

CCIE Warm-Up (Volume II)

Written by Thomas P. Larus
CCIE 10,014



Distributed by
IPexpert, Inc. (www.ipexpert.com)

For technical support
please visit www.CertificationTalk.com

Copyright© 2004 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®, CCNAT™, 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.net. 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.

Table of Contents

Preface		iv
Chapter 1	Scenario One	1
Chapter 2	Scenario Two	49
Chapter 3	Scenario Three	85
Chapter 4	Scenario Four	117
Chapter 5	Scenario Five	149

Preface

These five new scenarios teach many lessons that were not included in *CCIE Warm-Up: Advice and Learning Labs*. I have kept the labs at the same intermediate level of difficulty. Listed below are some points covered in these new scenarios. Some of these points were covered in Volume One, but many are new to this volume, or are configured using different syntax or in a different context:

- 1) Multi-area OSPF with several virtual links
- 2) Mutual route redistribution between OSPF and RIP
- 3) Summarizing external routes on an OSPF Autonomous System Boundary Router (ASBR)
- 4) Route summarization in RIP version 2
- 5) OSPF demand circuit
- 6) IP Multicast- PIM sparse mode and PIM dense mode
- 7) Configuring access lists so that they do not break applications or routing protocols.
- 8) Frame relay traffic shaping
- 9) Rate-limiting based on type of traffic
- 10) Class-based weighted-fair queuing based on type of traffic and based on DSCP value
- 11) Using IS-IS passive interface to get networks into IS-IS
- 12) Dialer watch with ISDN, legacy DDR and EIGRP
- 13) BGP Multi-exit-discriminator
- 14) Preventing an Autonomous System from being used as a transit AS
- 15) Basic DLSW+ using TCP encapsulation
- 16) TCP header compression configured as part of static frame relay map.
- 17) IS-IS interface mismatch
- 18) Route redistribution from IS-IS to another IGP, and the need to manually redistribute "connected" networks participating in IS-IS into the target IGP.
- 19) Hub-and-spoke frame relay network in EIGRP— the need to disable split-horizon for EIGRP on the hub's frame relay interface

- 20) Preventing route feedback in EIGRP-to-IS-IS mutual route redistribution on two routers by using route-maps referencing access-lists
- 21) BGP confederation with a route-reflector within a sub-AS
- 22) BGP confederation without a route reflector within a sub-AS
- 23) BGP weight attribute
- 24) AS-path prepend
- 25) ISDN backup using floating static routes
- 26) EIGRP authentication
- 27) RIP authentication
- 28) BGP authentication
- 29) IP accounting
- 30) Dialer-watch with ISDN and IS-IS
- 31) IS-IS interface authentication, area authentication, and domain authentication
- 32) BGP local-preference

These scenarios focus heavily on challenging and pervasive Layer Two and Layer Three issues that CCIE candidates have struggled with for years, and which are not likely to go away any time soon. They also cover QOS, IP multicast, DLSW+ and certain IP services, but these tasks are intended to build familiarity and confidence, rather than to test the knowledge of a candidate who is nearly ready to pass the Lab Exam.

There are two things I have done differently with this book from *CCIE Warm-Up: Advice and Learning Labs*. First, I have added routing tables for each router. Second, I do not specify DTE or DCE in the network diagrams. It does not matter much which side of a link is DCE or DTE in a lab setting, as long as you set a clock rate on whichever one is DCE. This is a minor point, but I want to remove one way in which a student's rack or a rental rack might diverge from the topology used in the book.

I am grateful to Technical Reviewers Joseph Rothstein and Richard Dumoulin for working through these five scenarios and suggesting improvements.

To discuss these scenarios, post a query on www.certificationtalk.com or e-mail me directly at tlarus@ipexpert.net.

Chapter One

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 of ethernet interfaces 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. In all the scenarios, only one physical interface on R7 will be needed -- a serial interface. This permits the reader to use a terminal server as R7. Some CCIE candidates like to connect their terminal server to the rest of their network, so that they can Telnet to it, rather than be limited to console access to it. Since only four serial interfaces on the frame switch router will be used for frame switching, any unit with five or more serial interfaces can be used 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.

R1 through R4 connect to each other over the frame relay cloud using their serial0 interfaces. The Visio diagrams indicate how the R1 through R4 serial1 interfaces connect, as well as the s0 and S1 interfaces of R5 and R6 and the s0 interfaces on R7. Different links are used in different scenarios, but the cabling will never change. To facilitate your initial wiring, here is a table showing the connections between the serial interfaces that are not connected to the frame relay cloud. Note that R3 s1 is not used in any scenario.

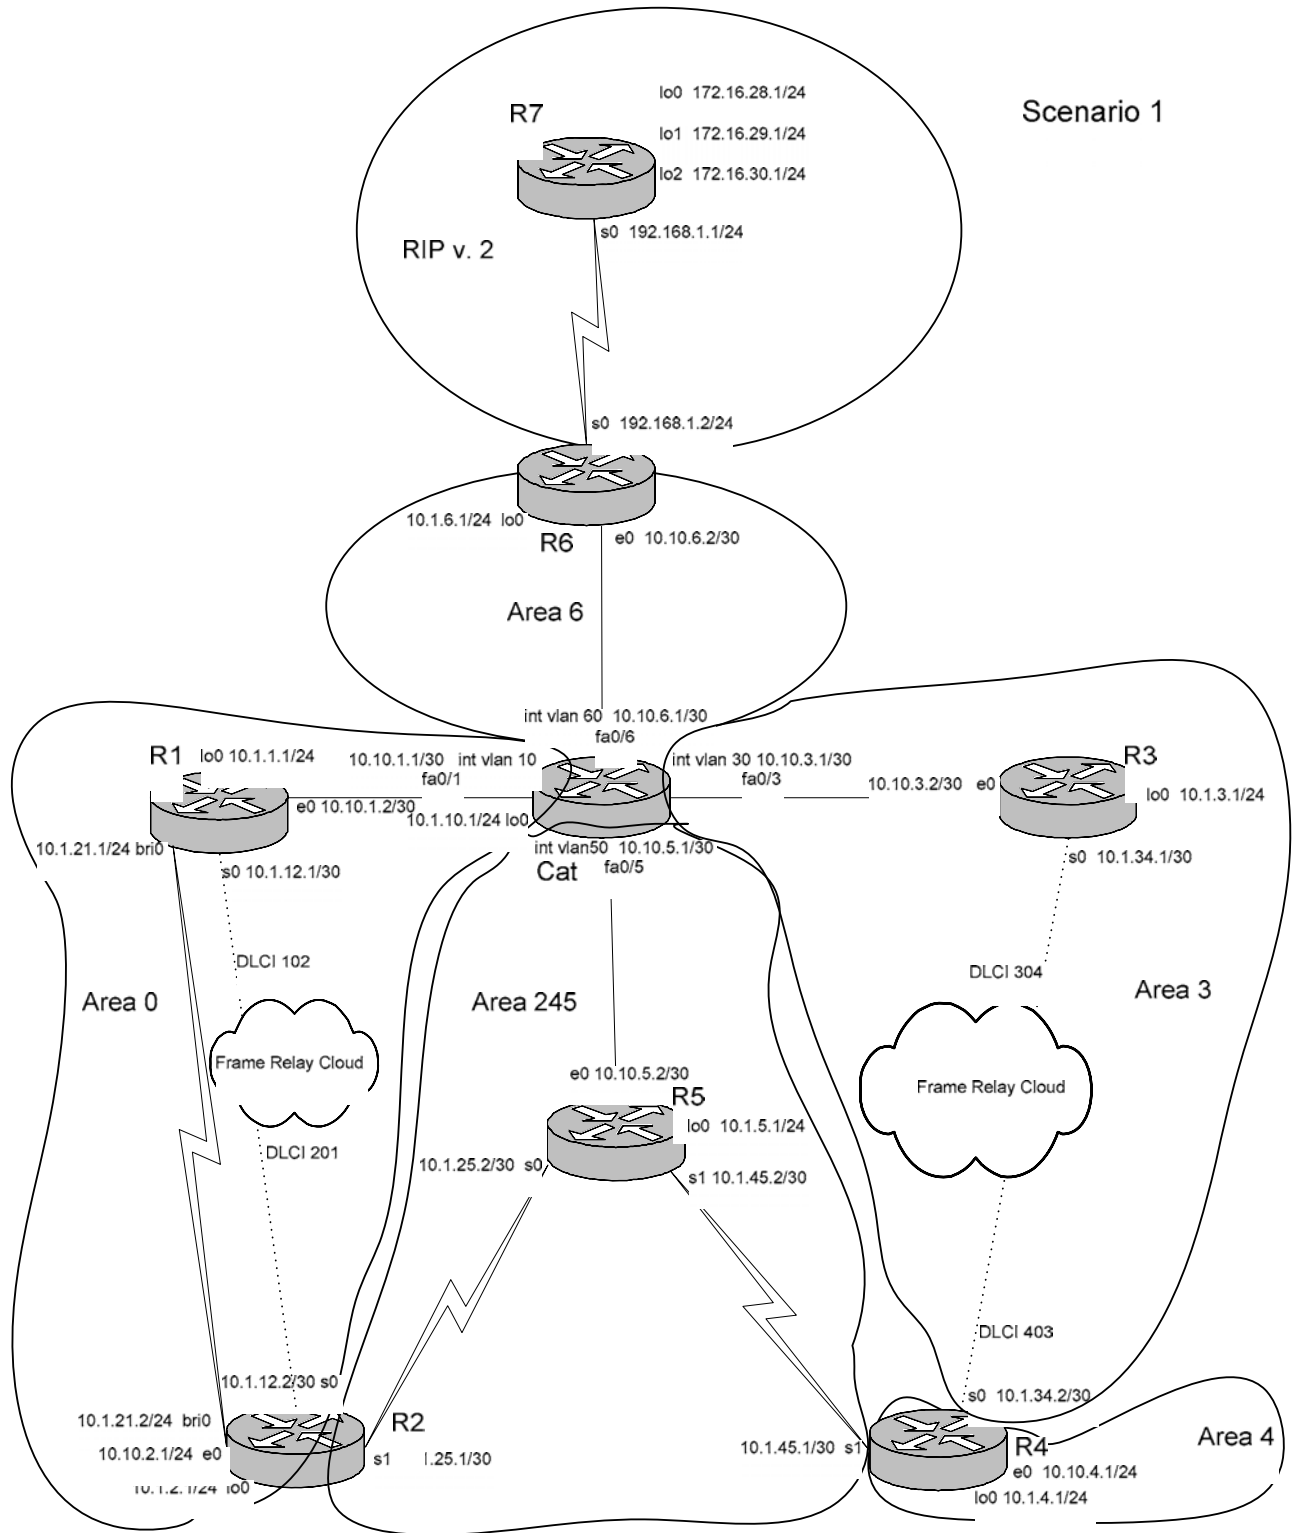
R1 s1	R6 s1
R2 s1	R5 s0
R4 s1	R3 s0
R6 s0	R7 s0

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. I like to stick to all lower-case in configuration scripts, where possible, to eliminate potential for error having to do with case-sensitivity. This practice departs from real-world practice, where route-map names are often capitalized for readability.

I like to use the Cat 3550 s "spanning-tree portfast default" feature whenever I can. This leaves spanning-tree in place, but dramatically shortens the spanning-tree delays for all the access ports (not the uplink ports).

Always disable DNS lookup on a router in a lab setting, using "no ip domain-lookup." This will seem ridiculously obvious to most of you, but I still forget to do it sometimes. It is also helpful to set the exec-timeout to 0 under "line con 0."

Scenario 1



Scenario One focuses on OSPF in a network with many redundant paths. Anyone who designed an OSPF production network that looks like this would be fired. In the real world, you might well put all these OSPF routers in one area, area 0. Also, as you look at the diagram, you see area 0 over on the left. It would seem to make more sense to put area 0 where area 245 is, so that all areas would connect directly to the backbone. CCIE practice scenarios are generally not designed as models of real-world "best practices," but rather to teach lessons.

Try not to be distracted too much by oddities like ethernet subnets with /30 masks. Some CCIE candidates with a lot of experience on production networks are uncomfortable with the odd and unrealistic situations in CCIE practice scenarios. It reminds me of how, when I was studying law and practicing law, I had trouble learning legal theories or rules that I found repugnant.

If a scenario seems awkward or ugly or ridiculous, focus on what you do not like about it. I was taught that what you do not like about a scenario may well be where you will find serious issues.

Tasks

1. Configure all Ethernet links and loopback addresses as shown in the diagram and according to the chart below. Configure appropriate VLANs on Cat, and enable the spanning tree portfast feature on all Cat switch access ports.

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

2. Insure that the switch ports specified in the chart above cannot become trunk ports.

3. Configure all frame-relay links as shown in the diagram. Do not use any frame relay inverse-arp or frame relay subinterfaces.

4. Configure all other serial links as shown in the diagram using PPP encapsulation. Set the clock rate on each serial link to 125 kbps. Configure link authentication on each serial PPP link using the password "ppp." R2 should authenticate with R5 as "r2s1."

5. The OSPF areas are shown in the diagram. Note that area 6 is a Not-So-Stubby-Area. Cat and each router in the OSPF domain will use the loopback0 IP address as the OSPF Router ID. Configure plaintext authentication in area 0, using password "clear," and MD5 authentication in all other areas, using password "crypt."
6. Configure OSPF demand circuit on the ISDN link between R1 and R2. Use PPP CHAP authentication with "cisco" as the password.
7. Configure virtual-links in the OSPF domain, and feel free to create some redundant virtual links.
8. R6 s0 and R7 s0, lo0, lo1, and lo2 will be part of the RIP domain, but only R7 s0 and R6 s0 will send RIP updates.
9. On R6, redistribute RIP into OSPF and OSPF into RIP. On R6, summarize the three routes for the R7 loopbacks to one network. Using a route-map but without referencing an access list, prevent this summary route from being redistributed into RIP. On R6, summarize all the 10.x.x.x routes to 10.0.0.0/8 before they reach R7. Insure that all interfaces in the OSPF domain (except ISDN interfaces, which could be idle) are reachable from R7, and that all interfaces on R7 are reachable from the OSPF domain.
10. Use R7 lo0 as the source of multicast packets. Enable multicast using a shared tree so that R4 lo0 and R2 lo0 can receive the multicast packets from R7 lo0, which will be addressed to the multicast group 239.0.0.1. Test your configuration by sending extended pings sourced from R7 lo0 to the multicast group.
11. Configure frame relay traffic shaping on R3 s0 using a MINCIR of 32 kbps, a CIR of 64, and the ability to burst to 128 kbps. The rate should drop no lower than the MINCIR in the event of congestion indicated by BECN messages.
12. Configure R7, R4 and R3 so that they can be accessed by means of Telnet. Configure an access list on R6 s0 (inbound and outbound) permitting all ICMP traffic, RIP routing protocol traffic, and Telnet traffic between R7 s loopbacks and R4 s loopback. Test your configuration to make sure that R7 can Telnet to R4, and vice-versa, but R7 cannot Telnet to R3.
13. Configure rate limiting inbound and outbound on R6 e0 so that Telnet traffic will not exceed 64 kbps.
14. Configure Low Latency Queueing on R6 s0 so that Telnet traffic going onto the R6-R7 link will enjoy a strict priority queue with guaranteed bandwidth of 16 kbps. Allocate 20 kbps to traffic that is tagged with a DSCP value of 8 (which corresponds to the old IP precedence of 1), and use random early detection for that traffic.

Configuration Scripts

R1

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

Current configuration : 1749 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 ospf network point-to-point
!
interface Ethernet0
 ip address 10.10.1.2 255.255.255.252
 ip pim sparse-mode
 ip ospf authentication
 ip ospf authentication-key clear
!
interface Serial0
 ip address 10.1.12.1 255.255.255.252
 ip pim sparse-mode
 encapsulation frame-relay
 ip ospf authentication
 ip ospf authentication-key clear
 ip ospf network point-to-point
 frame-relay map ip 10.1.12.1 102
 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
  ip ospf authentication
  ip ospf authentication-key clear
  ip ospf demand-circuit
  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
  no peer neighbor-route
  ppp authentication chap
!
router ospf 64
  log-adjacency-changes
  area 0 authentication
  network 0.0.0.0 255.255.255.255 area 0
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 192.168.1.1
!
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 : 2270 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
username r5 password 0 ppp
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
 ip ospf network point-to-point
 ip igmp join-group 239.0.0.1
!
interface Ethernet0
 ip address 10.10.2.1 255.255.255.0
!
interface Serial0
 ip address 10.1.12.2 255.255.255.252
 ip pim sparse-mode
 encapsulation frame-relay
 ip ospf authentication
 ip ospf authentication-key clear
 ip ospf network point-to-point
 frame-relay map ip 10.1.12.1 201 broadcast
 frame-relay map ip 10.1.12.2 201
 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
 encapsulation ppp
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 crypt
 clockrate 125000
 ppp authentication chap
 ppp chap hostname r2s1
 ppp chap password 7 12091507
!
interface BRI0
 ip address 10.1.21.2 255.255.255.0
```

```

encapsulation ppp
ip ospf authentication
ip ospf authentication-key clear
dialer map ip 10.1.21.1 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
area 245 authentication message-digest
area 245 virtual-link 10.1.4.1 authentication authentication-key clear
area 245 virtual-link 10.1.10.1 authentication authentication-key
clear
network 10.1.2.0 0.0.0.255 area 0
network 10.1.12.0 0.0.0.3 area 0
network 10.1.21.0 0.0.0.255 area 0
network 10.1.25.0 0.0.0.3 area 245
network 10.10.2.1 0.0.0.0 area 0
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 192.168.1.1
!
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

r2#

```

R3

```

r3#sh runn
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 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 ospf network point-to-point
!
interface Ethernet0
 ip address 10.10.3.2 255.255.255.252
 ip pim sparse-mode
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 crypt
!
interface Serial0
 ip address 10.1.34.1 255.255.255.252
 ip pim sparse-mode
 encapsulation frame-relay
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 crypt
 ip ospf network point-to-point
 no fair-queue
 frame-relay class r3s0frts
 frame-relay traffic-shaping
 frame-relay map ip 10.1.34.1 304
 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 3 authentication message-digest
 network 0.0.0.0 255.255.255.255 area 3
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 192.168.1.1

```

```

!
!
map-class frame-relay r3s0frts
  frame-relay cir 64000
  frame-relay bc 8000
  frame-relay be 8000
  frame-relay mincir 32000
  frame-relay adaptive-shaping becn
!
alias configure rr router rip
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  password cisco
  login
!
end

r3#

```

R4

```

r4#sh runn
Building configuration...

Current configuration : 1887 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
no logging console
!
username r5 password 0 ppp
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 239.0.0.1
!
interface Ethernet0
 ip address 10.10.4.1 255.255.255.0
!
interface Serial0
 ip address 10.1.34.2 255.255.255.252
 ip pim sparse-mode
 encapsulation frame-relay
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 crypt
 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
 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
 encapsulation ppp
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 crypt
 clockrate 125000
 ppp authentication chap
!
interface TokenRing0
 no ip address
 shutdown
!
router ospf 64
 router-id 10.1.4.1
 log-adjacency-changes
 area 3 authentication message-digest
 area 3 virtual-link 10.1.10.1 authentication authentication-key clear
 area 245 authentication message-digest
 area 245 virtual-link 10.1.2.1 authentication authentication-key clear
 area 245 virtual-link 10.1.10.1 authentication authentication-key
 clear
 network 10.1.4.0 0.0.0.255 area 4
 network 10.1.34.0 0.0.0.3 area 3
 network 10.1.45.0 0.0.0.3 area 245
 network 10.10.4.1 0.0.0.0 area 4
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 192.168.1.1
!
!
!
line con 0
 exec-timeout 0 0
 transport input none
line aux 0

```

```

line vty 0 4
  password cisco
  login
!
end

```

```
r4#
```

R5

```

r5#sh runn
Building configuration...

Current configuration : 1351 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
!
username r2s1 password 0 ppp
username r4 password 0 ppp
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 ospf network point-to-point
!
interface Ethernet0
  ip address 10.10.5.2 255.255.255.252
  ip pim sparse-mode
  ip ospf authentication message-digest
  ip ospf message-digest-key 1 md5 crypt
!
interface Serial0
  ip address 10.1.25.2 255.255.255.252
  ip pim sparse-mode
  encapsulation ppp
  ip ospf authentication message-digest
  ip ospf message-digest-key 1 md5 crypt
  no fair-queue
  ppp authentication chap

```

```

!
interface Serial1
 ip address 10.1.45.2 255.255.255.252
 ip pim sparse-mode
 encapsulation ppp
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 crypt
 ppp authentication chap
!
router ospf 64
 log-adjacency-changes
 area 245 authentication message-digest
 network 0.0.0.0 255.255.255.255 area 245
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 192.168.1.1
!
!
!
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 : 2767 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 r7 password 0 ppp
ip subnet-zero
no ip finger
no ip domain-lookup
!

```

```

ip multicast-routing
no ip dhcp-client network-discovery
!
class-map match-all telnet
  match access-group 106
class-map match-all cs1
  match ip dscp 8
!
!
policy-map s0policy
  class telnet
    priority 16
  class cs1
    bandwidth 20
    random-detect
!
!
!
!
!
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
  ip pim sparse-mode
  rate-limit input access-group 107 64000 12000 24000 conform-action
transmit exc
  eed-action drop
  rate-limit output access-group 107 64000 12000 24000 conform-action
transmit ex
  ceed-action drop
  ip ospf authentication message-digest
  ip ospf message-digest-key 1 md5 crypt
!
interface Ethernet1
  no ip address
  shutdown
!
interface Serial0
  ip address 192.168.1.2 255.255.255.0
  ip access-group 106 in
  ip access-group 106 out
  ip pim sparse-mode
  encapsulation ppp
  ip summary-address rip 10.0.0.0 255.0.0.0
  service-policy output s0policy
  no peer neighbor-route
  clockrate 125000
  ppp authentication chap
!
interface Serial1
  no ip address
  shutdown
!
router ospf 64

```

```

log-adjacency-changes
area 6 authentication message-digest
area 6 nssa
summary-address 172.16.28.0 255.255.252.0
redistribute rip metric-type 1 subnets
network 10.1.6.0 0.0.0.255 area 6
network 10.10.6.0 0.0.0.3 area 6
!
router rip
version 2
redistribute ospf 64 metric 3 route-map blocksummary
passive-interface default
no passive-interface Serial0
network 192.168.1.0
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 192.168.1.1
!
access-list 106 permit tcp host 10.1.4.1 172.16.28.1 0.0.3.0 eq telnet
access-list 106 permit tcp host 10.1.4.1 eq telnet 172.16.28.1 0.0.3.0
access-list 106 permit tcp 172.16.28.1 0.0.3.0 host 10.1.4.1 eq telnet
access-list 106 permit tcp 172.16.28.1 0.0.3.0 eq telnet host 10.1.4.1
access-list 106 permit udp any eq rip any
access-list 106 permit udp any any eq rip
access-list 106 permit icmp any any
access-list 106 permit pim any any
access-list 107 permit tcp host 10.1.4.1 172.16.28.1 0.0.3.0 eq telnet
access-list 107 permit tcp host 10.1.4.1 eq telnet 172.16.28.1 0.0.3.0
access-list 107 permit tcp 172.16.28.1 0.0.3.0 host 10.1.4.1 eq telnet
access-list 107 permit tcp 172.16.28.1 0.0.3.0 eq telnet host 10.1.4.1
route-map blocksummary deny 10
match interface Null0
!
route-map blocksummary permit 20
!
!
!
line con 0
transport input none
line aux 0
line vty 0 4
login
!
end

r6#

```

R7

```
r7#sh runn
```

Building configuration...

