exam, this book is worth looking at in preparing for the CCIE (Routing and Switching).

Parkhurst, William, *Cisco BGP-4 Command and Configuration Handbook*, Indianapolis, Indiana: Cisco Press; 2001. This was one of my favorite CCIE preparation books, because Dr. Parkhurst focuses heavily on command syntax. Doctor Parkhurst leads us through most of the important BGP commands and includes configuration examples. It is designed to be a reference book, but I enjoyed it so much I read it for fun. This book and the Halabi book make a great combination, since Parkhurst is strong on syntax and his style is very readable, while Halabi is strong on theory and his style is more formal.

Parkhurst, William, *Cisco OSPF Command and Configuration Handbook*, Indianapolis, Indiana: Cisco Press; 2001. I do not yet own this book, but if it anything like his new Cisco Press BGP book, it has to be good.

Parkhurst, William, *Cisco Router OSPF: Design and Implementation Guide*, New York: McGraw-Hill, 1998. This is Dr. Parkhurst's older book on OSPF. It was good, but he has come out with a later book on OSPF from Cisco Press, which is mentioned above.

Parkhurst, William, *Multicast Routing & Switching*, New York: McGraw-Hill, 1999. I used this book often, as I like Dr. Parkhurst's style and emphasis on hands-on configuration and syntax.

Perlman, Radia, *Interconnections: Bridges and Routers (Second Edition)*, Reading, Massachusetts: Addison-Wesley: 1999. I used the first edition of this book, but apparently, the Second Edition is even better. It is worth reading what she has to say about subjects like Spanning Tree Protocol and IS-IS, since they would not exist without her. Many people enjoy her witty commentary.

Raza, Khalid; Turner, Mark, *CCIE Professional Development: Large-Scale IP Network Solutions*, Indianapolis, Indiana: Cisco Press; 2000. This book offers basic-to-intermediate coverage of various IP routing protocols and matters such as SNMP and congestion-avoidance algorithms. It balances theory and syntax roughly evenly.

Sackett, George; Sackett, Nancy, *Internetworking SNA with Cisco Solutions*, Indianapolis, Indiana: Cisco Press; 1999. A good book for real-world SNA networking, but may not be necessary for the current CCIE, from which Token Ring has been removed. It contains coverage of DLSW+.

Slattery, Terry; Burton, Bill, *Advanced IP Routing in Cisco Networks (Second Edition)*, New York: McGraw-Hill, 2000. It is just what it purports to be, and is worth having in your library.

Solie, Karl, *CCIE Practical Studies (Volume I)*, Indianapolis, Indiana: Cisco Press; 2002. This is an absolute must-have book for CCIE Lab preparation. While it falls short of perfection due to errors, it is an unbeatable value because of its scope of coverage and number of learning labs.

Thomas, II, Thomas, *OSPF Network Design Solutions (Second Edition)*, Indianapolis, Indiana: Cisco Press; 2003. I used the earlier version of this book. I like this author's writing style, but the first edition did not focus strictly on OSPF, but rather covered numerous subjects that someone running an OSPF network might need to know about.

Tripod, Mark, *Cisco Router Configuration & Troubleshooting*. Indianapolis, Indiana: New Riders; 1999. This was a good general reference to have around for the "baby steps" phase of CCIE preparation.

Vegesna, Srinivas, *IP Quality of Service*, Indianapolis, Indiana: Cisco Press; 2001. An indispensable guide to numerous quality of service features. As good as this book is, I still recommend that you read broadly, and do not forget the IOS 12.2. configuration guide volume on QOS. QOS material is so difficult that it helps to see the same thing explained by different people.

Williamson, Beau. *Developing IP Multicast Networks*, Indianapolis, Indiana: Cisco Press; 2000. This is widely considered an excellent book on IP Multicast, and I wish I had purchased it before I began work on this book.