```
Current configuration : 1190 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
!
username r6 password 0 ppp
ip subnet-zero
no ip finger
no ip domain-lookup
!
ip multicast-routing
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
 ip address 172.16.28.1 255.255.255.0
 ip pim sparse-mode
!
interface Loopback1
 ip address 172.16.29.1 255.255.255.0
!
interface Loopback2
 ip address 172.16.30.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.1 255.255.255.0
 ip pim sparse-mode
 encapsulation ppp
 no peer neighbor-route
 no fair-queue
 ppp authentication chap
!
interface Serial1
 no ip address
 shutdown
!
router rip
 version 2
```

Copyright © 2004, Thomas P. Larus


```

passive-interface default
no passive-interface Serial0
network 172.16.0.0
network 192.168.1.0
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
ip pim rp-address 192.168.1.1
!
!
!
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 : 3195 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 Loopback0
  ip address 10.1.10.1 255.255.255.0
  ip ospf network point-to-point

```

```
!  
interface FastEthernet0/1  
  switchport access vlan 10  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/2  
  switchport access vlan 20  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/3  
  switchport access vlan 30  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/4  
  switchport access vlan 40  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/5  
  switchport access vlan 50  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/6  
  switchport access vlan 60  
  switchport mode access  
  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  
  no ip address  
!  
interface FastEthernet0/12  
  no ip address  
!  
interface FastEthernet0/13  
  no ip address  
!  
interface FastEthernet0/14  
  no ip address  
!  
interface FastEthernet0/15  
  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  
  shutdown  
!  
interface GigabitEthernet0/1  
  no ip address  
!  
interface GigabitEthernet0/2  
  no ip address  
!  
interface Vlan1  
  no ip address  
  shutdown  
!  
interface Vlan10  
  ip address 10.10.1.1 255.255.255.252  
  ip pim sparse-mode  
  ip ospf authentication  
  ip ospf authentication-key clear  
!  
interface Vlan30  
  ip address 10.10.3.1 255.255.255.252  
  ip pim sparse-mode  
  ip ospf authentication message-digest  
  ip ospf message-digest-key 1 md5 crypt  
!  
interface Vlan50  
  ip address 10.10.5.1 255.255.255.252  
  ip pim sparse-mode  
  ip ospf authentication message-digest  
  ip ospf message-digest-key 1 md5 crypt  
!
```

```
interface Vlan60
  ip address 10.10.6.1 255.255.255.252
  ip pim sparse-mode
  ip ospf authentication message-digest
  ip ospf message-digest-key 1 md5 crypt
!
router ospf 64
  log-adjacency-changes
  area 0 authentication
  area 3 authentication message-digest
  area 3 virtual-link 10.1.4.1 authentication authentication-key clear
  area 6 authentication message-digest
  area 6 nssa
  area 245 authentication message-digest
  area 245 virtual-link 10.1.2.1 authentication authentication-key clear
  area 245 virtual-link 10.1.4.1 authentication authentication-key clear
  network 10.1.10.0 0.0.0.255 area 0
  network 10.10.1.0 0.0.0.3 area 0
  network 10.10.3.0 0.0.0.3 area 3
  network 10.10.5.0 0.0.0.3 area 245
  network 10.10.6.0 0.0.0.3 area 6
!
ip classless
ip http server
ip pim rp-address 192.168.1.1
!
!
!
!
line con 0
  exec-timeout 0 0
line vty 0 4
  login
line vty 5 15
  login
!
end

cat#
```

Routing Tables

R1

```

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

      172.16.0.0/22 is subnetted, 1 subnets
O E1   172.16.28.0 [110/32] via 10.10.1.1, 00:13:30, Ethernet0
      10.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
O      10.1.10.0/24 [110/11] via 10.10.1.1, 00:13:30, Ethernet0
C      10.10.1.0/30 is directly connected, Ethernet0
O      10.10.2.0/24 [110/74] via 10.1.12.2, 00:13:30, Serial0
O IA   10.10.3.0/30 [110/11] via 10.10.1.1, 00:13:30, Ethernet0
O IA   10.10.4.0/24 [110/85] via 10.10.1.1, 00:13:30, Ethernet0
O IA   10.10.5.0/30 [110/11] via 10.10.1.1, 00:13:31, Ethernet0
O IA   10.10.6.0/30 [110/11] via 10.10.1.1, 00:13:31, Ethernet0
C      10.1.12.0/30 is directly connected, Serial0
O IA   10.1.3.0/24 [110/12] via 10.10.1.1, 00:13:31, Ethernet0
O      10.1.2.0/24 [110/65] via 10.1.12.2, 00:13:31, Serial0
C      10.1.1.0/24 is directly connected, Loopback0
O IA   10.1.6.0/24 [110/12] via 10.10.1.1, 00:13:31, Ethernet0
O IA   10.1.5.0/24 [110/12] via 10.10.1.1, 00:13:31, Ethernet0
O IA   10.1.4.0/24 [110/76] via 10.10.1.1, 00:13:31, Ethernet0
O IA   10.1.25.0/30 [110/75] via 10.10.1.1, 00:13:31, Ethernet0
C      10.1.21.0/24 is directly connected, BRI0
O IA   10.1.45.0/30 [110/75] via 10.10.1.1, 00:13:31, Ethernet0
O IA   10.1.34.0/30 [110/75] via 10.10.1.1, 00:13:31, Ethernet0
O E1  192.168.1.0/24 [110/32] via 10.10.1.1, 00:13:31, Ethernet0
r1#

```

R2

```

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

```

    172.16.0.0/22 is subnetted, 1 subnets
O E1   172.16.28.0 [110/96] via 10.1.25.2, 00:13:58, Serial1
        [110/96] via 10.1.12.1, 00:13:58, Serial0
    10.0.0.0/8 is variably subnetted, 19 subnets, 3 masks
O      10.1.10.0/24 [110/75] via 10.1.25.2, 00:13:58, Serial1
        [110/75] via 10.1.12.1, 00:13:58, Serial0
O      10.10.1.0/30 [110/74] via 10.1.12.1, 00:13:58, Serial0
C      10.10.2.0/24 is directly connected, Ethernet0
O IA   10.10.3.0/30 [110/75] via 10.1.25.2, 00:13:59, Serial1
        [110/75] via 10.1.12.1, 00:13:59, Serial0
O IA   10.10.4.0/24 [110/138] via 10.1.25.2, 00:13:59, Serial1
O      10.10.5.0/30 [110/74] via 10.1.25.2, 01:28:33, Serial1
O IA   10.10.6.0/30 [110/75] via 10.1.25.2, 00:13:59, Serial1
        [110/75] via 10.1.12.1, 00:13:59, Serial0
C      10.1.12.0/30 is directly connected, Serial0
O IA   10.1.3.0/24 [110/76] via 10.1.25.2, 00:13:59, Serial1
        [110/76] via 10.1.12.1, 00:14:00, Serial0
C      10.1.2.0/24 is directly connected, Loopback0
O      10.1.1.0/24 [110/65] via 10.1.12.1, 00:14:00, Serial0
O IA   10.1.6.0/24 [110/76] via 10.1.25.2, 00:14:00, Serial1
        [110/76] via 10.1.12.1, 00:14:00, Serial0
O      10.1.5.0/24 [110/65] via 10.1.25.2, 01:28:33, Serial1
O IA   10.1.4.0/24 [110/129] via 10.1.25.2, 00:14:00, Serial1
C      10.1.25.2/32 is directly connected, Serial1
C      10.1.25.0/30 is directly connected, Serial1
C      10.1.21.0/24 is directly connected, BRI0
O      10.1.45.0/30 [110/128] via 10.1.25.2, 01:28:34, Serial1
O IA   10.1.34.0/30 [110/139] via 10.1.25.2, 00:14:01, Serial1
        [110/139] via 10.1.12.1, 00:14:01, Serial0
O E1  192.168.1.0/24 [110/96] via 10.1.25.2, 00:14:01, Serial1
        [110/96] via 10.1.12.1, 00:14:01, Serial0
r2#

```

R3

```

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

```

172.16.0.0/22 is subnetted, 1 subnets
O E1    172.16.28.0 [110/32] via 10.10.3.1, 01:29:16, Ethernet0
10.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
O IA    10.10.1.0/30 [110/11] via 10.10.3.1, 00:31:31, Ethernet0
O IA    10.1.10.0/24 [110/11] via 10.10.3.1, 01:29:14, Ethernet0
O IA    10.10.2.0/24 [110/85] via 10.10.3.1, 01:28:54, Ethernet0
C       10.10.3.0/30 is directly connected, Ethernet0
O IA    10.10.4.0/24 [110/74] via 10.1.34.2, 01:29:09, Serial0
O IA    10.10.5.0/30 [110/11] via 10.10.3.1, 01:29:17, Ethernet0
O IA    10.10.6.0/30 [110/11] via 10.10.3.1, 01:29:15, Ethernet0
O IA    10.1.12.0/30 [110/75] via 10.10.3.1, 00:14:52, Ethernet0
C       10.1.3.0/24 is directly connected, Loopback0
O IA    10.1.2.0/24 [110/76] via 10.10.3.1, 01:28:55, Ethernet0
O IA    10.1.1.0/24 [110/12] via 10.10.3.1, 00:31:38, Ethernet0
O IA    10.1.6.0/24 [110/12] via 10.10.3.1, 01:29:15, Ethernet0
O IA    10.1.5.0/24 [110/12] via 10.10.3.1, 01:29:17, Ethernet0
O IA    10.1.4.0/24 [110/65] via 10.1.34.2, 01:29:11, Serial0
O IA    10.1.25.0/30 [110/75] via 10.10.3.1, 01:29:15, Ethernet0
O IA    10.1.21.0/24 [110/1573] via 10.10.3.1, 00:31:38, Ethernet0
O IA    10.1.45.0/30 [110/75] via 10.10.3.1, 01:29:17, Ethernet0
C       10.1.34.0/30 is directly connected, Serial0
O E1    192.168.1.0/24 [110/32] via 10.10.3.1, 01:29:17, Ethernet0
r3#

```

R4

```

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 not set

```

172.16.0.0/22 is subnetted, 1 subnets
O E1    172.16.28.0 [110/96] via 10.1.34.1, 00:15:09, Serial0
        [110/96] via 10.1.45.2, 00:15:09, Serial1
10.0.0.0/8 is variably subnetted, 19 subnets, 3 masks
O       10.10.1.0/30 [110/75] via 10.1.45.2, 00:15:09, Serial1
        [110/75] via 10.1.34.1, 00:15:09, Serial0
O       10.1.10.0/24 [110/75] via 10.1.34.1, 00:15:09, Serial0
        [110/75] via 10.1.45.2, 00:15:09, Serial1
O       10.10.2.0/24 [110/138] via 10.1.45.2, 00:15:10, Serial1
O       10.10.3.0/30 [110/74] via 10.1.34.1, 01:29:43, Serial0
C       10.10.4.0/24 is directly connected, Ethernet0
O       10.10.5.0/30 [110/74] via 10.1.45.2, 01:29:43, Serial1
O IA    10.10.6.0/30 [110/75] via 10.1.34.1, 00:15:10, Serial0
        [110/75] via 10.1.45.2, 00:15:10, Serial1
O       10.1.12.0/30 [110/139] via 10.1.45.2, 00:15:10, Serial1

```

```

[110/139] via 10.1.34.1, 00:15:10, Serial0
O 10.1.3.0/24 [110/65] via 10.1.34.1, 01:29:43, Serial0
O 10.1.2.0/24 [110/129] via 10.1.45.2, 00:15:10, Serial1
O 10.1.1.0/24 [110/76] via 10.1.45.2, 00:15:10, Serial1
[110/76] via 10.1.34.1, 00:15:10, Serial0
O IA 10.1.6.0/24 [110/76] via 10.1.34.1, 00:15:10, Serial0
[110/76] via 10.1.45.2, 00:15:10, Serial1
O 10.1.5.0/24 [110/65] via 10.1.45.2, 01:29:43, Serial1
C 10.1.4.0/24 is directly connected, Loopback0
O 10.1.25.0/30 [110/128] via 10.1.45.2, 01:29:43, Serial1
O 10.1.21.0/24 [110/1637] via 10.1.45.2, 00:15:11, Serial1
[110/1637] via 10.1.34.1, 00:15:11, Serial0
C 10.1.45.2/32 is directly connected, Serial1
C 10.1.45.0/30 is directly connected, Serial1
C 10.1.34.0/30 is directly connected, Serial0
O E1 192.168.1.0/24 [110/96] via 10.1.34.1, 00:15:11, Serial0
[110/96] via 10.1.45.2, 00:15:11, Serial1
r4#

```

R5

```

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

```

172.16.0.0/22 is subnetted, 1 subnets
O E1 172.16.28.0 [110/32] via 10.10.5.1, 01:30:33, Ethernet0
10.0.0.0/8 is variably subnetted, 20 subnets, 3 masks
O IA 10.10.1.0/30 [110/11] via 10.10.5.1, 00:16:07, Ethernet0
O IA 10.1.10.0/24 [110/11] via 10.10.5.1, 01:30:33, Ethernet0
O IA 10.10.2.0/24 [110/74] via 10.1.25.1, 01:30:20, Serial0
O IA 10.10.3.0/30 [110/11] via 10.10.5.1, 01:30:33, Ethernet0
O IA 10.10.4.0/24 [110/74] via 10.1.45.1, 01:30:20, Serial1
C 10.10.5.0/30 is directly connected, Ethernet0
O IA 10.10.6.0/30 [110/11] via 10.10.5.1, 01:30:34, Ethernet0
O IA 10.1.12.0/30 [110/75] via 10.10.5.1, 00:16:18, Ethernet0
O IA 10.1.3.0/24 [110/12] via 10.10.5.1, 01:30:34, Ethernet0
O IA 10.1.2.0/24 [110/65] via 10.1.25.1, 01:30:21, Serial0
O IA 10.1.1.0/24 [110/12] via 10.10.5.1, 00:16:08, Ethernet0
O IA 10.1.6.0/24 [110/12] via 10.10.5.1, 01:30:34, Ethernet0
C 10.1.5.0/24 is directly connected, Loopback0
O IA 10.1.4.0/24 [110/65] via 10.1.45.1, 01:30:21, Serial1
C 10.1.25.0/30 is directly connected, Serial0
C 10.1.25.1/32 is directly connected, Serial0
O IA 10.1.21.0/24 [110/1573] via 10.10.5.1, 00:33:04, Ethernet0

```



```

C      10.1.45.0/30 is directly connected, Serial1
C      10.1.45.1/32 is directly connected, Serial1
O IA   10.1.34.0/30 [110/75] via 10.10.5.1, 01:30:34, Ethernet0
O E1  192.168.1.0/24 [110/32] via 10.10.5.1, 01:30:34, Ethernet0
r5#

```

R6

```

r6#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, 4 subnets, 2 masks
O      172.16.28.0/22 is a summary, 01:35:42, Null0
R      172.16.28.0/24 [120/1] via 192.168.1.1, 00:00:12, Serial0
R      172.16.29.0/24 [120/1] via 192.168.1.1, 00:00:12, Serial0
R      172.16.30.0/24 [120/1] via 192.168.1.1, 00:00:12, Serial0
       10.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
O IA   10.10.1.0/30 [110/11] via 10.10.6.1, 00:37:37, Ethernet0
O IA   10.1.10.0/24 [110/11] via 10.10.6.1, 01:35:34, Ethernet0
O IA   10.10.2.0/24 [110/85] via 10.10.6.1, 01:35:26, Ethernet0
O IA   10.10.3.0/30 [110/11] via 10.10.6.1, 01:35:31, Ethernet0
O IA   10.10.4.0/24 [110/85] via 10.10.6.1, 01:35:16, Ethernet0
O IA   10.10.5.0/30 [110/11] via 10.10.6.1, 01:35:31, Ethernet0
C      10.10.6.0/30 is directly connected, Ethernet0
O IA   10.1.12.0/30 [110/75] via 10.10.6.1, 00:20:58, Ethernet0
O IA   10.1.3.0/24 [110/12] via 10.10.6.1, 01:35:35, Ethernet0
O IA   10.1.2.0/24 [110/76] via 10.10.6.1, 01:35:27, Ethernet0
O IA   10.1.1.0/24 [110/12] via 10.10.6.1, 00:37:44, Ethernet0
C      10.1.6.0/24 is directly connected, Loopback0
O IA   10.1.5.0/24 [110/12] via 10.10.6.1, 01:35:35, Ethernet0
O IA   10.1.4.0/24 [110/76] via 10.10.6.1, 01:35:17, Ethernet0
O IA   10.1.25.0/30 [110/75] via 10.10.6.1, 01:35:35, Ethernet0
O IA   10.1.21.0/24 [110/1573] via 10.10.6.1, 00:37:44, Ethernet0
O IA   10.1.45.0/30 [110/75] via 10.10.6.1, 01:35:35, Ethernet0
O IA   10.1.34.0/30 [110/75] via 10.10.6.1, 01:35:35, Ethernet0
C      192.168.1.0/24 is directly connected, Serial0
r6#

```

R7

```

r7#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, 3 subnets
C    172.16.28.0 is directly connected, Loopback0
C    172.16.29.0 is directly connected, Loopback1
C    172.16.30.0 is directly connected, Loopback2
R    10.0.0.0/8 [120/2] via 192.168.1.2, 00:00:03, Serial0
C    192.168.1.0/24 is directly connected, Serial0
r7#

```

Cat

```

cat#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

```

172.16.0.0/22 is subnetted, 1 subnets
O N1  172.16.28.0 [110/22] via 10.10.6.2, 00:19:07, Vlan60
10.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
C    10.10.1.0/30 is directly connected, Vlan10
C    10.1.10.0/24 is directly connected, Loopback0
O    10.10.2.0/24 [110/75] via 10.10.5.2, 00:19:07, Vlan50
      [110/75] via 10.10.1.2, 00:19:07, Vlan10
C    10.10.3.0/30 is directly connected, Vlan30
O IA  10.10.4.0/24 [110/75] via 10.10.5.2, 00:19:07, Vlan50
      [110/75] via 10.10.3.2, 00:19:08, Vlan30
C    10.10.5.0/30 is directly connected, Vlan50
C    10.10.6.0/30 is directly connected, Vlan60
O    10.1.12.0/30 [110/65] via 10.10.1.2, 00:19:08, Vlan10
O    10.1.3.0/24 [110/2] via 10.10.3.2, 01:33:46, Vlan30
O    10.1.2.0/24 [110/66] via 10.10.5.2, 00:19:08, Vlan50
      [110/66] via 10.10.1.2, 00:19:08, Vlan10
O    10.1.1.0/24 [110/2] via 10.10.1.2, 00:19:08, Vlan10
O    10.1.6.0/24 [110/2] via 10.10.6.2, 01:33:57, Vlan60
O    10.1.5.0/24 [110/2] via 10.10.5.2, 01:33:47, Vlan50
O IA  10.1.4.0/24 [110/66] via 10.10.5.2, 00:19:09, Vlan50
      [110/66] via 10.10.3.2, 00:19:09, Vlan30
O    10.1.25.0/30 [110/65] via 10.10.5.2, 01:33:47, Vlan50

```

```
O      10.1.21.0/24 [110/1563] via 10.10.1.2, 00:19:09, Vlan10
O      10.1.45.0/30 [110/65] via 10.10.5.2, 01:33:47, Vlan50
O      10.1.34.0/30 [110/65] via 10.10.3.2, 01:33:47, Vlan30
O N1 192.168.1.0/24 [110/22] via 10.10.6.2, 00:19:09, Vlan60
cat#
```

Explanation

1. Configure all Ethernet links and loopback addresses as shown in the diagram and according to the chart below. Configure appropriate VLANs on Cat, and enable the spanning tree portfast feature on all Cat switch access ports.

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

The configuration of the switch is straightforward. Each active switch port is in a separate VLAN and four ports connect to four different routers.

The beauty of VLAN-capable switches in lab scenarios is that you can connect just about any router to any other router in the scenario. We need not worry about losing the LAN interface, because we can always use loopback interfaces as substitutes for LANs. Using loopbacks works fine for most purposes, except when it comes to testing something like DLSW+, where you really need a LAN with actual hosts.

2. Insure that the switch ports specified in the chart above cannot become trunk ports.

It is a good idea to nail down a switch port as a switch port using "switchport mode access." I did not adopt this policy until very recently, and I will probably forget to do it from time to time, but I can assure you it is a good policy. The alternative is to merely set the VLAN for the port using syntax like "switchport access vlan 20." The problem with this is that if you hook this port to a trunk port, you could end up trunking without meaning to trunk.

It is often a good idea to nail things down in CCIE practice scenarios so that you know exactly what is going on. Some people think that you should hard code speed and duplex settings on 10/100 interfaces to prevent problems with auto-negotiation. I am not convinced that this is necessarily a good practice in an all-Cisco environment, but auto-negotiation is often a problem when connecting 10/100 cards from different vendors. I would also hesitate to hard-code speed and duplex on switch ports intended for individual users of a production network. It could create more service tickets than it prevents. That Vice-President who, for whatever reason, has a Notebook NIC set to

100Full may be unhappy that he or she cannot connect to the network at all because you hard coded the switchport to match the 10Half NIC he or she used until today.

On the other hand, you might gain performance improvements by hard-coding the speed/duplex settings of your various servers' NICs, but there again, make sure your server folks are aware of what you have done on the switch.

3. Configure all frame-relay links as shown in the diagram. Do not use any frame relay inverse-arp or frame relay subinterfaces.

This task implicitly requires us to use frame relay map statements to nail down what inverse-arp would otherwise have provided. When you see frame relay configured directly on physical interfaces, one problem that should leap to mind is using only the DLCIs that you are supposed to be using. Two other issues that should spring to mind are potential interface-type mismatches in OSPF or IS-IS, and split horizon issues in RIP and EIGRP. Neither of these latter issues are implicated in this scenario. Often in a CCIE practice lab setting you will have a full-mesh frame-relay network available, but the scenario involves a hub-and-spoke. Here we have no hub-and-spoke network, but we do not want R2 to try to talk to R4 using DLCI 204.

When configuring frame relay, it is safer to get everything configured on all frame relay interfaces on all routers before you bring up a single frame relay interface. If you make a mistake and see a mapping to 0.0.0.0/0 when you do "sh frame map," you will probably need to reboot your routers to get rid of it. If I remember correctly, one way this 0.0.0.0/0 route is created is by bringing up a frame relay interface before putting an IP address on it. Remember to include "no frame inverse-arp" under each frame relay interface. Rebooting takes time, which is why you should to be disciplined about the way you bring up your frame relay network and look over your configs before bringing up a single frame relay interface.

In this scenario, there is no language making it absolutely necessary that you provide a frame relay map statement so that you can ping the frame relay interface from the local router. Still, it is something I tend to do as a matter of course. I do not consider this sort of "over-configuration" a problem, but I am not a proctor. Also, I included an unnecessary "broadcast" at the end of the frame relay maps referring to the local interface itself. This reminds me of the question of whether to use 0.0.0.0 mask when configuring OSPF network statements. If you have the presence of mind to do everything precisely and never make any mistakes, then you should probably go ahead and demonstrate your superior command of detail.

4. Configure all other serial links as shown in the diagram using PPP encapsulation. Set the clock rate on each serial link to 125 kbps. Configure link authentication on each serial PPP link using the password "ppp." R2 should authenticate with R5 as "r2s1."

When you use "ppp chap hostname r2s1" on r2, you are holding out "r2s1" as the ppp hostname that this router will present on this link.

You may observe in the solution configs that the PPP password configured under interface s1 on r2 is encrypted, while the password configured globally on r5 to match this password is unencrypted. This is not a problem, and the link authentication will still work. If you want all passwords in a config file to show up encrypted, you can globally enable "service password-encryption." Encrypting passwords in config files is a good security policy, particularly in a network where the same small group of passwords are used throughout the network.

5. The OSPF areas are shown in the diagram. Note that area 6 is a Not-So-Stubby-Area. Cat and each router in the OSPF domain will use the loopback0 IP address as the OSPF Router ID. Configure plaintext authentication in area 0, using password "clear," and MD5 authentication in all other areas, using password "crypt."

In this scenario, configuring OSPF is divided among several tasks. If you look at R1's config script, you will see the "**network 0.0.0.0 255.255.255.255 area 0**" under **router ospf 64**. Since all interfaces on R1 participate in OSPF, and are in the same OSPF area, we can use this one line to bring all interfaces into OSPF. I do not hold this out as a best practice, but merely as something that can be done. Please note that e0 will attempt to form an OSPF adjacency, and will never be able to form one because there is no other OSPF router on the ethernet with which to form an adjacency. If you think to make the e0 interface passive under OSPF, you are very astute.

I set the OSPF network type on the frame relay links to point-to-point, since they are point-to-point links, with no hub-and-spoke situations in the entire scenario. Try to remember the default OSPF network-types for each kind of link. The default OSPF network type for a frame relay physical interface is non-broadcast. You can remember this because frame relay is non-broadcast multi-access (NBMA). You will find that non-broadcast as a network type is troublesome and you will often want to change it to either point-to-multipoint, point-to-point, or, if the scenario requires it, broadcast. If you are required to use the broadcast network type, then remember to deal with any potential DR/BDR election issues. **If the network type is broadcast, but the topology is a hub-and-spoke frame relay network, make sure that you set the spokes' priority to 0, so that neither spoke can become DR or BDR. It is not enough to insure that the hub becomes DR. You need to make sure that neither spoke can be DR or BDR.**

There are numerous points to remember about OSPF authentication. One is that you configure it under the interface and also enable it per area under OSPF. Also, **on the virtual-links use whatever kind of authentication you are using in area 0**. A virtual link is a connection to area 0, so you will use "area 0 authentication" under OSPF on a router that might look very distant from area 0 on a logical diagram. Finally, as with any passwords and authentication-keys, watch out for extra spaces at the end. Also, you enable authentication under OSPF but put the authentication details under each interface (except in the case of a virtual-link, where all authentication details relating to the virtual link are configured under OSPF).

One of the easiest mistakes to make is failing to put an interface in a routing protocol, or to put one in the wrong routing protocol. You can check which interfaces are in OSPF using "**show ip ospf interfaces**":

```

r3#sh ip ospf int
Serial0 is up, line protocol is up
  Internet Address 10.1.34.1/30, Area 3
  Process ID 64, Router ID 10.1.3.1, Network Type POINT_TO_POINT, Cost:
64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:03
  Index 3/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 19
  Last flood scan time is 0 msec, maximum is 8 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.4.1
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    No key configured, using default key id 0
Loopback0 is up, line protocol is up
  Internet Address 10.1.3.1/24, Area 3
  Process ID 64, Router ID 10.1.3.1, Network Type POINT_TO_POINT, Cost:
1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    No key configured, using default key id 0
Ethernet0 is up, line protocol is up
  Internet Address 10.10.3.2/30, Area 3
  Process ID 64, Router ID 10.1.3.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 10.1.10.1, Interface address 10.10.3.1
  Backup Designated router (ID) 10.1.3.1, Interface address 10.10.3.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:06
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 17
  Last flood scan time is 0 msec, maximum is 8 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.10.1 (Designated Router)
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    No key configured, using default key id 0
r3#

```

On R6, we have both OSPF and RIP running, so first let's use "**sh ip ospf interfaces**":

```

r6#sh ip ospf int
Ethernet0 is up, line protocol is up

```

```

Internet Address 10.10.6.2/30, Area 6
Process ID 64, Router ID 10.1.6.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 10.1.10.1, Interface address 10.10.6.1
Backup Designated router (ID) 10.1.6.1, Interface address 10.10.6.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:02
Index 2/2, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 3, maximum is 3
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 10.1.10.1 (Designated Router)
Suppress hello for 0 neighbor(s)
Message digest authentication enabled
  No key configured, using default key id 0
Loopback0 is up, line protocol is up
Internet Address 10.1.6.1/24, Area 6
Process ID 64, Router ID 10.1.6.1, Network Type POINT_TO_POINT, Cost:
1
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 0, maximum is 0
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 0, Adjacent neighbor count is 0
Suppress hello for 0 neighbor(s)
Message digest authentication enabled
  No key configured, using default key id 0
r6#

```

We need to use a different command for RIP— "**sh ip protocols**" -- and it will also give us a look at what is going on with other protocols.

```

r6#sh ip protocols
Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 19 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: rip, ospf 64 (internal, external 1 & 2, nssa-external
1 & 2)

  Default version control: send version 2, receive version 2
  Interface          Send  Recv  Triggered RIP  Key-chain
  Serial0           2    2
Automatic network summarization is not in effect
Routing for Networks:
  192.168.1.0
Passive Interface(s):
  Ethernet0
  Ethernet1
  Loopback0
  Serial1
Routing Information Sources:
  Gateway           Distance      Last Update

```



```

192.168.1.1      120      00:00:09
Distance: (default is 120)

```

```

Routing Protocol is "ospf 64"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: rip
  Address Summarization:
  Routing for Networks:
    10.1.6.0/24
    10.10.6.0/30
  Routing Information Sources:
    Gateway          Distance      Last Update
    10.1.10.1         110          01:21:57
    (this router)     110          01:22:27
  Distance: (default is 110)

```

```
r6#
```

6. Configure OSPF demand circuit on the ISDN link between R1 and R2. Use PPP CHAP authentication with "cisco" as the password.

This is a great feature which lets an OSPF adjacency remain in place over the ISDN link without the link staying up. It suspends the sending of OSPF hellos over the ISDN link. The link should go up once or twice when you first configure it. You should also expect it to go up every time you make a change in the network. Once you have finished configuring the network, it should stay down for very long time periods, as shown below:

```
r1#sh dialer
```

```
BRI0 - dialer type = ISDN
```

```

Dial String      Successes  Failures  Last DNIS  Last status
4082222222      1        0        01:47:17  successful
0 incoming call(s) have been screened.
0 incoming call(s) rejected for callback.

```

```

BRI0:1 - dialer type = ISDN
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re-enable (15 secs)
Dialer state is idle

```

```

BRI0:2 - dialer type = ISDN
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re-enable (15 secs)
Dialer state is idle
r1#

```

We can look at how long the router had been up using "**sh ver**", and learn that the last time the dialer interface was activated was soon after the router was booted up, which is what we want to see if we have not made any changes in the network.

```

r1#sh ver
Cisco Internetwork Operating System Software

```

IOS (tm) 2500 Software (C2500-JK8OS-L), Version 12.2(1d), RELEASE SOFTWARE (fcl)

Copyright (c) 1986-2002 by cisco Systems, Inc.
 Compiled Sun 03-Feb-02 22:01 by srani
 Image text-base: 0x0307EEEE0, data-base: 0x00001000

ROM: System Bootstrap, Version 11.0(10c), SOFTWARE
 BOOTFLASH: 3000 Bootstrap Software (IGS-BOOT-R), Version 11.0(10c),
 RELEASE SOFT
 WARE (fcl)

r1 uptime is 1 hour, 51 minutes

System returned to ROM by power-on
 System image file is "flash:/c2500-jk8os-l.122-1d.bin"

cisco 2500 (68030) processor (revision N) with 14336K/2048K bytes of memory.

Processor board ID 09756366, with hardware revision 00000001

Bridging software.

X.25 software, Version 3.0.0.

SuperLAT software (copyright 1990 by Meridian Technology Corp).

TN3270 Emulation software.

Basic Rate ISDN software, Version 1.1.

1 Ethernet/IEEE 802.3 interface(s)

2 Serial network interface(s)

1 ISDN Basic Rate interface(s)

32K bytes of non-volatile configuration memory.

16384K bytes of processor board System flash (Read ONLY)

Configuration register is 0x2102

r1#

ISDN can be tricky to get running right the first few times, and you will get very good at using the various debugging tools before you feel really comfortable with ISDN. I spent an enormous amount of time configuring ISDN and troubleshooting ISDN in the very early stages of my hands-on preparation for the CCIE lab exam, and I can remember reviewing some advanced points about ISDN early in the morning before my first lab attempt. (I do not recommend studying right up to the lab exam.) I have read that ISDN can be particularly helpful for data networkers to learn about because it allows us to learn a bit about telephone networks. This knowledge can be helpful in troubleshooting WAN links in general and in learning about IP telephony.

Let's see what happens when we bring down the frame relay link on R1:

r1#debug dialer

Dial on demand events debugging is on

r1#conf t

Enter configuration commands, one per line. End with CNTL/Z.

r1(config)#**int s0**

r1(config-if)#**shu**

r1(config-if)#

```

01:57:11: %OSPF-5-ADJCHG: Process 64, Nbr 10.1.2.1 on Serial0 from FULL
to DOWN,
  Neighbor Down: Interface down or detached
01:57:12: BR0 DDR: Dialing cause ip (s=10.1.21.1, d=224.0.0.5)
01:57:12: BR0 DDR: Attempting to dial 4082222222
01:57:12: %LINK-3-UPDOWN: Interface BRI0:1, changed state to up
01:57:13: %LINK-5-CHANGED: Interface Serial0, changed state to
administratively
down
01:57:14: BR0:1 DDR: dialer protocol up
01:57:14: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0,
changed state
to down
01:57:15: %LINEPROTO-5-UPDOWN: Line protocol on Interface BRI0:1,
changed state to up

```

You can see from the multicast destination IP address 224.0.0.5 that OSPF brought the ISDN link up. We want this to happen now, but if you see this happening pretty often, you probably have a problem. When redistribution is taking place on a router with an ISDN link, it is sometimes necessary to configure some sort of route-filtering on account of the ISDN link.

In this case, the ISDN link went up so that R1 could inform R2 of the change in link status. Data traffic did not use the ISDN link until I also brought down R1 e0. Things could work differently in your network if you used a low speed serial interface between R5 and R2.

Here is what my routing table looked like after I brought down both R1 s0 and R1 e0:

```

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

      172.16.0.0/22 is subnetted, 1 subnets
O E1    172.16.28.0 [110/1658] via 10.1.21.2, 00:02:06, BRIO
      10.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
O       10.1.10.0/24 [110/1637] via 10.1.21.2, 00:02:06, BRIO
O       10.10.1.0/30 [110/1637] via 10.1.21.2, 00:02:06, BRIO
O       10.10.2.0/24 [110/1572] via 10.1.21.2, 00:02:06, BRIO
O IA    10.10.3.0/30 [110/1637] via 10.1.21.2, 00:02:06, BRIO
O IA    10.10.4.0/24 [110/1700] via 10.1.21.2, 00:02:06, BRIO
O IA    10.10.5.0/30 [110/1636] via 10.1.21.2, 00:02:07, BRIO
O IA    10.10.6.0/30 [110/1637] via 10.1.21.2, 00:02:07, BRIO
O       10.1.12.0/30 [110/1626] via 10.1.21.2, 00:02:07, BRIO
O IA    10.1.3.0/24 [110/1638] via 10.1.21.2, 00:02:07, BRIO
O       10.1.2.0/24 [110/1563] via 10.1.21.2, 00:02:07, BRIO
C       10.1.1.0/24 is directly connected, Loopback0
O IA    10.1.6.0/24 [110/1638] via 10.1.21.2, 00:02:07, BRIO

```

```

O IA 10.1.5.0/24 [110/1627] via 10.1.21.2, 00:02:07, BRI0
O IA 10.1.4.0/24 [110/1691] via 10.1.21.2, 00:02:07, BRI0
O IA 10.1.25.0/30 [110/1690] via 10.1.21.2, 00:02:07, BRI0
C 10.1.21.0/24 is directly connected, BRI0
O IA 10.1.45.0/30 [110/1690] via 10.1.21.2, 00:02:07, BRI0
O IA 10.1.34.0/30 [110/1701] via 10.1.21.2, 00:02:07, BRI0
O E1 192.168.1.0/24 [110/1658] via 10.1.21.2, 00:02:07, BRI0
r1#

```

You can see that all data traffic will now use the ISDN link to go anywhere.

Before I brought down R1 e0, I wanted to make sure that an OSPF adjacency had formed over the ISDN link, so I did the following:

```
r2#sh ip ospf neigh
```

```

Neighbor ID      Pri   State           Dead Time   Address
Interface
10.1.1.1         1     FULL/ -         -           10.1.21.1   BRI0
10.1.5.1         1     FULL/ -         00:00:39   10.1.25.2
Serial1
r2#

```

Note that there is no dead time for the demand circuit. I have found that ISDN, particularly OSPF demand circuit, is a lot more trouble-free than it used to be. Perhaps part of this is that I understand ISDN better than I used to, but I gather that Cisco has made IOS improvements regarding ISDN.

If you are having trouble getting ISDN to go up, or it seems to go up for a moment but quickly drops again, check your PPP authentication ("**debug ppp authentication**") and/or PPP negotiation ("**debug ppp negotiation**"). Here is the progress of PPP authentication as the ISDN link goes up:

```

02:08:16: %LINK-3-UPDOWN: Interface BRI0:1, changed state to up
02:08:16: %ISDN-6-CONNECT: Interface BRI0:1 is now connected to
4082222222
r1(config-if)#
02:08:16: BR0:1 PPP: Treating connection as a callout
r1(config-if)#
r1(config-if)#
r1(config-if)#
r1(config-if)#
02:08:18: BR0:1 CHAP: O CHALLENGE id 4 len 23 from "r1"
02:08:18: BR0:1 CHAP: I CHALLENGE id 4 len 23 from "r2"
02:08:18: BR0:1 CHAP: O RESPONSE id 4 len 23 from "r1"
02:08:18: BR0:1 CHAP: I SUCCESS id 4 len 4
02:08:18: BR0:1 CHAP: I RESPONSE id 4 len 23 from "r2"
02:08:18: BR0:1 CHAP: O SUCCESS id 4 len 4
02:08:19: %LINEPROTO-5-UPDOWN: Line protocol on Interface BRI0:1,
changed state to up
02:08:24: %OSPF-5-ADJCHG: Process 64, Nbr 10.1.2.1 on Serial0 from
LOADING to FULL, Loading Done

```

The key line you want to see is the one with "success," obviously. The router also shows the OSPF state change to "**FULL**," which is what we want to see.

7. Configure virtual-links in the OSPF domain, and feel free to create some redundant virtual links.

Any area that is not directly connected to area 0 requires a virtual link to area 0. But what if the path used by one redundant link is lost? Perhaps you would want another virtual-link to connect that area to area 0?

In this scenario, I experimented some with redundant virtual-links, and found that some were not as necessary as I had initially thought they would be. I found OSPF to be remarkably resilient, probably because there were so many redundant paths. The config scripts show all the virtual links that I configured, even though one or more may not be strictly necessary, or would require many links to fail before it would become necessary. I tend to err on the side of "over-configuring" things. In this case, I simply want the reader to experiment with virtual-links, not to come up with some ideal, "best practices" configuration of virtual-links for this hideously ugly OSPF domain.

I configured redundant virtual-links so that various links could fail and there would be still be routes from, say, R1 in area 0 to R4's interfaces that lie in area 4. I experimented with configuring virtual links and then bringing down interfaces to force a particular virtual link to be used, and then tested further by disabling the one remaining "up" virtual-link. I found that it was pretty hard to break connectivity in the OSPF domain. Bringing down just one physical link (except the Cat-R6 link) will often not break much. Even when you manage to cut area 4 off from area 0 by bringing down all virtual links connecting area 0 and area 4, you may not spot the problem right away, because there will remain connectivity from area 0 to the two interfaces on R4 that do not lie in area 4. When experimenting, watch for the two networks that lie in area 4, 10.10.4.0/24, and 10.1.4.0/24.

Remember that **the virtual link references the OSPF Router ID of the router at the other end of the virtual link**. Using a loopback address as a Router-ID or manually setting a Router-ID makes configuration of virtual links much easier than using the highest physical address on a router. As noted above, remember to use the same authentication for virtual links that you use for area 0, since the virtual link is part of area 0.

When configuring a virtual-link, you may experience annoying and repetitive error messages relating to the virtual link before it is configured on both sides. This can be very distracting and you may save your sanity by turning off logging to the console. The danger inherent in doing this is that you may forget to turn logging back on, and then never get notified that your ISDN link keeps going up, or that some other problem exists on that router.

You can quickly view the status of virtual-links configured on a given OSPF router (or Cat 3550 switch functioning as an OSPF router) as follows:

```
cat#sh ip ospf virtual-links
Virtual Link OSPF_VL2 to router 10.1.2.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 245, via interface Vlan50, Cost of using 65
```

```

Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:09
  Adjacency State FULL (Hello suppressed)
  Index 3/6, retransmission queue length 0, number of retransmission
0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec
  Simple password authentication enabled
Virtual Link OSPF_VL1 to router 10.1.4.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 245, via interface Vlan50, Cost of using 65
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:08
  Adjacency State FULL (Hello suppressed)
  Index 4/7, retransmission queue length 0, number of retransmission
0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec
  Simple password authentication enabled
Virtual Link OSPF_VL0 to router 10.1.4.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 3, via interface Vlan30, Cost of using 65
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:05
  Adjacency State FULL (Hello suppressed)
  Index 2/5, retransmission queue length 0, number of retransmission
1
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec
  Simple password authentication enabled
cat#

```

Note that these are treated as demand circuits, and that the simple password authentication used in area 0 is configured.

8. R6 s0 and R7 s0, lo0, lo1, and lo2 will be part of the RIP domain, but only R7 s0 and R6 s0 will send RIP updates.

If you are running RIP on one interface only, and some other routing protocol on other interfaces which would be included in the RIP network statement's IP address range, it makes sense to make all interfaces passive in RIP, and disable passive interface on the

one interface on which you wish to send RIP updates. It is very easy to end up actively running RIP on interfaces that you did not intend for RIP to run on.

I noticed while checking over this scenario that I used passive interface on R6 when it was, strictly speaking, unnecessary. There is only one interface that falls in the classful range specified under RIP, so there was no need to nail things down the way I did. Still, I would rather adopt habits that result in occasional over-configuration than err in the other direction.

9. On R6, redistribute RIP into OSPF and OSPF into RIP. On R6, summarize the three routes for the R7 loopbacks to one network. Using a route-map but without referencing an access list, prevent this summary route from being redistributed into RIP. On R6, summarize all the 10.x.x.x routes to 10.0.0.0/8 before they reach R7. Insure that all interfaces in the OSPF domain (except ISDN interfaces, which could be idle) are reachable from R7, and that all interfaces on R7 are reachable from the OSPF domain.

This task involves mutual redistribution between OSPF and RIP version 2 on an ASBR that is in an NSSA. If you look at Cat's routing table, you will see two N1 routes. These are an NSSA's equivalent of E1 external routes. I prefer Type 1 routes to Type 2 routes because Type 1 routes take into account the internal distance to get to the ASBR, and are therefore likely to be more accurate, in general.

On an OSPF ASBR, you can easily summarize routes that are redistributed into OSPF using the "**summary-address**" command under "router ospf." To summarize the 10.x.x.x routes, you need to do it under the R6 s0 interface that participates in RIP.

If you do not do any route filtering, you may see on R7 a summary route to the R7 loopbacks. These loopbacks will never become unreachable from R7, but having a summary route to them pointing to a next-hop on R6 seems like bad form, so we need to keep that summary route from being redistributed back into RIP. There is more than one way to accomplish this goal, but I thought it would be interesting to use a route-map referencing the "bit bucket" null0, since it is the 172.16.28.0/22 summary route's next-hop address on R6 where redistribution occurs. Do not be concerned if you did not see this way to do this particular part of the task, because I just thought of it while developing this scenario. One beta tester used a prefix-list, and it worked for him.

10. Use R7 lo0 as the source of multicast packets. Enable multicast using a shared tree so that R4 lo0 and R2 lo0 can receive the multicast packets from R7 lo0, which will be addressed to the multicast group 239.0.0.1. Test your configuration by sending extended pings sourced from R7 lo0 to the multicast group.

This means enabling multicast on all interfaces that lead anywhere between these points, and on the loopbacks involved. The words "shared tree" call for a rendezvous point and "**pim sparse-mode**." It requires manual configuration of the rp-address on each router, but is very straightforward in its configuration. **One potential gotcha in this scenario is that you will need to include both icmp and pim in your access-list 106 applied inbound and outbound on R6 e0 in a later task.**

We can determine what pim neighbors we are communicating with as follows:

```

cat#sh ip pim neighbor
PIM Neighbor Table
Neighbor Address  Interface          Uptime    Expires    Ver
Mode
10.10.1.2         Vlan10            02:41:55  00:01:34  v2
(DR)
10.10.3.2         Vlan30            02:41:58  00:01:28  v2
(DR)
10.10.5.2         Vlan50            02:41:58  00:01:29  v2
(DR)
10.10.6.2         Vlan60            02:41:55  00:01:30  v2
(DR)
cat#

```

Please note that this is configured on Cat in its capacity as a router, and that pim configuration is done under the VLAN interfaces, and could be configured under router interfaces if we were using fastethernet interfaces configured as router interfaces.

You may have seen examples of pim sparse-mode where the router serving as the rendezvous point does not have the rp-address specified in its own configuration. It used to be that it was not necessary to specify the rp-address on the rendezvous point itself, but some of us who tend to over-configure things would do so anyway. I have read that in the more recent IOS revisions it is now required, so you should specify rp-address on all sparse-mode routers.

11. Configure frame relay traffic shaping on R3 s0 using a MINCIR of 32 kbps, a CIR of 64, and the ability to burst to 128 kbps. The rate should drop no lower than the MINCIR in the event of congestion indicated by BECN messages.

Unfortunately, we cannot test this properly because a router serving as a frame relay switch will not send BECN messages. We can, however, see the following output showing the various settings:

```

r3#sh traffic-shape s0

Interface  Se0
Access Target  Byte  Sustain  Excess  Interval  Increment  Adapt
List  Rate  Limit  bits/int  bits/int  (ms)      (bytes)    Active
302      64000  2000  8000    8000    125      1000      BECN
304      64000  2000  8000    8000    125      1000      BECN
301      64000  2000  8000    8000    125      1000      BECN
r3#

```

A few observations spring to mind about the output shown above. First, note that since everything is being done on the physical interface, rather than on subinterfaces, any traffic shaping is applied to all DLCIs, including DLCIs that are not being used in this scenario.

The interval is 125 milliseconds, which is 1/8 of a second. This is the default interval and is considered appropriate for data traffic. The byte limit, 2,000, is equal to 128,000 bits per second (the rate we can burst to), divided by 8 to get 16,000 bytes per second, divided

by 8 equals 2,000 bytes per interval. The normal bytes per interval is 1,000, which is 64,000 divided by 8 and then divided by 8 again.

You can see in last column that adaptive shaping based on BECNs is configured. Remember that you can do frame relay traffic shaping that is not adaptive.

We can also see how traffic-shaping is being viewed by the router as follows:

```
r3#sh frame pvc 304
PVC Statistics for interface Serial0 (Frame Relay DTE)
DLCI = 304, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0

  input pkts 2777          output pkts 2915          in bytes 266924
  out bytes 273148        dropped pkts 0           in FECN pkts 0
  in BECN pkts 0         out FECN pkts 0         out BECN pkts 0
  in DE pkts 0           out DE pkts 0
  out bcast pkts 1276    out bcast bytes 102272
Shaping adapts to BECN
pvc create time 01:15:55, last time pvc status changed 01:07:13
cir 64000    bc 8000    be 8000    byte limit 2000    interval 125
mincir 32000    byte increment 1000    Adaptive Shaping BECN
pkts 2915    bytes 273148    pkts delayed 72    bytes delayed 7472
shaping inactive
traffic shaping drops 0
Queueing strategy: fifo
Output queue 0/40, 0 drop, 72 dequeued
r3#
```

12. Configure R7, R4 and R3 so that they can be accessed by means of Telnet. Configure an access list on R6 s0 (inbound and outbound) permitting all ICMP traffic, RIP routing protocol traffic, and Telnet traffic between R7 s loopbacks and R4 s loopback. Test your configuration to make sure that R7 can Telnet to R4, and vice-versa, but R7 cannot Telnet to R3.

In this scenario, I specified that you need to permit RIP protocol traffic. In another practice scenario this requirement could be implied, and you will need to infer it from the circumstances. **Be very careful that any access-list you use does not stop a routing protocol or a necessary application from being able to run.**

Note that in this scenario the ACL is applied inbound and outbound. If the access-list had been only inbound, the access-list would have had fewer lines.

We can Telnet to R4 lo0 from R7 lo0:

```
r7#telnet 10.1.4.1 /source lo0
Trying 10.1.4.1 ... Open
```

User Access Verification

```
Password:
r4>
```

We cannot Telnet to R3 lo0 from R7 lo0:

```
r7#telnet 10.1.3.1 /source lo0
Trying 10.1.3.1 ...
% Destination unreachable; gateway or host down

r7#
```

We can Telnet to R7 lo2 from R4 lo0:

```
r4#telnet 172.16.30.1 /source lo0
Trying 172.16.30.1 ... Open
```

User Access Verification

```
Password:
r7>
```

We can Telnet to R3 lo0 from R4 lo0, because no access-list blocks the way:

```
r4#telnet 10.1.3.1
Trying 10.1.3.1 ... Open
```

User Access Verification

```
Password:
r3>
```

We had to manually set the source interface for Telnet, just as we often manually set the source interface from which we ping a destination address.

13. Configure rate limiting inbound and outbound on R6 e0 so that Telnet traffic will not exceed 64 kbps.

A better real-world application for rate-limiting would be rate-limiting FTP traffic, which can hog bandwidth. I specified Telnet traffic in this scenario because it is something we can actually test in a lab, even a rental lab accessed over the internet. You will probably never need to limit Telnet traffic to save bandwidth, because it is not greedy for bandwidth. In fact, it tends to require very little bandwidth. We all know about using Telnet to configure network equipment, but there are still a few important real-world applications that use Telnet. One example is the on-line database for the US Copyright Office, locis.loc.gov. For months, I Telneted to this database fairly often to see whether my Copyright Registration for my first book had gone into the database.

I Telneted from R4 lo0 to R7 lo0, and then looked at rate-limiting under the interface:

```
r6#sh interface e0 rate-limit
Ethernet0
```

Copyright © 2004, Thomas P. Larus

```

Input
  matches: access-group 107
    params: 64000 bps, 12000 limit, 24000 extended limit
    conformed 23 packets, 1284 bytes; action: transmit
    exceeded 0 packets, 0 bytes; action: drop
    last packet: 22968ms ago, current burst: 0 bytes
    last cleared 00:00:51 ago, conformed 0 bps, exceeded 0 bps
Output
  matches: access-group 107
    params: 64000 bps, 12000 limit, 24000 extended limit
    conformed 20 packets, 1166 bytes; action: transmit
    exceeded 0 packets, 0 bytes; action: drop
    last packet: 22980ms ago, current burst: 0 bytes
    last cleared 00:00:51 ago, conformed 0 bps, exceeded 0 bps
r6##

```

It looks like the command "sh interface e0 rate-limit hit the CPU 22-23 seconds after the last Telnet traffic came through and matched the access-list.

14. Configure Low Latency Queueing on R6 s0 so that Telnet traffic going onto the R6-R7 link will enjoy a strict priority queue with guaranteed bandwidth of 16 kbps. Allocate 20 kbps to traffic that is tagged with a DSCP value of 8 (which corresponds to the old IP precedence of 1), and use random early detection for that traffic.

Again, we use Telnet here because it lends itself to testing in a lab environment. A more typical application would be to put voice payload traffic into a considerably larger priority queue than this.

I appreciate Cisco's providing options to specify DSCP according to the IP precedence value.

Take a look at the following interactive help output:

```

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 dscp ?
<0-63>      Differentiated services codepoint value
af11       Match packets with AF11 dscp (001010)
af12       Match packets with AF12 dscp (001100)
af13       Match packets with AF13 dscp (001110)
af21       Match packets with AF21 dscp (010010)
af22       Match packets with AF22 dscp (010100)
af23       Match packets with AF23 dscp (010110)
af31       Match packets with AF31 dscp (011010)
af32       Match packets with AF32 dscp (011100)
af33       Match packets with AF33 dscp (011110)
af41       Match packets with AF41 dscp (100010)
af42       Match packets with AF42 dscp (100100)
af43       Match packets with AF43 dscp (100110)
cs1        Match packets with CS1(precedence 1) dscp (001000)
cs2        Match packets with CS2(precedence 2) dscp (010000)

```

```

cs3      Match packets with CS3(precedence 3) dscp (011000)
cs4      Match packets with CS4(precedence 4) dscp (100000)
cs5      Match packets with CS5(precedence 5) dscp (101000)
cs6      Match packets with CS6(precedence 6) dscp (110000)
cs7      Match packets with CS7(precedence 7) dscp (111000)
default  Match packets with default dscp (000000)
ef       Match packets with EF dscp (101110)

```

The value we want is "cs1." We can test by initiating a Telnet session from R4 lo0 to a loopback address on R7. Then we can look at the service policy on R6 s0 and see that Telnet packets matched a policy map on R6 s0:

```

r6#sh policy int s0 output
Serial0

```

Service-policy output: **s0policy**

```

Class-map: telnet (match-all)
 17 packets, 784 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match: access-group 106
Weighted Fair Queueing
  Strict Priority
  Output Queue: Conversation 264
  Bandwidth 16 (kbps) Burst 400 (Bytes)
  (pkts matched/bytes matched) 17/784
  (total drops/bytes drops) 0/0

```

```

Class-map: cs1 (match-all)
 0 packets, 0 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match: ip dscp 8
Weighted Fair Queueing
  Output Queue: Conversation 265
  Bandwidth 20 (kbps)
  (pkts matched/bytes matched) 0/0
  (depth/total drops/no-buffer drops) 0/0/0
  exponential weight: 9
  mean queue depth: 0

```

Class (Prec)	Random drop pkts/bytes	Tail drop pkts/bytes	Minimum threshold	Maximum threshold	Mark probability
0	0/0	0/0	20	40	1/10
1	0/0	0/0	22	40	1/10
2	0/0	0/0	24	40	1/10
3	0/0	0/0	26	40	1/10
4	0/0	0/0	28	40	1/10
5	0/0	0/0	30	40	1/10
6	0/0	0/0	32	40	1/10
7	0/0	0/0	34	40	1/10
rsvp	0/0	0/0	36	40	1/10

```

Class-map: class-default (match-any)
 516 packets, 62464 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match: any

```

r6# [smaller font size used here to permit columns to line up correctly].

Note that the class **class-default** is created automatically, and does not show up in the configuration unless you manually configure parameters on it, such as "random-detect."

No traffic has matched the "cs1" class-map because I do not have any host that is setting IP precedence to 1. If you have Cisco IP phones, voice payload traffic is set to IP precedence 5. Also, you might also put voice traffic into a priority queue based on RTP (Real-Time Protocol) ports used for voice traffic.

Chapter Two

Scenario Two

Scenario Two should not take long to configure if you follow the instructions carefully. Four Autonomous Systems peer with each other over a common ethernet (Cat VLAN 10), but there is also EBGP and IBGP peering using WAN links. Peering over a common ethernet permits us to have routers with one available ethernet interface peer with each other in a sort of logical full-mesh. You may find it helpful to draw your own BGP diagram to more explicitly show how each AS is peering with other autonomous systems.

It is worth noting that you do not need a Cat 3550-EMI switch to configure this scenario. One VLAN-capable switch will do fine, and you could even use simple hubs or switches.

Tasks

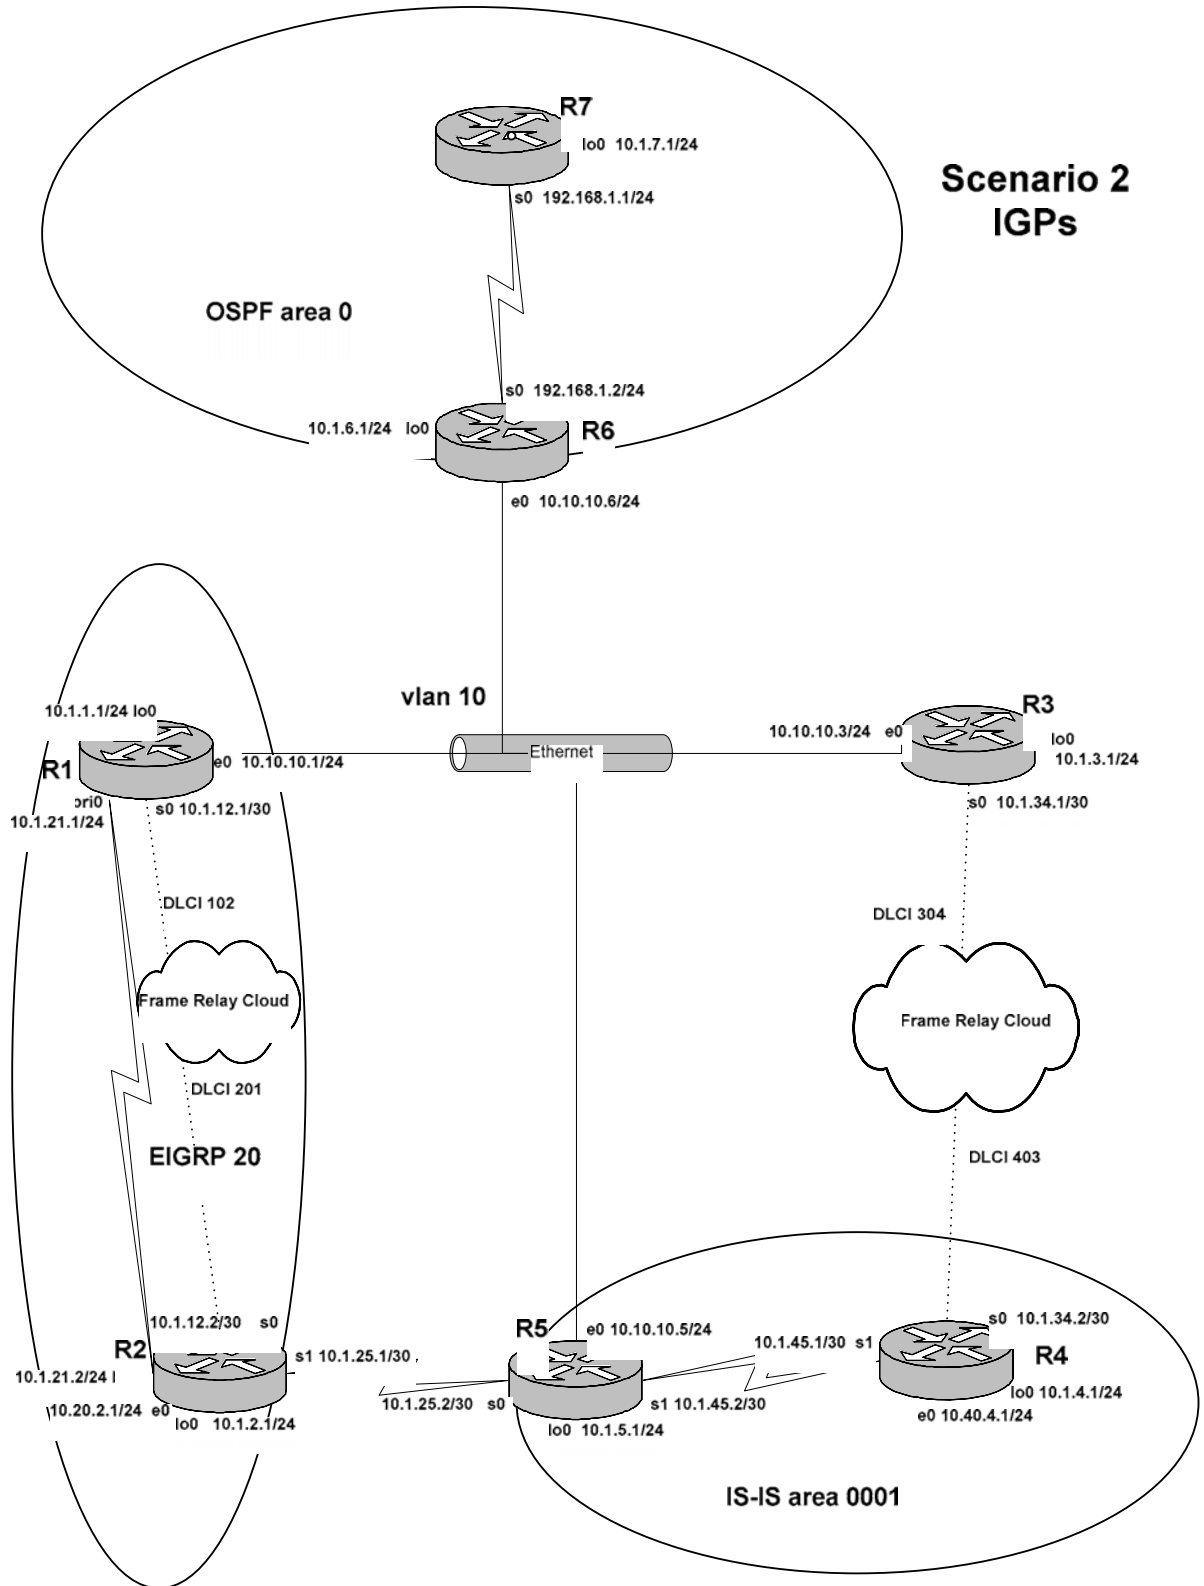
1. Configure the pertinent Cat 10/100 interfaces and router Ethernet interfaces as shown in the diagram and chart.

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

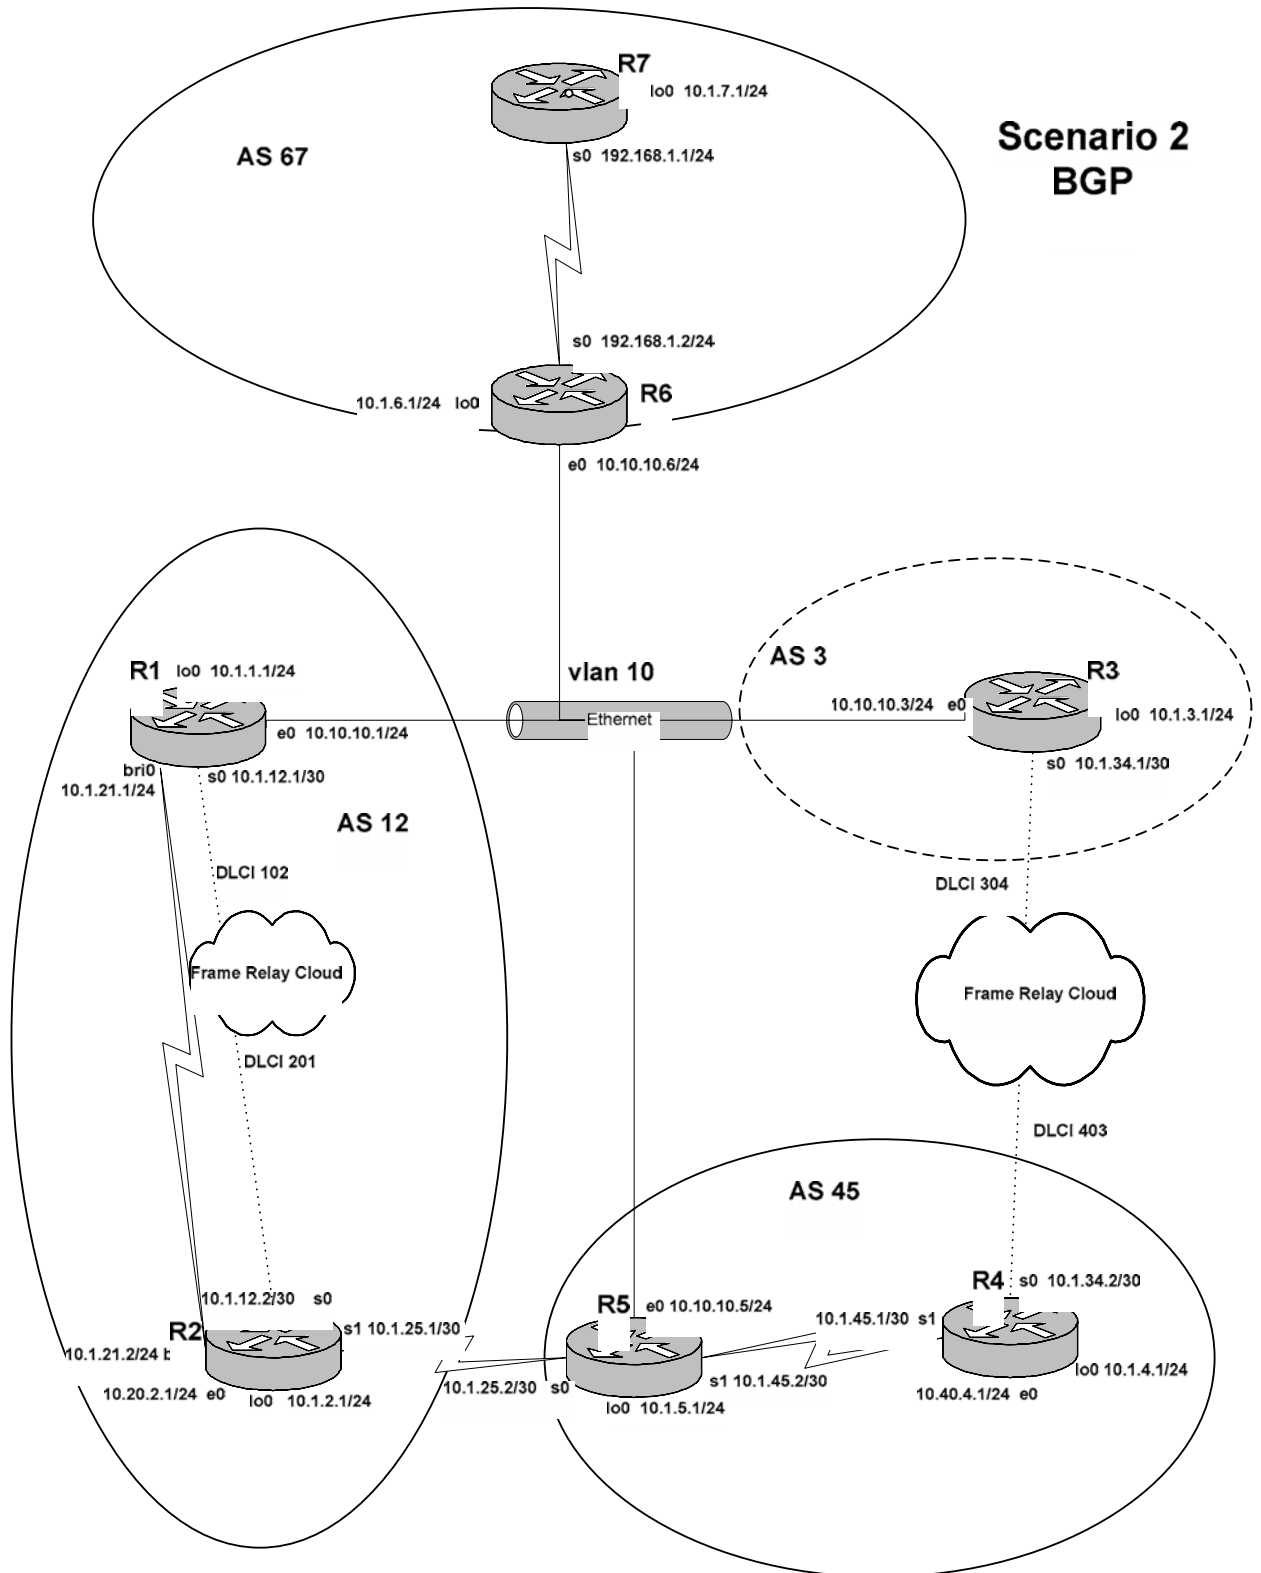
2. Configure all serial links as shown in the diagram using default encapsulation.
3. Configure all frame relay interfaces as shown in the diagram.
4. Configure R1 s0, e0, lo0, and bri0 in EIGRP AS 20. Configure R2 s0, s1, lo0, e0 and bri0 in EIGRP AS 20. Make interfaces passive where appropriate.
5. Configure R5 s1, e0 and lo0, and R4 s0, s1, e0, and lo0 in IS-IS area 0001. Only R5 s1 and R4 s1 will send IS-IS hellos.
6. Configure the ISDN link between R1 and R2 so that R1 will call R2 when the route to 10.20.2.0/24 is no longer in R1's routing table. Use legacy dial-on-demand rather than dialer profiles.
7. Do not configure any internal routing protocol on R3.
8. Configure R7 s0 and lo0 in OSPF area 0, as well as R6 s0, e0, and lo0.
9. Do not configure any route redistribution in this scenario.
10. Configure IBGP in all Autonomous Systems. Do not use any loopback addresses for peering.

11. Configure EBGP peering using the shared ethernet among the various autonomous systems.
12. Configure EBGP peering using the serial link between AS 12 and AS 45 and the frame relay link between AS 45 and AS 3.
13. Advertise in BGP the loopback0 networks on all routers, and also R4 e0. Do not advertise in BGP R2 e0 s network.
14. On R5, configure the MED attribute so that traffic entering AS 45 from AS 12 will normally prefer the R1-R5 Ethernet link. Configure the weight attribute on R5 so that traffic entering AS 12 from AS 45 will prefer the R1-R5 Ethernet link.
15. AS 45 does not want to be used as a transit AS. Configure BGP in AS 45 so that AS 45 will never be used as a transit AS. Test your configuration by shutting down R3 e0 and testing to make sure that you cannot ping from R3 lo0 to R1 lo0.
16. Configure DLSW+ on R2 and R4 so that Netbeui-only hosts attached to R2 e0 will be able to reach Netbeui-only hosts attached to R4 e0.
17. Configure TCP header-compression on the frame-relay link between R3 and R4. Do this using a frame-relay map statement.
18. Configure TCP header compression on the serial link between R2 and R5.

**Scenario 2
IGPs**



**Scenario 2
BGP**



Configuration Scripts

R1

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

Current configuration : 2482 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
 ip address 10.10.10.1 255.255.255.0
!
interface Serial0
 ip address 10.1.12.1 255.255.255.252
 encapsulation frame-relay
 frame-relay map ip 10.1.12.1 102
 frame-relay map ip 10.1.12.2 102 broadcast
 no frame-relay inverse-arp
!
interface Serial1
 no ip address
 shutdown
!
interface BRI0
 ip address 10.1.21.1 255.255.255.0
 encapsulation ppp
 dialer map ip 10.20.2.0 name r2 broadcast 4082222222
 dialer map ip 10.1.21.2 name r2 broadcast 4082222222
 dialer watch-group 1
 dialer-group 1
```

```

isdn switch-type basic-ni
isdn spid1 4081111111 4081111111
isdn spid2 4081111112 4081111111
cdapi buffers regular 0
cdapi buffers raw 0
cdapi buffers large 0
ppp authentication chap
!
router eigrp 20
  passive-interface default
  no passive-interface BRI0
  no passive-interface Serial0
  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 10.10.10.0 0.0.0.255
  no auto-summary
  no eigrp log-neighbor-changes
!
router bgp 12
  no synchronization
  bgp log-neighbor-changes
  bgp deterministic-med
  network 10.1.1.0 mask 255.255.255.0
  neighbor 10.1.12.2 remote-as 12
  neighbor 10.10.10.3 remote-as 3
  neighbor 10.10.10.5 remote-as 45
  neighbor 10.10.10.6 remote-as 67
  no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 130 deny eigrp any any
access-list 130 permit ip any any
dialer watch-list 1 ip 10.20.2.0 255.255.255.0
dialer-list 1 protocol ip list 130
!
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 : 2444 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
!
!
dlsw local-peer peer-id 10.1.2.1
dlsw remote-peer 0 tcp 10.1.4.1
dlsw bridge-group 1
!
!
interface Loopback0
 ip address 10.1.2.1 255.255.255.0
!
interface Ethernet0
 ip address 10.20.2.1 255.255.255.0
 bridge-group 1
!
interface Serial0
 ip address 10.1.12.2 255.255.255.252
 encapsulation frame-relay
 frame-relay map ip 10.1.12.1 201 broadcast
 frame-relay map ip 10.1.12.2 201
 no frame-relay inverse-arp
!
interface Serial1
 ip address 10.1.25.1 255.255.255.252
 ip tcp header-compression
 clockrate 1300000
!
interface BRI0
 ip address 10.1.21.2 255.255.255.0
 encapsulation ppp
 dialer-group 1
 isdn switch-type basic-ni
 isdn spid1 4082222221 4082222222
 isdn spid2 4082222222 4082222222
 cdapi buffers regular 0
 cdapi buffers raw 0
 cdapi buffers large 0

```

```

    ppp authentication chap
    !
router eigrp 20
  passive-interface default
  no passive-interface BRI0
  no passive-interface Serial0
  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
  network 10.20.2.0 0.0.0.255
  no auto-summary
  no eigrp log-neighbor-changes
  !
router bgp 12
  no synchronization
  bgp log-neighbor-changes
  bgp deterministic-med
  network 10.1.2.0 mask 255.255.255.0
  neighbor 10.1.12.1 remote-as 12
  neighbor 10.1.25.2 remote-as 45
  no auto-summary
  !
ip kerberos source-interface any
ip classless
ip http server
  !
access-list 130 deny    eigrp any any
access-list 130 permit ip any any
dialer-list 1 protocol ip list 130
  !
bridge 1 protocol ieee
  !
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 : 1581 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 10.10.10.3 255.255.255.0
!
interface Serial0
 ip address 10.1.34.1 255.255.255.252
 encapsulation frame-relay
 frame-relay map ip 10.1.34.1 304 tcp header-compression active
 frame-relay map ip 10.1.34.2 304 broadcast tcp header-compression
 active
 no frame-relay inverse-arp
!
interface Serial1
 no ip address
 shutdown
!
router bgp 3
 no synchronization
 bgp log-neighbor-changes
 network 10.1.3.0 mask 255.255.255.0
 neighbor 10.1.34.2 remote-as 45
 neighbor 10.10.10.1 remote-as 12
 neighbor 10.10.10.5 remote-as 45
 neighbor 10.10.10.6 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
```

r3#

R4

r4#sh runn
Building configuration...

```
Current configuration : 1961 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
!
!
dlsw local-peer peer-id 10.1.4.1
dlsw remote-peer 0 tcp 10.1.2.1
dlsw bridge-group 1
!
!
interface Loopback0
 ip address 10.1.4.1 255.255.255.0
!
interface Ethernet0
 ip address 10.40.4.1 255.255.255.0
 bridge-group 1
!
interface Serial0
 ip address 10.1.34.2 255.255.255.252
 encapsulation frame-relay
 frame-relay map ip 10.1.34.1 403 broadcast tcp header-compression
 active
 frame-relay map ip 10.1.34.2 403 tcp header-compression active
 no frame-relay inverse-arp
!
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
  passive-interface Ethernet0
  passive-interface Loopback0
  passive-interface Serial0
  net 49.0001.4444.4444.4444.00
!
router bgp 45
  no synchronization
  bgp log-neighbor-changes
  network 10.1.4.0 mask 255.255.255.0
  network 10.40.4.0 mask 255.255.255.0
  neighbor 10.1.34.1 remote-as 3
  neighbor 10.1.34.1 filter-list 150 out
  neighbor 10.1.45.2 remote-as 45
  no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
ip as-path access-list 150 permit ^$
!
!
bridge 1 protocol ieee
!
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 : 2038 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 10.10.10.5 255.255.255.0
!
interface Serial0
 ip address 10.1.25.2 255.255.255.252
 ip tcp header-compression
!
interface Serial1
 ip address 10.1.45.2 255.255.255.252
 ip router isis
!
router isis
 passive-interface Ethernet0
 passive-interface Loopback0
 passive-interface Serial0
 net 49.0001.5555.5555.5555.00
!
router bgp 45
 no synchronization
 bgp log-neighbor-changes
 network 10.1.5.0 mask 255.255.255.0
 neighbor 10.1.25.1 remote-as 12
 neighbor 10.1.25.1 weight 1
 neighbor 10.1.25.1 route-map setmed25 out
 neighbor 10.1.25.1 filter-list 150 out
 neighbor 10.1.45.1 remote-as 45
 neighbor 10.10.10.1 remote-as 12
 neighbor 10.10.10.1 weight 32767
 neighbor 10.10.10.1 route-map setmed out
 neighbor 10.10.10.1 filter-list 150 out
 neighbor 10.10.10.3 remote-as 3
 neighbor 10.10.10.3 filter-list 150 out
 neighbor 10.10.10.6 remote-as 67
 neighbor 10.10.10.6 filter-list 150 out
 no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
ip as-path access-list 150 permit ^$
!
route-map setmed25 permit 10
 set metric 1000
!
route-map setmed permit 10
 set metric 1
```

```

!
!
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 : 1654 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 10.10.10.6 255.255.255.0
!
interface Ethernet1
  no ip address
  shutdown
!
interface Serial0
  ip address 192.168.1.2 255.255.255.0
  clockrate 1300000
!
interface Serial1

```

```

no ip address
shutdown
!
router ospf 64
 log-adjacency-changes
 network 10.1.6.1 0.0.0.0 area 0
 network 10.10.10.6 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
 network 10.1.6.0 mask 255.255.255.0
 neighbor 10.10.10.1 remote-as 12
 neighbor 10.10.10.3 remote-as 3
 neighbor 10.10.10.5 remote-as 45
 neighbor 192.168.1.1 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

r6#

```

R7

```

r7#sh runn
Building configuration...

Current configuration : 1484 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 10.1.7.1 255.255.255.0
 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
!
interface Serial1
 no ip address
 shutdown
!
router ospf 64
 log-adjacency-changes
 passive-interface Loopback0
 network 10.1.7.1 0.0.0.0 area 0
 network 192.168.1.1 0.0.0.0 area 0
!
router bgp 67
 no synchronization
 bgp log-neighbor-changes
 network 10.1.7.0 mask 255.255.255.0
 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#sh runn
Building configuration...

Current configuration : 2404 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
  switchport mode access
  no ip address
!
interface FastEthernet0/2
  switchport access vlan 20
  switchport mode access
  no ip address
!
interface FastEthernet0/3
  switchport access vlan 10
  switchport mode access
  no ip address
!
interface FastEthernet0/4
  switchport access vlan 40
  switchport mode access
  no ip address
!
interface FastEthernet0/5
  switchport access vlan 10
  switchport mode access
  no ip address
!
interface FastEthernet0/6
  switchport access vlan 10
  switchport mode access
  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
no ip address
!
interface FastEthernet0/12
no ip address
!
interface FastEthernet0/13
no ip address
!
interface FastEthernet0/14
no ip address
!
interface FastEthernet0/15
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
switchport access vlan 20
switchport mode access
no ip address
!
interface FastEthernet0/21
no ip address
!
interface FastEthernet0/22
no ip address
!
interface FastEthernet0/23
no ip address
!
interface FastEthernet0/24
switchport access vlan 40
switchport mode access
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#
```


Routing Tables

R1

```
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

    10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
C       10.1.12.0/30 is directly connected, Serial0
B       10.1.3.0/24 [20/0] via 10.10.10.3, 00:29:48
D       10.1.2.0/24 [90/2297856] via 10.1.12.2, 00:29:48, Serial0
C       10.1.1.0/24 is directly connected, Loopback0
C       10.10.10.0/24 is directly connected, Ethernet0
B       10.1.7.0/24 [20/0] via 10.10.10.6, 00:29:48
B       10.1.6.0/24 [20/0] via 10.10.10.6, 00:29:48
B       10.1.5.0/24 [20/1] via 10.10.10.5, 00:29:48
B       10.1.4.0/24 [20/1] via 10.10.10.5, 00:29:49
D       10.1.25.0/30 [90/2681856] via 10.1.12.2, 00:29:49, Serial0
D       10.20.2.0/24 [90/2195456] via 10.1.12.2, 00:29:49, Serial0
C       10.1.21.0/24 is directly connected, BRI0
B       10.40.4.0/24 [20/1] via 10.10.10.5, 00:29:49
r1#
```

R2

```
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

    10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
C       10.1.12.0/30 is directly connected, Serial0
B       10.1.3.0/24 [200/0] via 10.10.10.3, 00:24:01
C       10.1.2.0/24 is directly connected, Loopback0
D       10.10.10.0/24 [90/2195456] via 10.1.12.1, 00:30:44, Serial0
```

```

D    10.1.1.0/24 [90/2297856] via 10.1.12.1, 00:30:44, Serial0
B    10.1.7.0/24 [200/0] via 10.10.10.6, 00:24:01
B    10.1.6.0/24 [200/0] via 10.10.10.6, 00:24:01
B    10.1.5.0/24 [200/1] via 10.10.10.5, 00:24:02
B    10.1.4.0/24 [200/1] via 10.10.10.5, 00:24:03
C    10.1.25.0/30 is directly connected, Serial1
C    10.20.2.0/24 is directly connected, Ethernet0
C    10.1.21.0/24 is directly connected, BRI0
B    10.40.4.0/24 [200/1] via 10.10.10.5, 00:24:03
r2#

```

R3

```

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

```

      10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
C    10.1.3.0/24 is directly connected, Loopback0
B    10.1.2.0/24 [20/0] via 10.10.10.1, 00:24:17
B    10.1.1.0/24 [20/0] via 10.10.10.1, 01:14:32
C    10.10.10.0/24 is directly connected, Ethernet0
B    10.1.7.0/24 [20/0] via 10.10.10.6, 01:15:07
B    10.1.6.0/24 [20/0] via 10.10.10.6, 01:15:07
B    10.1.5.0/24 [20/0] via 10.10.10.5, 01:14:29
B    10.1.4.0/24 [20/0] via 10.1.34.2, 01:14:42
B    10.40.4.0/24 [20/0] via 10.1.34.2, 01:14:43
C    10.1.34.0/30 is directly connected, Serial0
r3#

```

R4

```

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 not set

```

    10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
B       10.1.3.0/24 [20/0] via 10.1.34.1, 01:14:37
B       10.1.2.0/24 [200/0] via 10.10.10.1, 00:24:37
B       10.1.1.0/24 [200/0] via 10.10.10.1, 01:15:03
i L1    10.10.10.0/24 [115/10] via 10.1.45.2, Serial1
B       10.1.7.0/24 [20/0] via 10.1.34.1, 01:15:08
B       10.1.6.0/24 [20/0] via 10.1.34.1, 01:15:07
i L1    10.1.5.0/24 [115/10] via 10.1.45.2, Serial1
C       10.1.4.0/24 is directly connected, Loopback0
i L1    10.1.25.0/30 [115/10] via 10.1.45.2, Serial1
C       10.1.45.0/30 is directly connected, Serial1
C       10.40.4.0/24 is directly connected, Ethernet0
C       10.1.34.0/30 is directly connected, Serial0
r4#

```

R5

```
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

```

    10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
B       10.1.3.0/24 [20/0] via 10.10.10.3, 01:14:58
B       10.1.2.0/24 [20/0] via 10.10.10.1, 00:25:05
B       10.1.1.0/24 [20/0] via 10.10.10.1, 01:15:26
C       10.10.10.0/24 is directly connected, Ethernet0
B       10.1.7.0/24 [20/0] via 10.10.10.6, 01:15:56
B       10.1.6.0/24 [20/0] via 10.10.10.6, 01:15:56
C       10.1.5.0/24 is directly connected, Loopback0
i L1    10.1.4.0/24 [115/10] via 10.1.45.1, Serial1
C       10.1.25.0/30 is directly connected, Serial0
i L1    10.40.4.0/24 [115/10] via 10.1.45.1, Serial1
C       10.1.45.0/30 is directly connected, Serial1
i L1    10.1.34.0/30 [115/10] via 10.1.45.1, Serial1
r5#

```

R6

```

r6#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

```

      10.0.0.0/24 is subnetted, 9 subnets
B       10.1.3.0 [20/0] via 10.10.10.3, 01:15:53
B       10.1.2.0 [20/0] via 10.10.10.1, 00:25:33
B       10.1.1.0 [20/0] via 10.10.10.1, 01:16:03
C       10.10.10.0 is directly connected, Ethernet0
O       10.1.7.0 [110/65] via 192.168.1.1, 01:16:36, Serial0
C       10.1.6.0 is directly connected, Loopback0
B       10.1.5.0 [20/0] via 10.10.10.5, 01:15:32
B       10.1.4.0 [20/0] via 10.10.10.5, 01:16:02
B       10.40.4.0 [20/0] via 10.10.10.5, 01:16:02
C       192.168.1.0/24 is directly connected, Serial0
r6#

```

R7

```

r7#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

```

      10.0.0.0/24 is subnetted, 9 subnets
B       10.1.3.0 [200/0] via 10.10.10.3, 01:16:26
B       10.1.2.0 [200/0] via 10.10.10.1, 00:26:05
B       10.1.1.0 [200/0] via 10.10.10.1, 01:16:36
O       10.10.10.0 [110/74] via 192.168.1.2, 01:17:01, Serial0
C       10.1.7.0 is directly connected, Loopback0
O       10.1.6.0 [110/65] via 192.168.1.2, 01:17:01, Serial0
B       10.1.5.0 [200/0] via 10.10.10.5, 01:16:04
B       10.1.4.0 [200/0] via 10.10.10.5, 01:16:31
B       10.40.4.0 [200/0] via 10.10.10.5, 01:16:33
C       192.168.1.0/24 is directly connected, Serial0
r7#

```


Explanation

1. Configure the pertinent Cat 10/100 interfaces and router Ethernet interfaces as shown in the diagram and chart.

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

As I mentioned earlier, there is no need for a Cat 3550-EMI for this scenario. The switch is being used only for basic layer two functions.

2. Configure all serial links as shown in the diagram using default encapsulation.

Default encapsulation is HDLC. I have deliberately chosen not to weigh this scenario down with tasks involving link level authentication, multilink PPP, and the like.

3. Configure all frame relay interfaces as shown in the diagram.

Since we are using frame relay physical interfaces, we should disable inverse arp and use static frame relay mappings for IP. If you read through all the tasks before beginning to configure this one, you will see the frame-relay TCP header compression task and save a little time by incorporating that into this task.

4. Configure R1 s0, e0, lo0, and bri0 in EIGRP AS 20. Configure R2 s0, s1, lo0, e0 and bri0 in EIGRP AS 20. Make interfaces passive where appropriate.

Of course, you can manually make an interface passive, or you can make all interfaces passive by default and then make exceptions for individual interfaces.

5. Configure R5 s1, e0 and lo0, and R4 s0, s1, e0, and lo0 in IS-IS area 0001. Only R5 s1 and R4 s1 will send IS-IS hellos.

The passive interface command for IS-IS is strange. By simply specifying which interfaces are passive under the routing protocol you bring that interface's network into IS-IS. Also, you cannot use the "passive-interface" command for a particular interface until you have configured at least one other interface to actively run IS-IS.

There has been no instruction as to what MAC address to use in your NET. I like to use a MAC address based on the router number, but other practice scenarios may handle this

differently. I have read that it is good form to start the NET with "49" in lab scenarios, because 49 denotes a private address range. Do not forget to include the final ".00" at the end of the NET-- the NSAP Selector. This value is a bit like a TCP or UDP port number, and when it is set to 00 refers to the device itself, rather than some service running on the device.

6. Configure the ISDN link between R1 and R2 so that R1 will call R2 when the route to 10.20.2.0/24 is no longer in R1's routing table. Use legacy dial-on-demand rather than dialer profiles.

The main thing to remember about configuring dialer-watch with legacy dial-on-demand (using dialer map statements rather than dialer profiles) is that you need a dialer map pointing to the **network** that is being watched. This is a dramatic departure from the normal situation where your dialer map points to an IP address. Another odd thing about dialer-watch is that the dialer watch-list uses a normal mask (e.g. 255.255.255.0, rather than a wildcard mask (e.g. 0.0.0.255)).

I went to R1, enabled dialer debugging, and brought down s0, and saw this:

```
r1(config)#int s0
r1(config-if)#shu
r1(config-if)#
00:18:45: DDR: Dialer Watch: watch-group = 1
00:18:45: DDR: network 10.20.2.0/255.255.255.0 DOWN,
00:18:45: DDR: primary DOWN
00:18:45: DDR: Dialer Watch: Dial Reason: Primary of group 1 DOWN
00:18:45: DDR: Dialer Watch: watch-group = 1,
00:18:45: DDR: dialing secondary by dialer map 10.20.2.0 on BR0
00:18:45: BR0 DDR: Attempting to dial 4082222222
00:18:45: DDR: Dialer Watch: watch-group = 1
00:18:45: DDR: network 10.20.2.0/255.255.255.0 DOWN,
00:18:45: DDR: primary DOWN
00:18:45: DDR: Dialer Watch: Dial Reason: Primary of group 1 DOWN
00:18:45: DDR: Dialer Watch: watch-group = 1,
00:18:45: %LINK-3-UPDOWN: Interface BRI0:1, changed state to upexi
r1(config)#
00:18:45: BR0:1 DDR: Dialer Watch: resetting call in progress
00:18:47: %LINK-5-CHANGED: Interface Serial0, changed state to
administratively down
00:18:47: BR0:1 DDR: dialer protocol up
00:18:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0,
changed state to down
00:18:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface BRI0:1,
changed state to up
r1#
```

If we check the routing table, we will see that EIGRP worked very quickly and that the ISDN link is now used instead of the frame relay link for several routes:

```
r1#sh ip r
```

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

```

10.0.0.0/8 is variably subnetted, 14 subnets, 3 masks
D    10.1.12.0/30 [90/41024000] via 10.1.21.2, 00:00:09, BRI0
B    10.1.3.0/24 [20/0] via 10.10.10.3, 00:15:52
D    10.1.2.0/24 [90/40640000] via 10.1.21.2, 00:00:09, BRI0
C    10.1.1.0/24 is directly connected, Loopback0
C    10.10.10.0/24 is directly connected, Ethernet0
B    10.1.7.0/24 [20/0] via 10.10.10.6, 00:15:46
B    10.1.6.0/24 [20/0] via 10.10.10.6, 00:15:46
B    10.1.5.0/24 [20/1] via 10.10.10.5, 00:15:47
B    10.1.4.0/24 [20/1] via 10.10.10.5, 00:15:48
D    10.1.25.0/30 [90/41024000] via 10.1.21.2, 00:00:10, BRI0
D    10.20.2.0/24 [90/40537600] via 10.1.21.2, 00:00:10, BRI0
C    10.20.2.0/32 is directly connected, BRI0
C    10.1.21.0/24 is directly connected, BRI0
B    10.40.4.0/24 [20/1] via 10.10.10.5, 00:15:54
r1#

```

Since we brought down the frame relay link by shutting down R1 s0, there is still a route to the 10.1.12.0/30 network, because R2 s0 is still up. We can actually ping R2 s0 and the traffic uses the ISDN link.

If you have trouble getting dialer watch to work, you might want to reboot the router on which it is configured. If the watched route is not in the routing table when the router boots up, dialer watch may not bring up the ISDN link when the primary link goes down. Because of timing issues having to do with booting the routers up, you might even want to experiment with the command "**dialer watch-list group-number delay route-check initial seconds.**" For me, the problem usually turns out to be more fundamental, and is not something that will be cured by the delaying command mentioned here. **Be careful when you remove a line of configuration involving a dialer watch-list, because you may well remove the dialer watch-list itself.**

7. Do not configure any internal routing protocol on R3.

It looks odd to have no routing internal routing protocol running on this router, but it brings home the point that we will be relying on BGP for much of the connectivity in this scenario.

8. Configure R7 s0 and lo0 in OSPF area 0, as well as R6 S0, e0, and lo0.

It is very important to make sure that R7 can reach the next-hop address for BGP prefixes, which means it is essential that R7 have an internal route to the ethernet 10.10.10.0/24 network. You also need to make sure that you disable synchronization on R7, or it will not put many BGP prefixes into its routing table, because it will lack an internal route to them. (In practice scenarios, I would disable synchronization whenever I can.) If you remember late in the game to disable synchronization, you should then reset BGP on R7 and then you will be able to immediately ping between R7 lo0 and R4 lo0, for example:

```
r7#ping
Protocol [ip]:
Target IP address: 10.1.4.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.1.7.1
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 5, 100-byte ICMP Echos to 10.1.4.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/13/24 ms
r7#
```

I know we all love the BGP soft resets ("clear ip bgp * soft in," etc., but **in a lab setting** it is sometimes simpler and cleaner to simply go around to each router and reset BGP altogether.

9. Do not configure any route redistribution in this scenario.

This requirement simplifies the scenario considerably. It also means that there will not be any-to-any connectivity. When you test connectivity between a network advertised in BGP in one AS and a network advertised by BGP in another AS, you will need to use extended ping and specify the source interface. I can tell you that I still forget this sometimes, and it can cause unnecessary worry and loss of time.

10. Configure IBGP in all Autonomous Systems. Do not use any loopback addresses for peering.

We will be advertising in BGP only loopback addresses and R2 e0 and R4 e0.

11. Configure EBGP peering using the shared ethernet among the various autonomous systems.

The ethernet provides so many possible peerings that it would be very easy to leave one out. You may want to pencil in arrows on the BGP diagram to show the logic of the various peering relationships. In the real lab exam, you may or may not want to draw your own detailed diagram (there are good arguments both ways), but you probably

should draw up simple diagrams that show the logical relations among routers for the purposes of BGP, DLSW+, and the like.

A convenient way to see what neighbor relations are formed is to use the following command:

```
r1#sh ip bgp summary
BGP router identifier 10.1.1.1, local AS number 12
BGP table version is 11, main routing table version 11
8 network entries and 17 paths using 1388 bytes of memory
10 BGP path attribute entries using 600 bytes of memory
7 BGP AS-PATH entries using 168 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 9/65 prefixes, 21/4 paths, scan interval 15 secs
```

Neighbor State/PfxRcd	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	
10.1.12.2	4	12	19	33	11	0	0	00:00:50	1
10.10.10.3	4	3	63	61	11	0	0	00:51:30	6
10.10.10.5	4	45	56	61	11	0	0	00:51:37	3
10.10.10.6	4	67	61	60	11	0	0	00:51:47	6

```
r1#
```

You can see how long the BGP peering has been up, and how many messages have been received, and how many prefixes have been received. You can see that the BGP peering with 10.1.12.2 was recently brought up or reset.

12. Configure EBGP peering using the serial link between AS 12 and AS 45 and the frame relay link between AS 45 and AS 3.

No explanation needed.

13. Advertise in BGP the loopback0 networks on all routers, and also R4 e0. Do not advertise in BGP R2 e0 s network.

This task was crafted to avoid making this scenario too complicated. It also means that you cannot expect to be able to ping from any interface to any other interface. When testing connectivity between BGP-advertised networks in different autonomous systems, you will need to use extended ping and specify source and destination networks.

I did a quick test by looking at the BGP table on R7 and then doing an extended ping from R7 lo0 to R4 lo0. You can see that we have a BGP route to all the loopbacks and R4 e0:

```
r7#sh ip bgp
BGP table version is 10, local router ID is 10.1.7.1
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
*>i10.1.1.0/24    10.10.10.1        0      100     0 12 i
*>i10.1.2.0/24    10.10.10.1        0      100     0 12 i
*>i10.1.3.0/24    10.10.10.3        0      100     0 3 i
*>i10.1.4.0/24    10.10.10.5        0      100     0 45 i
*>i10.1.5.0/24    10.10.10.5        0      100     0 45 i
*>i10.1.6.0/24    192.168.1.2       0      100     0 i
*> 10.1.7.0/24    0.0.0.0           0              32768 i
*>i10.40.4.0/24   10.10.10.5        0      100     0 45 i
r7#ping
Protocol [ip]:
Target IP address: 10.1.4.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.1.7.1
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 5, 100-byte ICMP Echos to 10.1.4.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/13/20 ms
r7#

```

As I said before, it is easy to forget to use extended ping and specify the source address.

14. On R5, configure the MED attribute so that traffic entering AS 45 from AS 12 will normally prefer the R1-R5 Ethernet link. Configure the weight attribute on R5 so that traffic entering AS 12 from AS 45 will prefer the R1-R5 Ethernet link.

This task shows two ways to manipulate attributes to dictate the preferred path. The weight adjustment is a powerful yet crude tool, as you may have noticed.

We set weight on R5 so that this one router will favor the ethernet link. Since any traffic from AS 45 to AS 12 will normally go through this router, because it connects directly to AS 12 at two points and BGP will prefer the shortest AS-path (rather than the path through AS 3).

Weight is an extremely powerful Cisco-specific BGP attribute, in terms of the "weight" the router gives it when deciding which BGP prefix is best, but weight is local to one router only. Since weight outranks AS-path length in the BGP decision process, the specified weight adjustment on R5 will actually make R5 prefer a prefix advertised by AS 12 for traffic destined for AS 67:

```

r5#sh ip bgp
BGP table version is 14, 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
* 10.1.1.0/24 10.10.10.1    0      32767 12 i
*             10.10.10.1    0      67 12 i
*>            10.10.10.1    0      32767 12 i
*             10.1.25.1    1      12 i
* 10.1.2.0/24 10.10.10.1    0      32767 12 i
*>            10.10.10.1    32767 12 i
*             10.10.10.1    0      67 12 i
*             10.1.25.1    0      12 i
* 10.1.3.0/24 10.1.25.1     0      12 3 i
*             10.10.10.3    0      67 3 i
* i           10.1.34.1     0     100    0 3 i
*>            10.10.10.3    32767 12 3 i
*             10.10.10.3    0      0 3 i
*>i10.1.4.0/24 10.1.45.1     0     100    0 i
*> 10.1.5.0/24 0.0.0.0       0      32768 i
* 10.1.6.0/24 10.1.25.1     1     12 67 i
* i           10.1.34.1     0     100    0 3 67 i
*>           10.10.10.6     32767 12 67 i
Network      Next Hop      Metric LocPrf Weight Path
*             10.10.10.6    0      0 3 67 i
*             10.10.10.6    0      0 67 i
* 10.1.7.0/24 10.1.25.1     1     12 67 i
* i           10.1.34.1     0     100    0 3 67 i
*>           10.10.10.6     32767 12 67 i
*             10.10.10.6    0      0 3 67 i
*             10.10.10.6    0      0 67 i
*>i10.4.0.4.0/24 10.1.45.1     0     100    0 i
r5#

```

The fact that AS 12 appears in the AS-path to 10.1.6.0/24 is disturbing, because it suggests that a very sub-optimal route will be used, but note the next-hop address – 10.10.10.6 (R6 e0). Traffic to AS 67 from AS 45 will still go directly from R5 e0 to R6 e0, because of BGP's handling of next hop addresses on a shared medium like Ethernet.

This behavior of leaving the next-hop unchanged helped us in this case, but it can cause problems in a hub-and-spoke frame relay network. In the frame-relay hub-and-spoke situation, using "next-hop-self" on the hub router is a solution to the problem.

We set MED on R5, using two different route-maps applied to two separate peerings with AS 12 routers.

We set the MED on outgoing routes, and this value will be advertised to AS 12 and will go no further. We set the MED low for BGP prefixes advertised over the ethernet link to R1 and high for BGP prefixes advertised to over the serial link to R2. To help me to remember that low MED is preferred, I remember that MED is a "metric," and low metrics are preferred by routing protocols. Note that when you configure a route-map to manipulate the MED value, you refer to the MED as "metric."

The MED (multi-exit discriminator) attribute is a pretty weak attribute to manipulate, in the grand scheme of BGP attribute manipulation. Local-preference has high precedence in

the decision process, and is shared throughout an autonomous system. One can prepend AS numbers to an AS-path to make one path look longer to a router, even one located more than one autonomous system away from yours.

If you need to configure an attribute on a BGP speaker at the edge of your AS that will influence decisions in the next AS about which path into your AS is better, MED can do this. Look out for instructions indicating where attribute manipulation is to take place. Also, when dealing with BGP, think carefully about the direction of the traffic flows you are trying to influence, and the direction your advertisements are moving. It complicates matters that you can do attribute manipulation on incoming prefixes and on outgoing prefixes. The logic can get confusing, and is easy to configure the exact opposite of what is specified in a BGP task, as I have done myself in configuring practice labs.

I used "bgp deterministic-med" under "router bgp 12" on R1 and R2 to be on the safe side and make sure that MED is handled correctly and that the router's decision does not depend on the order in which prefixes are received.

If you look at R5's BGP table, you see that the best route to R1's loopback is the one with the next-hop of R1's ethernet. This is as things would be without any attribute manipulation. More tellingly, the best route to R2's loopback is the one with the next-hop of R1's ethernet:

```
r5#sh ip bgp
```

```
BGP table version is 14, 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
* 10.1.1.0/24	10.10.10.1			0 3 12	i
*	10.10.10.1			0 67 12	i
*>	10.10.10.1	0		32767 12	i
*	10.1.25.1			1 12	i
* 10.1.2.0/24	10.10.10.1			0 3 12	i
*>	10.10.10.1			32767 12	i
*	10.10.10.1			0 67 12	i
*	10.1.25.1	0		1 12	i
* 10.1.3.0/24	10.1.25.1			1 12 3	i
*	10.10.10.3			0 67 3	i
* i	10.1.34.1	0	100	0 3	i
*>	10.10.10.3			32767 12 3	i
*	10.10.10.3	0		0 3	i
*>i	10.1.45.1	0	100	0	i
*> 10.1.5.0/24	0.0.0.0	0		32768	i
* 10.1.6.0/24	10.1.25.1			1 12 67	i
* i	10.1.34.1		100	0 3 67	i
*>	10.10.10.6			32767 12 67	i
Network	Next Hop	Metric	LocPrf	Weight	Path
*	10.10.10.6			0 3 67	i
*	10.10.10.6	0		0 67	i
* 10.1.7.0/24	10.1.25.1			1 12 67	i
* i	10.1.34.1		100	0 3 67	i
*>	10.10.10.6			32767 12 67	i
*	10.10.10.6			0 3 67	i
*	10.10.10.6			0 67	i

```
*>i10.40.4.0/24      10.1.45.1          0    100    0 i
r5#
```

When we look at R2's BGP table, we see that the best routes to R4's loopback and ethernet LAN both use the R1-R5 ethernet link, and the best route to R5's loopback uses the R1-R5 ethernet link.

```
r2#sh ip bgp
```

```
BGP table version is 12, local router ID is 10.1.2.1
```

```
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
*>i10.1.1.0/24	10.1.12.1	0	100	0	i
*> 10.1.2.0/24	0.0.0.0	0		32768	i
*>i10.1.3.0/24	10.10.10.3	0	100	0	3 i
*>i10.1.4.0/24	10.10.10.5	1	100	0	45 i
*	10.1.25.2	1000		0	45 i
*>i10.1.5.0/24	10.10.10.5	1	100	0	45 i
*	10.1.25.2	1000		0	45 i
*>i10.1.6.0/24	10.10.10.6	0	100	0	67 i
*>i10.1.7.0/24	10.10.10.6		100	0	67 i
*>i10.40.4.0/24	10.10.10.5	1	100	0	45 i
*	10.1.25.2	1000		0	45 i

```
r2#
```

I tested using traceroute specifying the source address as well as the destination address:

```
r1#trace
```

```
Protocol [ip]:
```

```
Target IP address: 10.1.4.1
```

```
Source address: 10.1.1.1
```

```
Numeric display [n]:
```

```
Timeout in seconds [3]:
```

```
Probe count [3]:
```

```
Minimum Time to Live [1]:
```

```
Maximum Time to Live [30]:
```

```
Port Number [33434]:
```

```
Loose, Strict, Record, Timestamp, Verbose[none]:
```

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

```
Tracing the route to 10.1.4.1
```

```

  1 10.10.10.5 4 msec 8 msec 4 msec
  2 10.1.45.1 8 msec * 4 msec
```

```
r2#trace
```

```
Protocol [ip]:
```

```
Target IP address: 10.1.4.1
```

```
Source address: 10.1.2.1
```

```
Numeric display [n]:
```

```
Timeout in seconds [3]:
```

```
Probe count [3]:
```

```
Minimum Time to Live [1]:
```

```
Maximum Time to Live [30]:
```

```

Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 10.1.4.1

```

```

  1 10.1.12.1 8 msec 8 msec 4 msec
  2 10.10.10.5 8 msec 8 msec 4 msec
  3 10.1.45.1 12 msec * 8 msec

```

```
r2#
```

```
r4#trace
```

```

Protocol [ip]:
Target IP address: 10.1.2.1
Source address: 10.1.4.1
Numeric display [n]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 10.1.2.1

```

```

  1 10.1.45.2 8 msec 8 msec 4 msec
  2 10.10.10.1 4 msec 8 msec 8 msec
  3 10.1.12.2 8 msec * 8 msec

```

```
r4#
```

Note that the traceroute from R4 is influenced by the weight configured on R5, since R4 is using the path through R5 due to its having a shorter AS-path than the path through R3 (AS 3).

When manipulating attributes to influence best path decisions, remember the importance of **next-hop reachability**, because **a BGP prefix can have high local preference, low MED, high weight, and a short as-path, and yet be rejected if the next-hop for that path is not reachable**.

15. AS 45 does not want to be used as a transit AS. Configure BGP in AS 45 so that AS 45 will never be used as a transit AS. Test your configuration by shutting down R3 e0 and testing to make sure that you cannot ping from R3 lo0 to R1 lo0.

A simple way to keep an AS from being used as a transit AS is to permit only BGP routes (prefixes) with an empty as-path to be advertised from your AS to neighboring autonomous systems. The only BGP prefixes with an empty as-path will be prefixes that originate in your own AS.

You use "ip as-path access-list 1 permit ^\$" and apply it to a neighbor relationship using either a filter-list or a route-map.

Sometimes you may see an example of an ip as-path access-list where a second line, denying everything else, is included. The second line is not necessary, but can be included for logical clarity.

I used a filter-list here, but could just as easily have used a route-map. An IP as-path access-list can be numbered between 1 and 199.

16. Configure DLSW+ on R2 and R4 so that Netbeui-only hosts attached to R2 e0 will be able to reach Netbeui-only hosts attached to R4 e0.

This is a very simple DLSW+ task, to show how easy it is to get basic DLSW+ up and running. You should not be scared of DLSW+. I was intimidated by DLSW+ at first, but became pretty comfortable with it through practice with real hosts. Some of the DLSW+ features are simply not available in an environment containing no Token Ring LANs. To verify your DLSW+ configuration, it is good to have actual hosts running Netbeui only. The first thing I would do is check to see if the peer relationship between the routers has formed:

```
r2#sh dlsw peer
Peers:
state      pkts_rx  pkts_tx  type  drops  ckts  TCP  uptime
TCP 10.1.4.1    CONNECT    269      327  conf      0     1    0 0:23:52

Total number of connected peers: 1
Total number of connections: 1

r2#
```

The peer relationship will form immediately upon configuration, and you do not need to try to browse for Netbeui hosts or "net use" a Netbeui resource first. However, if you want to see what actual Netbeui resources you can reach using DLSW+, you will need to browse or attempt to "net use" a resource first. I used the Netbeui hosts DELL800, a desktop, and DELLNOTE, the notebook I use for writing my books. After browsing, I saw the following reachability cache:

```
r2#sh dlsw reach
DLSw Local MAC address reachability cache list
Mac Addr      status      Loc.      port      rif
0002.6902.9337  FOUND      LOCAL    TBridge-001  --no rif--
0040.5407.18b1  FOUND      LOCAL    TBridge-001  --no rif--

DLSw Remote MAC address reachability cache list
Mac Addr      status      Loc.      peer
00d0.db05.161a  FOUND      REMOTE  10.1.4.1 (2065) max-1f (1500)

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  10.1.4.1 (2065) max-1f (1500)
```


r2#

DLSW+ worked fine using TCP even though the desktop was connected to the rest of the lab network by a switch connected to a wireless bridge with a hub at the other end of the wireless bridge. If you have a choice of what encapsulation to use in a lab situation, you should use TCP because it is easy to make work reliably.

17. Configure TCP header-compression on the frame-relay link between R3 and R4. Do this using a frame-relay map statement.

TCP header compression needs to be configured on both sides of the link. I know I addressed this "gotcha" in the first book, but this task is different in that it requires configuration of frame relay TCP header compression in the frame relay map.

Ping testing will not uncover a problem in this regard. What would reveal the existence of the problem would be the failure of your BGP peering between R3 and R4.

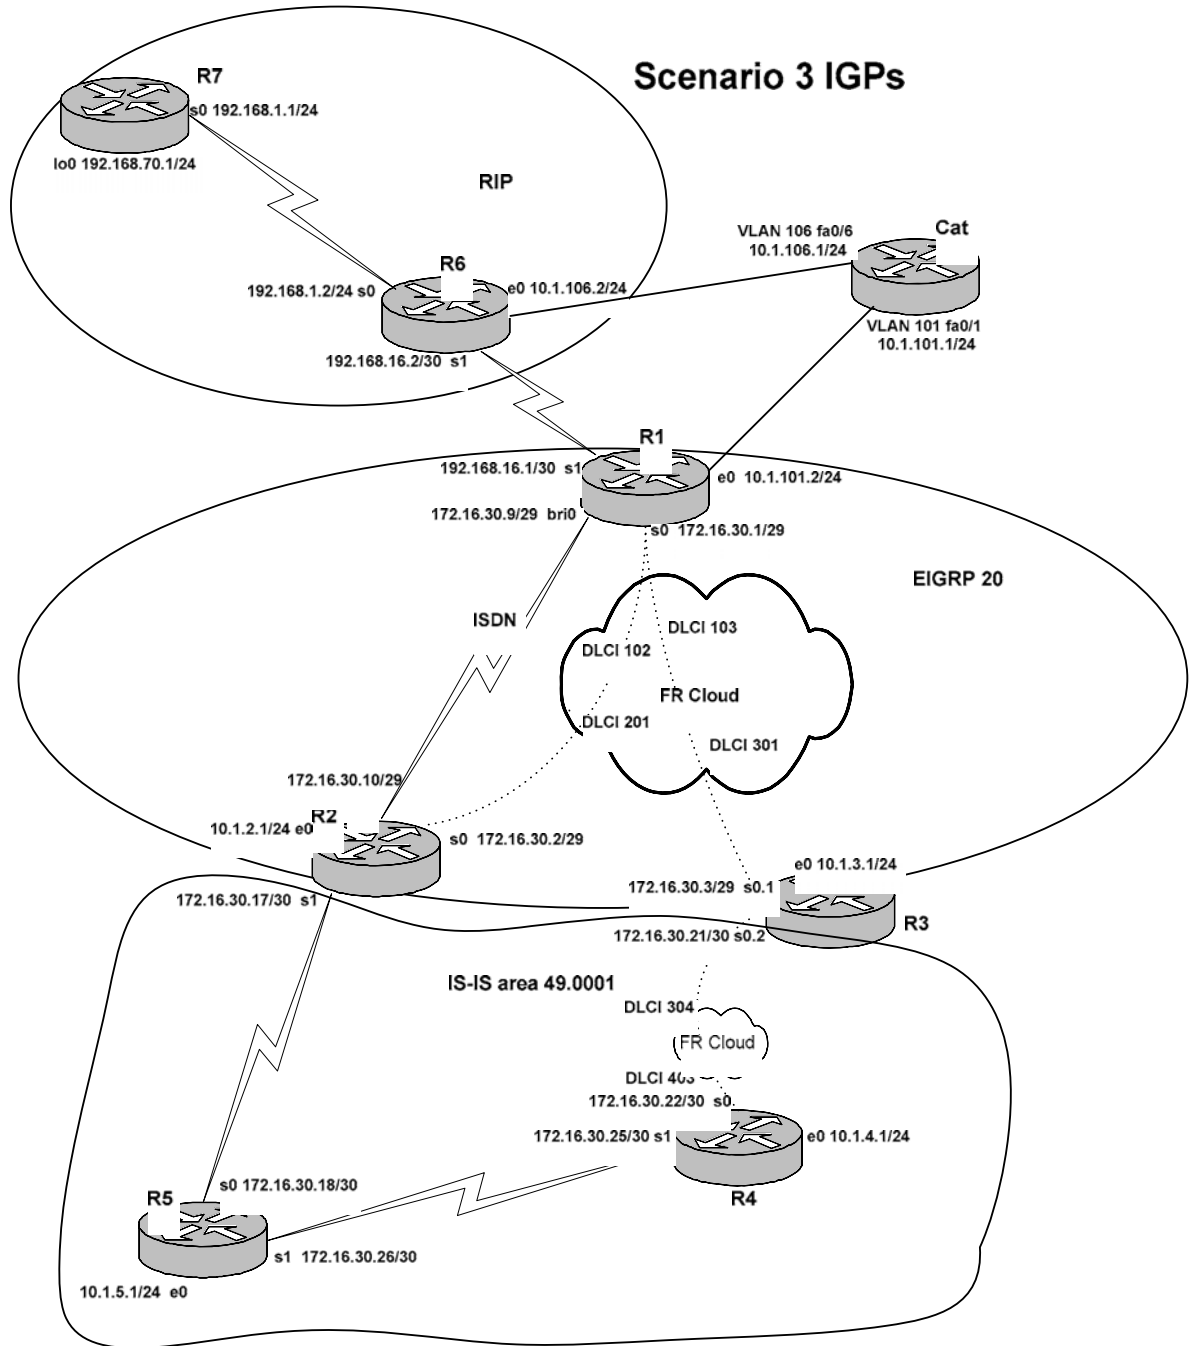
I tested this by configuring TCP header compression on one side of the R3-R4 link, and waited for the BGP session to close, and it did pretty soon. Once TCP header compression was configured on both sides, the peering relationship came back up very quickly.

18. Configure TCP header compression on the serial link between R2 and R5.

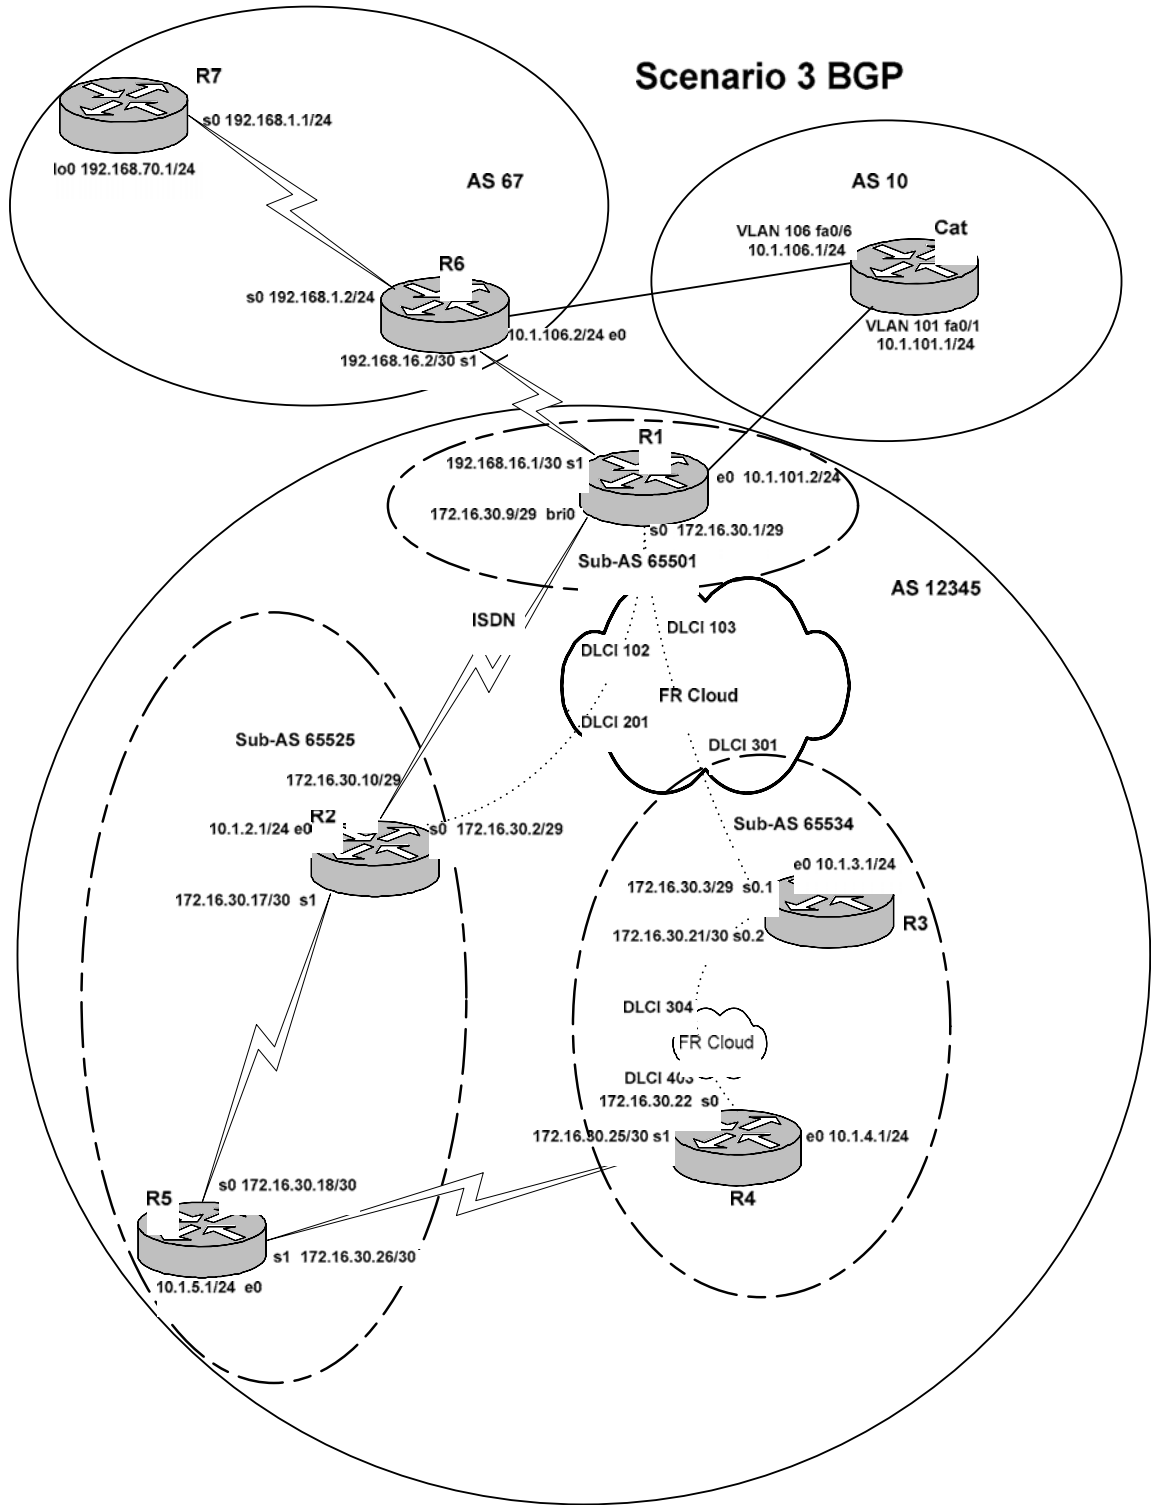
I include this task simply to show how easy it is to configure TCP header compression on a serial link. It is such a little thing to configure, but it can save a great deal of bandwidth on a link. Make sure to configure it on both sides of the link or BGP will fail between R2 and R5.

Chapter Three

Scenario Three



Scenario 3 BGP



Tasks

1. Configure frame relay links. Subinterfaces can only be used on R3.
2. Configure ethernet links as shown in the diagram and in the chart below. Note that you will configure on Cat an interface VLAN 101 and an interface VLAN 106.

Router	Interface	VLAN #	Switch port # (fastether 0/x)
R1	e0	101	1
R2	e0	20	2
R3	e0	30	3
R4	e0	40	4
R5	e0	50	5
R6	e0	106	6

3. Configure the serial interfaces using PPP. Configure CHAP authentication. Eliminate unnecessary host routes on PPP serial links.
4. Configure the EIGRP domain. Configure dialer watch so that R1 will bring up the ISDN link if 10.1.2.0/24 disappears from its routing table.
5. Configure IS-IS as shown in the diagram. You are free to use L1, L2, or L1/L2 as you like. Use passive interface to get networks into IS-IS over which no IS-IS adjacency could form.
6. Configure the RIP domain.
7. Cat will not participate in any internal routing protocol.
8. Configure route redistribution between EIGRP and IS-IS. Use route-maps and access-lists to prevent route feedback.
9. Configure BGP within AS 12345 without using any route reflector. R1 will form its own sub-AS (65501), and R2 and R5 will form their own sub-AS (65525) and R3 and R4 will form their own sub-AS (65534).
10. Configure BGP peering among AS 12345, AS 67 and AS 10.
11. Advertise in BGP all networks except R2 e0 s network and the networks used for peering.
12. R1 wants to use the ethernet link to AS 10 for user traffic. Configure BGP on R6 so that all destinations advertised from AS 12345 to AS 67 on the serial link will seem to AS 67 to be two Autonomous Systems further away than they really are.

13. Configure a BGP attribute on R1 so that it will always prefer the ethernet link to AS 10 for traffic destined for AS 10 or AS 67.

14. Permit AS 67 to carry traffic between AS 10 and AS 12345 in the unlikely event that the ethernet link between R1 and Cat were to fail but the ethernet link between R6 and Cat were to remain up.

Configuration Scripts

R1

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

Current configuration : 2161 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
username r6 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 10.1.101.2 255.255.255.0
!
interface Serial0
 ip address 172.16.30.1 255.255.255.248
 encapsulation frame-relay
 no ip split-horizon eigrp 20
 frame-relay map ip 172.16.30.1 102
 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 192.168.16.1 255.255.255.252
 encapsulation ppp
 no peer neighbor-route
 ppp authentication chap
!
interface BRI0
 ip address 172.16.30.9 255.255.255.248
 encapsulation ppp
 dialer map ip 10.1.2.0 name r2 broadcast 4082222222
```

```

dialer map ip 172.16.30.10 name r2 broadcast 4082222222
dialer watch-group 1
dialer-group 1
  isdn switch-type basic-ni
  isdn spid1 40811111111 4081111111
  isdn spid2 40811111112 4081111111
  cdapi buffers regular 0
  cdapi buffers raw 0
  cdapi buffers large 0
  no peer neighbor-route
  ppp authentication chap
!
router eigrp 20
  passive-interface Serial1
  network 10.1.101.0 0.0.0.255
  network 172.16.30.0 0.0.0.7
  network 172.16.30.8 0.0.0.7
  network 192.168.16.0 0.0.0.3
  no auto-summary
  no eigrp log-neighbor-changes
!
router bgp 65501
  no synchronization
  bgp log-neighbor-changes
  bgp confederation identifier 12345
  bgp confederation peers 65525 65534
  neighbor 10.1.101.1 remote-as 10
  neighbor 10.1.101.1 weight 65535
  neighbor 172.16.30.2 remote-as 65525
  neighbor 172.16.30.3 remote-as 65534
  neighbor 192.168.16.2 remote-as 67
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 101 deny eigrp any any
access-list 101 permit ip any any
dialer watch-list 1 ip 10.1.2.0 255.255.255.0
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 : 2547 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
username r5 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 split-horizon
 frame-relay map ip 172.16.30.1 201 broadcast
 frame-relay map ip 172.16.30.2 201
 frame-relay map ip 172.16.30.3 201 broadcast
 no frame-relay inverse-arp
!
interface Serial1
 ip address 172.16.30.17 255.255.255.252
 ip router isis
 encapsulation ppp
 no peer neighbor-route
 clockrate 1300000
 ppp authentication chap
!
interface BRI0
 ip address 172.16.30.10 255.255.255.248
 encapsulation ppp
 dialer map ip 172.16.30.9 name r2 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 eigrp 20
 redistribute connected route-map slonly
 redistribute isis level-2 route-map isistoeigrp
 passive-interface Ethernet0
 network 10.1.2.0 0.0.0.255
 network 172.16.30.0 0.0.0.7
 network 172.16.30.8 0.0.0.7
 default-metric 1500 20 255 1 1500
 no auto-summary
 no eigrp log-neighbor-changes
!
router isis
 redistribute eigrp 20 route-map eigrptoisis
 net 49.0001.2222.2222.2222.00
!
router bgp 65525
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 12345
 bgp confederation peers 65501 65534
 neighbor 172.16.30.1 remote-as 65501
 neighbor 172.16.30.18 remote-as 65525
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 50 permit 172.16.30.0 0.0.0.7
access-list 50 permit 10.1.3.0 0.0.0.255
access-list 50 permit 10.1.101.0 0.0.0.255
access-list 50 permit 192.168.16.0 0.0.0.3
access-list 50 permit 172.16.30.8 0.0.0.7
access-list 50 permit 10.1.2.0 0.0.0.255
access-list 101 deny eigrp any any
access-list 101 permit ip any any
dialer-list 1 protocol ip list 101
route-map slonly permit 10
 match interface Serial1
!
route-map isistoeigrp deny 10
 match ip address 50
!
route-map isistoeigrp permit 20
!
route-map eigrptoisis permit 10
 match ip address 50
!
!
!
line con 0
 transport input none
line aux 0

```

```

line vty 0 4
  login
!
end

```

```
r2#
```

R3

```

r3#sh runn
Building configuration...

Current configuration : 2162 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
!
interface Serial0
  no ip address
  encapsulation frame-relay
  no frame-relay inverse-arp
!
interface Serial0.1 multipoint
  ip address 172.16.30.3 255.255.255.248
  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
!
interface Serial0.2 multipoint
  ip address 172.16.30.21 255.255.255.252
  ip router isis
  frame-relay map clns 304 broadcast
  frame-relay map ip 172.16.30.21 304
  frame-relay map ip 172.16.30.22 304 broadcast
!
interface Serial1
  no ip address

```

```

shutdown
!
router eigrp 20
 redistribute connected route-map subint2only
 redistribute isis level-2 route-map isistoeigrp
 passive-interface Ethernet0
 network 10.1.3.0 0.0.0.255
 network 172.16.30.0 0.0.0.7
 default-metric 1500 20 255 1 1500
 no auto-summary
 no eigrp log-neighbor-changes
!
router isis
 redistribute eigrp 20 route-map eigrptois
 net 49.0001.3333.3333.00
!
router bgp 65534
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 12345
 bgp confederation peers 65501 65525
 network 10.1.3.0 mask 255.255.255.0
 neighbor 172.16.30.1 remote-as 65501
 neighbor 172.16.30.22 remote-as 65534
!
ip kerberos source-interface any
ip classless
ip http server
!
access-list 50 permit 172.16.30.0 0.0.0.7
access-list 50 permit 10.1.3.0 0.0.0.255
access-list 50 permit 10.1.101.0 0.0.0.255
access-list 50 permit 192.168.16.0 0.0.0.3
access-list 50 permit 172.16.30.8 0.0.0.7
access-list 50 permit 10.1.2.0 0.0.0.255
route-map subint2only permit 10
 match interface Serial0.2
!
route-map isistoeigrp deny 10
 match ip address 50
!
route-map isistoeigrp permit 20
!
route-map eigrptois permit 10
 match ip address 50
!
!
!
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 : 1348 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
!
username r5 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Ethernet0
 ip address 10.1.4.1 255.255.255.0
!
interface Serial0
 ip address 172.16.30.22 255.255.255.252
 ip router isis
 encapsulation frame-relay
 frame-relay map clns 403 broadcast
 frame-relay map ip 172.16.30.21 403 broadcast
 frame-relay map ip 172.16.30.22 403
 no frame-relay inverse-arp
!
interface Serial1
 ip address 172.16.30.25 255.255.255.252
 ip router isis
 encapsulation ppp
 no peer neighbor-route
 clockrate 1300000
 ppp authentication chap
!
interface TokenRing0
 no ip address
 shutdown
!
router isis
```

```

passive-interface Ethernet0
net 49.0001.4444.4444.4444.00
!
router bgp 65534
no synchronization
bgp log-neighbor-changes
bgp confederation identifier 12345
bgp confederation peers 65501 65525
network 10.1.4.0 mask 255.255.255.0
neighbor 172.16.30.21 remote-as 65534
neighbor 172.16.30.26 remote-as 65525
!
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#sh runn
Building configuration...

Current configuration : 1193 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
!
username r4 password 0 cisco
username r2 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!

```

```

!
!
interface Ethernet0
 ip address 10.1.5.1 255.255.255.0
!
interface Serial0
 ip address 172.16.30.18 255.255.255.252
 ip router isis
 encapsulation ppp
 no peer neighbor-route
 ppp authentication chap
!
interface Serial1
 ip address 172.16.30.26 255.255.255.252
 ip router isis
 encapsulation ppp
 no peer neighbor-route
 ppp authentication chap
!
router isis
 passive-interface Ethernet0
 net 49.0001.5555.5555.5555.00
!
router bgp 65525
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 12345
 bgp confederation peers 65501 65534
 network 10.1.5.0 mask 255.255.255.0
 neighbor 172.16.30.17 remote-as 65525
 neighbor 172.16.30.25 remote-as 65534
!
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

r5#

```

R6

```

r6#sh runn
Building configuration...

```

```
Current configuration : 1319 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 r1 password 0 cisco
username r7 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Ethernet0
 ip address 10.1.106.2 255.255.255.0
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.2 255.255.255.0
 encapsulation ppp
 no peer neighbor-route
 clockrate 1300000
!
interface Serial1
 ip address 192.168.16.2 255.255.255.252
 encapsulation ppp
 no peer neighbor-route
 clockrate 1300000
 ppp authentication chap
!
router rip
 version 2
 passive-interface default
 no passive-interface Serial0
 network 10.0.0.0
 network 192.168.1.0
 network 192.168.16.0
 no auto-summary
!
router bgp 67
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.1.106.1 remote-as 10
 neighbor 192.168.1.1 remote-as 67
```



```

neighbor 192.168.16.1 remote-as 12345
neighbor 192.168.16.1 route-map prep in
!
ip kerberos source-interface any
ip classless
ip http server
!
route-map prep permit 10
  set as-path prepend 12345 12345
!
!
!
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 : 1039 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
!
username r6 password 0 cisco
ip subnet-zero
no ip finger
no ip domain-lookup
!
no ip dhcp-client network-discovery
!
!
!
!
interface Loopback0
  ip address 192.168.70.1 255.255.255.0
!
interface Ethernet0

```

```
ip address 192.168.0.15 255.255.255.0
shutdown
!
interface Ethernet1
no ip address
shutdown
!
interface Serial0
ip address 192.168.1.1 255.255.255.0
encapsulation ppp
no peer neighbor-route
ppp authentication chap
!
interface Serial1
no ip address
shutdown
!
router rip
version 2
network 192.168.1.0
network 192.168.70.0
no auto-summary
!
router bgp 67
no synchronization
bgp log-neighbor-changes
network 192.168.70.0
neighbor 192.168.1.2 remote-as 67
!
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 : 2067 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 101
  switchport mode access
  no ip address
!
interface FastEthernet0/2
  switchport access vlan 20
  switchport mode access
  no ip address
!
interface FastEthernet0/3
  switchport access vlan 30
  switchport mode access
  no ip address
!
interface FastEthernet0/4
  switchport access vlan 40
  switchport mode access
  no ip address
!
interface FastEthernet0/5
  switchport access vlan 50
  switchport mode access
  no ip address
!
interface FastEthernet0/6
  switchport access vlan 106
  switchport mode access
  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
no ip address
!
interface FastEthernet0/12
no ip address
!
interface FastEthernet0/13
no ip address
!
interface FastEthernet0/14
no ip address
!
interface FastEthernet0/15
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 Vlan101
ip address 10.1.101.1 255.255.255.0
!
interface Vlan106
ip address 10.1.106.1 255.255.255.0
!
router bgp 10
```

```
no synchronization
bgp log-neighbor-changes
neighbor 10.1.101.2 remote-as 12345
neighbor 10.1.106.2 remote-as 67
!
ip classless
ip http server
!
!
!
!
line con 0
  exec-timeout 0 0
line vty 0 4
  login
line vty 5 15
  login
!
end

cat#
```

Routing Tables

R1

```

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, 5 subnets, 2 masks
C         172.16.30.0/29 is directly connected, Serial0
C         172.16.30.8/29 is directly connected, BRI0
D EX     172.16.30.16/30 [170/2223616] via 172.16.30.3, 02:04:12,
Serial0
D EX     172.16.30.20/30 [170/2223616] via 172.16.30.2, 02:04:12,
Serial0
D EX     172.16.30.24/30 [170/2223616] via 172.16.30.3, 02:04:12,
Serial0
                                                [170/2223616] via 172.16.30.2, 02:04:12,
Serial0
        10.0.0.0/24 is subnetted, 5 subnets
D         10.1.3.0 [90/2195456] via 172.16.30.3, 02:04:13, Serial0
D         10.1.2.0 [90/2195456] via 172.16.30.2, 02:04:13, Serial0
D EX     10.1.5.0 [170/2223616] via 172.16.30.3, 02:04:13, Serial0
                                                [170/2223616] via 172.16.30.2, 02:04:13, Serial0
D EX     10.1.4.0 [170/2223616] via 172.16.30.3, 02:04:13, Serial0
                                                [170/2223616] via 172.16.30.2, 02:04:13, Serial0
C         10.1.101.0 is directly connected, Ethernet0
        192.168.16.0/30 is subnetted, 1 subnets
C         192.168.16.0 is directly connected, Serial1
B        192.168.70.0/24 [20/0] via 10.1.101.1, 02:04:20
r1#

```

R2

```

r2#sh ip rou
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B -
BGP
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B -
BGP
        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
C       172.16.30.0/29 is directly connected, Serial0
C       172.16.30.8/29 is directly connected, BRI0
C       172.16.30.16/30 is directly connected, Serial1
i L1    172.16.30.20/30 [115/30] via 172.16.30.18, Serial1
i L1    172.16.30.24/30 [115/20] via 172.16.30.18, Serial1
    10.0.0.0/24 is subnetted, 5 subnets
D       10.1.3.0 [90/2707456] via 172.16.30.1, 02:04:48, Serial0
C       10.1.2.0 is directly connected, Ethernet0
i L1    10.1.5.0 [115/10] via 172.16.30.18, Serial1
i L1    10.1.4.0 [115/20] via 172.16.30.18, Serial1
D       10.1.101.0 [90/2195456] via 172.16.30.1, 02:04:49, Serial0
    192.168.16.0/30 is subnetted, 1 subnets
D       192.168.16.0 [90/2681856] via 172.16.30.1, 02:04:49, Serial0
B       192.168.70.0/24 [200/0] via 10.1.101.1, 02:04:26
r2#

```

R3

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, 5 subnets, 2 masks
C       172.16.30.0/29 is directly connected, Serial0.1
D       172.16.30.8/29 [90/41024000] via 172.16.30.1, 02:05:48,
Serial0.1
i L1    172.16.30.16/30 [115/30] via 172.16.30.22, Serial0.2
C       172.16.30.20/30 is directly connected, Serial0.2
i L1    172.16.30.24/30 [115/20] via 172.16.30.22, Serial0.2
    10.0.0.0/24 is subnetted, 5 subnets
C       10.1.3.0 is directly connected, Ethernet0
D       10.1.2.0 [90/2707456] via 172.16.30.1, 02:05:49, Serial0.1
i L1    10.1.5.0 [115/20] via 172.16.30.22, Serial0.2
i L1    10.1.4.0 [115/10] via 172.16.30.22, Serial0.2
D       10.1.101.0 [90/2195456] via 172.16.30.1, 02:05:49, Serial0.1
    192.168.16.0/30 is subnetted, 1 subnets
D       192.168.16.0 [90/2681856] via 172.16.30.1, 02:05:49, Serial0.1
B       192.168.70.0/24 [200/0] via 10.1.101.1, 02:05:29
r3#

```

R4

```
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 not set

```
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
i L2   172.16.30.0/29 [115/10] via 172.16.30.21, Serial0
i L2   172.16.30.8/29 [115/10] via 172.16.30.21, Serial0
i L1   172.16.30.16/30 [115/20] via 172.16.30.26, 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, 5 subnets
i L2   10.1.3.0 [115/10] via 172.16.30.21, Serial0
i L2   10.1.2.0 [115/10] via 172.16.30.21, Serial0
i L1   10.1.5.0 [115/10] via 172.16.30.26, Serial1
C      10.1.4.0 is directly connected, Ethernet0
i L2   10.1.101.0 [115/10] via 172.16.30.21, Serial0
      192.168.16.0/30 is subnetted, 1 subnets
i L2   192.168.16.0 [115/10] via 172.16.30.21, Serial0
B      192.168.70.0/24 [200/0] via 10.1.101.1, 02:05:55
r4#
```

R5

```
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
i L2   172.16.30.0/29 [115/10] via 172.16.30.17, Serial0
i L2   172.16.30.8/29 [115/10] via 172.16.30.17, Serial0
C      172.16.30.16/30 is directly connected, Serial0
i L1   172.16.30.20/30 [115/20] via 172.16.30.25, Serial1
C      172.16.30.24/30 is directly connected, Serial1
```



```

10.0.0.0/24 is subnetted, 5 subnets
i L2    10.1.3.0 [115/10] via 172.16.30.17, Serial0
i L2    10.1.2.0 [115/10] via 172.16.30.17, Serial0
C       10.1.5.0 is directly connected, Ethernet0
i L1    10.1.4.0 [115/10] via 172.16.30.25, Serial1
i L2    10.1.101.0 [115/10] via 172.16.30.17, Serial0
192.168.16.0/30 is subnetted, 1 subnets
i L2    192.168.16.0 [115/10] via 172.16.30.17, Serial0
B       192.168.70.0/24 [200/0] via 10.1.101.1, 02:06:25
r5#

```

R6

```

r6#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

```

10.0.0.0/24 is subnetted, 4 subnets
B       10.1.3.0 [20/0] via 10.1.106.1, 02:06:59
B       10.1.5.0 [20/0] via 10.1.106.1, 02:06:30
B       10.1.4.0 [20/0] via 10.1.106.1, 02:06:30
C       10.1.106.0 is directly connected, Ethernet0
192.168.16.0/30 is subnetted, 1 subnets
C       192.168.16.0 is directly connected, Serial1
C       192.168.1.0/24 is directly connected, Serial0
R       192.168.70.0/24 [120/1] via 192.168.1.1, 00:00:12, Serial0
r6#

```

R7

```

r7#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

```

10.0.0.0/24 is subnetted, 4 subnets
B       10.1.3.0 [200/0] via 10.1.106.1, 02:07:22

```

```

B      10.1.5.0 [200/0] via 10.1.106.1, 02:06:53
B      10.1.4.0 [200/0] via 10.1.106.1, 02:06:53
R      10.1.106.0 [120/1] via 192.168.1.2, 00:00:19, Serial0
      192.168.16.0/30 is subnetted, 1 subnets
R      192.168.16.0 [120/1] via 192.168.1.2, 00:00:19, Serial0
C      192.168.1.0/24 is directly connected, Serial0
C      192.168.70.0/24 is directly connected, Loopback0
r7#

```

Cat

```

cat#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

```

      10.0.0.0/24 is subnetted, 5 subnets
B      10.1.3.0 [20/0] via 10.1.101.2, 02:08:15
B      10.1.5.0 [20/0] via 10.1.101.2, 02:07:48
B      10.1.4.0 [20/0] via 10.1.101.2, 02:07:48
C      10.1.106.0 is directly connected, Vlan106
C      10.1.101.0 is directly connected, Vlan101
B      192.168.70.0/24 [20/0] via 10.1.106.2, 02:08:19
cat#

```

Explanation

1. Configure frame relay links. Subinterfaces can only be used on R3.

Be careful when you create your subinterfaces to make sure that they are the correct type of subinterfaces. If you look at the frame relay link between R3 and R4 and say to yourself "That is a point-to-point link, so I will set it up as a point-to-point subinterface," you will have stepped into a trap. IS-IS will be easier to configure if we create s0.2 as a multipoint subinterface. Configuring a subinterface as the wrong type is a very time-consuming error to remedy, because you have to get rid of any subinterfaces under the same physical interface, save the config, and reboot the router.

This brings up a related point about exam approach and psychology. No matter how well-prepared you feel, do not get overconfident as you configure a scenario or a real lab exam. When you feel like you are really moving along at a rapid clip, that is precisely when you are likely to make a mistake. At least that is the way it works with me.

2. Configure ethernet links as shown in the diagram and in the chart below. Note that you will configure on Cat an interface VLAN 101 and an interface VLAN 106.

Router	Interface	VLAN #	Switch port # (fastether 0/x)
R1	e0	101	1
R2	e0	20	2
R3	e0	30	3
R4	e0	40	4
R5	e0	50	5
R6	e0	106	6

If you look at the diagram, you will see that you create interface VLAN 106 and interface VLAN 101 and configure IP addresses for them. The switch will not permit you to do this unless you have first enabled IP routing on it.

3. Configure the serial interfaces using PPP. Configure CHAP authentication. Eliminate unnecessary host routes on PPP serial links.

This is an interesting situation, because on certain routers, you will be using authentication on more than one link, and will need to make it work with more than one router at a time. A simple way to handle this was to use the same password for all CHAP authentication, and simply configure a username with that password for EACH router with which the router be doing CHAP authentication. Thus, on R6, we have username R1 and username R7, and each username has the same password.

I originally did this scenario without any use of "no peer neighbor route." One of my technical reviewers suggested that I clean up my routing tables by using "no peer

neighbor route" on the PPP serial links as well as on the ISDN link, so I changed the task. I found it very difficult to get rid of the host routes once I added "no peer neighbor-route." Only after rebooting the routers were all the lingering host routes removed.

4. Configure the EIGRP domain. Configure dialer watch so that R1 will bring up the ISDN link if 10.1.2.0/24 disappears from its routing table.

There is a split-horizon issue here, because of the hub-and-spoke frame relay network. You need to disable split horizon for R1 s0. The interesting twist here is that EIGRP is not affected by the simple "ip split- horizon" or "no ip split-horizon" configured under an interface, so you need to disable split-horizon specifically for a particular EIGRP AS number.

I decided to see how the routing tables looked on R2 and R3 before I disabled split-horizon on R1, and it was just as you would expect. Each spoke could see the hub, but had no route to the other spoke.

To verify that dialer watch was functioning correctly, I typed "debug ip rout" and "debug dialer," and then brought down R1 s0. Here is the output:

```
00:45:55: is_up: 0 state: 6 sub state: 1 line: 0
00:45:55: RT: delete route to 10.1.2.0 via 172.16.30.2, eigrp metric
[90/2195456
]
00:45:55: RT: no routes to 10.1.2.0
00:45:55: DDR: Dialer Watch: watch-group = 1
00:45:55: DDR: network 10.1.2.0/255.255.255.0 DOWN,
00:45:55: DDR: primary DOWN
00:45:55: DDR: Dialer Watch: Dial Reason: Primary of group 1 DOWN
00:45:55: DDR: Dialer Watch: watch-group = 1,
00:45:55: RT: add 10.1.2.0/24 via 172.16.30.10, eigrp metric
[90/40537600]
00:45:55: RT: delete route to 172.16.30.18 via 172.16.30.2, eigrp
metric [170/26
81856]
00:45:55: RT: no routes to 172.16.30.18
00:45:55: RT: add 172.16.30.18/32 via 172.16.30.10, eigrp metric
[170/41024000]
00:45:55: RT: delete route to 172.16.30.20 via 172.16.30.2, eigrp
metric [170/22
23616]
00:45:55: RT: no routes to 172.16.30.20
00:45:55: RT: add 172.16.30.20/30 via 172.16.30.10, eigrp metric
[170/40517120]
00:45:55: RT: delete route to 172.16.30.26 via 172.16.30.2, eigrp
metric [170/22
2361xi
r1#6]
00:45:55: RT: no routes to 172.16.30.26
00:45:55: RT: add 172.16.30.26/32 via 172.16.30.10, eigrp metric
[170/40517120]
00:45:55: RT: delete route to 172.16.30.24 via 172.16.30.2, eigrp
metric [170/22
```

```

23616]
00:45:55: RT: no routes to 172.16.30.24
00:45:55: RT: add 172.16.30.24/30 via 172.16.30.10, eigrp metric
[170/40517120]
00:45:55: RT: delete route to 172.16.30.25 via 172.16.30.2, eigrp
metric [170/22
23616]
00:45:55: RT: no routes to 172.16.30.25
00:45:55: RT: add 172.16.30.25/32 via 172.16.30.10, eigrp metric
[170/40517120]
00:45:55: %BGP-5-ADJCHANGE: neighbor 172.16.30.2 Down Interface flap
00:45:55: %BGP-5-ADJCHANGE: neighbor 172.16.30.3 Down Interface flap
00:45:55: RT: interface Serial0 removed from routing table
00:45:55: RT: del 172.16.30.0/29 via 0.0.0.0, connected metric [0/0]
00:45:55: RT: delete subnet route to 172.16.30.0/29
00:45:55: RT: delete route to 10.1.3.0 via 172.16.30.3, Serial0
00:45:55: RT: no routes to 10.1.3.0, flushing
00:45:55: RT: del 10.1.4.0/24 via 172.16.30.25, bgp metric [200/0]
00:45:55: RT: delete subnet route to 10.1.4.0/24
00:45:55: RT: del 10.1.5.0/24 via 172.16.30.26, bgp metric [200/0]
00:45:55: RT: delete subnet route to 10.1.5.0/24
00:45:55: RT: add 172.16.30.0/29 via 172.16.30.10, eigrp metric
[90/41024000]
00:45:55: RT: add 10.1.3.0/24 via 172.16.30.10, eigrp metric
[90/41561600]
00:45:57: %SYS-5-CONFIG_I: Configured from console by console
00:45:57: %LINK-5-CHANGED: Interface Serial0, changed state to
administratively
down
00:45:57: is_up: 0 state: 6 sub state: 1 line: 0
00:45:58: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0,
changed state
to down
00:45:58: is_up: 0 state: 6 sub state: 1 line: 0

```

This ISDN backup restores EIGRP routing between R1 and R2, but will not restore the BGP session we lost when we brought down R1 s0.

I specified that R2 would not advertise its e0 network in BGP because if there is a BGP prefix, the existence of that BGP prefix will delay the operation of dialer watch on R1, because the BGP prefix will remain in R1's routing table longer than the EIGRP route.

If you are having trouble with dialer watch, here are some things to do to troubleshoot:

- 1) Make sure that the ISDN link works by pinging from one side of the link to the other.
- 2) Check the network address and mask of the watched route(s). The mask is an ordinary mask, not a wildcard mask. It should match the route for the network precisely.
- 2) Reboot the router on which dialer watch is configured so that the primary route will be in the routing table to be watched. **If you have shut down the primary interface to test dialer watch, you need to bring it back up and save the config.** If the watched route is not in the routing table when the router boots up with dialer-watch configured, do not expect dialer-watch to work.

5. Configure IS-IS as shown in the diagram. You are free to use L1, L2, or L1/L2 as you like. Use passive interface to get networks into IS-IS over which no IS-IS adjacency could form.

There is the potential for an IS-IS interface-type mismatch problem. If you configured R3 s0.2 as a point-to-point subinterface, there will be an IS-IS interface-type mismatch error with R4, since R4 is a physical interface. A frame relay physical interface will match up nicely with a frame relay multipoint subinterface, so the simplest way to make things work is to make R3 s0.2 a multipoint interface. A tunnel can also be used to get around the problem, but is a more complicated solution. As of IOS 12.2(8)T, the command "isis network-type point-to-point" provides another way to resolve some IS-IS interface mismatch problems, by having a broadcast interface (or multipoint frame relay interface) send point-to-point hellos like a point-to-point interface.

On R2 and R3, I redistributed into EIGRP the one "**connected**" network that runs IS-IS. You may have chosen another approach, but this is what I did because IS-IS did not redistribute it into EIGRP. IS-IS has an odd way of dealing with connected routes, which often makes this sort of workaround necessary when doing IGP redistribution on a router running IS-IS.

Passive interface is really neat in IS-IS, because it instantly puts the network into IS-IS without your needing to put anything under the interface. It also deletes "ip router isis" from under the interface that has been made passive after having first been configured to actively participate in IS-IS.

6. Configure the RIP domain.

Remember to get R6 s1 and R6 e0 into RIP, even though they will not be sending RIP updates. You can redistribute connected networks into RIP on R6 to get them in or simply use network statements to put them into RIP and then make them passive interfaces.

7. Cat will not participate in any internal routing protocol.

Cat will be running BGP as its only routing protocol.

8. Configure route redistribution between EIGRP and IS-IS. Use route-maps and access-lists to prevent route feedback.

Redistribution among routing protocols is not something to be done likely in production networks, and it is generally considered wise to use filters to prevent route feedback. Here we are redistributing from EIGRP to IS-IS and vice-versa on two routers, and we have also had to disable split horizon under EIGRP for a serial interface on R1, so it would seem like a good idea to use filters. Even so, it is possible that in this scenario filters to prevent route feedback are not absolutely necessary, due to EIGRP's different administrative distance values for internal routes (90) and external routes (170). I did not test whether problems would occur in this scenario without any filters.

To prevent any possible route feedback, I created an access-list containing all the routes that belong to EIGRP. I then used route-maps with redistribution so that only these routes are redistributed into IS-IS from EIGRP, and these same routes are not permitted into EIGRP from IS-IS but all other routes are permitted.

Remember to specify a seed metric where necessary (either using "default-metric" or specifying the seed metric with the "redistribute" command). For example, if you configure redistribution from IS-IS to EIGRP on R2 and R3, but neglect to specify a seed metric, redistribution into EIGRP simply will not occur because **EIGRP requires a seed metric**. Now let us imagine a hypothetical case in which you specified a seed metric on R2 but not on R3. If you were to check R1's routing table, you would see redistributed routes to the IS-IS networks 10.1.5.0/24 and 10.0.4.0/24, but you would see only one route for each where you should see two equal cost routes. Testing using ICMP would not uncover this error.

9. Configure BGP within AS 12345 without using any route reflector. R1 will form its own sub-AS (65501), and R2 and R5 will form their own sub-AS (65525) and R3 and R4 will form their own sub-AS (65534).

Confederations look really complicated at first, but they are really not very complicated at all. Just remember that your sub-AS is your AS for configuration purposes, and that you need to specify your "real" AS number as the confederation ID, and also need to specify the other sub-ASes in the confederation.

The numbers you can choose for your sub-ASes come from the private range of 64512-65535. The number 65534 that I used here for R3 and R4 lies just barely within the range.

10. Configure BGP peering among AS 12345, AS 67 and AS 10.

This is very simple, except that you need to make sure that when you peer with AS 12345 from other autonomous systems, you use the real AS number, rather than the private AS number for the sub-AS.

You may have noticed that I have not included "no auto-summary" under "router bgp" on every router in this scenario. While I have a habit of disabling auto-summary wherever I can, it is worth mentioning that failing to do so does not necessarily cause problems. According to the CCO article "BGP:Frequently Asked Questions," in order for auto-summarization to occur when it is enabled, you need to have a BGP network statement with a classful network and classful mask, and advertise a local subnet within the classful network. For example, if you have "network 10.0.0.0 mask 255.0.0.0" and have a local network 10.0.1.0/24, and auto-summary is enabled, your local network which you advertised in BGP using "network 10.0.1.0 mask 255.255.255.0" will end up being advertised as 10.0.0.0/8. The requirement that you need to advertise the classful network before auto-summarization will occur would seem to limit the auto-summarization problem to the person who foolishly advertises an entire classful network mask as his own.

Another situation where auto-summary can occur would be where you are redistributing routes from an IGP into BGP. Dr. William Parkhurst in discussing BGP "auto-summary" focuses on this situation. *Cisco BGP-4 Command and Configuration Handbook*, Indianapolis, Indiana: Cisco Press; 2001.

The upshot of all this is that it is generally a good practice to disable automatic summarization in BGP, failing to do in a particular scenario may not cause any problems. Failing to disable synchronization in a network that is not carefully designed for synchronization is far more likely to cause problems.

You can see why BGP would be designed to summarize by default. Doing so keeps routing tables from growing too large. But consider what could happen in a production network if you own a small subset of a class B network, and your routes end up being advertised in BGP as if you owned the entire class B network. Traffic intended for other companies would go to your autonomous system instead, and would be black-holed.

11. Advertise in BGP all networks except R2 e0 s network and the networks used for peering.

This means that you will end up with something less than any-to-any connectivity, and you will need to use extended ping for testing connectivity between, say, R7 lo0 and R5 e0.

12. R1 wants to use the ethernet link to AS 10 for user traffic. Configure BGP on R6 so that all destinations advertised from AS 12345 to AS 67 on the serial link will seem to AS 67 to be two Autonomous Systems further away than they really are.

I prepended 12345 twice to all prefixes received from R6 on the serial link. Here is R6 s BGP table, showing AS 12345 prepended twice in front of itself:

```
r6#sh ip bgp
BGP table version is 27, local router ID is 192.168.16.2
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
*> 10.1.3.0/24      10.1.106.1        0 10 12345 i
*                   192.168.16.1      0 12345 12345 12345
i

*> 10.1.4.0/24      10.1.106.1        0 10 12345 i
*                   192.168.16.1      0 12345 12345 12345
i

*> 10.1.5.0/24      10.1.106.1        0 10 12345 i
*                   192.168.16.1      0 12345 12345 12345
i

*>i192.168.70.0     192.168.1.1       0 100 0 i
r6#
```


13. Configure a BGP attribute on R1 so that it will always prefer the ethernet link to AS 10 for traffic destined for AS 10 or AS 67.

I set the weight on incoming BGP prefixes from Cat to a very high value. Note that setting weight does not even require a route-map. R1's BGP table shows the weight value for the only prefix involved:

```
r1#sh ip bgp
BGP table version is 27, local router ID is 192.168.16.1
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
*> 10.1.3.0/24      172.16.30.3         0      100      0 (65534) i
*> 10.1.4.0/24      172.16.30.25        0      100      0 (65525 65534) i
*                   172.16.30.22        0      100      0 (65534) i
* 10.1.5.0/24       172.16.30.18        0      100      0 (65525) i
*>                   172.16.30.26        0      100      0 (65534 65525) i
*> 192.168.70.0    10.1.101.1         65535 10 67 i
*                   192.168.16.2        0      100      0 67 i
r1#
```

The prefix with the higher weight is preferred over the prefix with the shorter AS-path.

14. Permit AS 67 to carry traffic between AS 10 and AS 12345 in the unlikely event that the ethernet link between R1 and Cat were to fail but the ethernet link between R6 and Cat were to remain up.

This should happen by default.

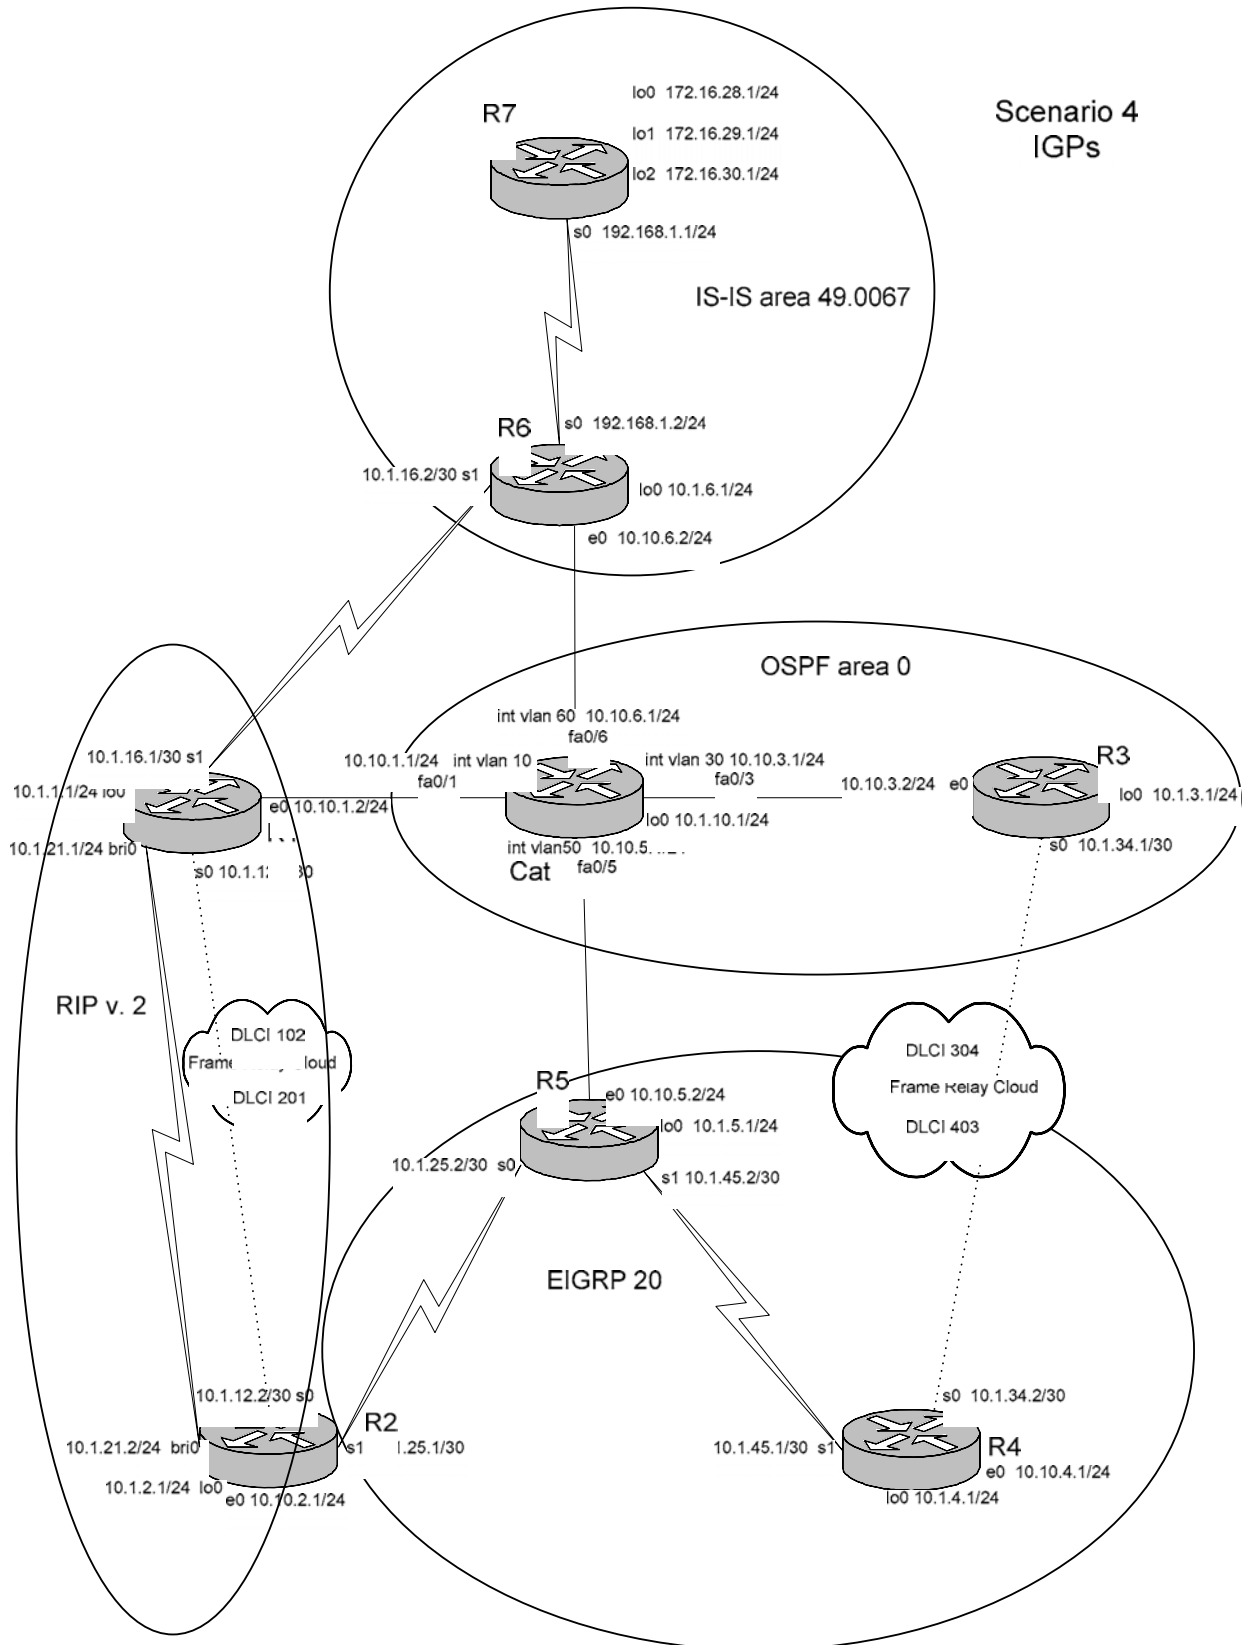
Remember that when testing connectivity in a scenario like this, where only a few selected routes are advertised in BGP, you will need to use extended ping and specify source and destination. I did a spot test using extended ping between two interfaces far removed from each other, but whose networks are advertised in BGP:

```
r7#ping
Protocol [ip]:
Target IP address: 10.1.5.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 192.168.70.1
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 5, 100-byte ICMP Echos to 10.1.5.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/16/28 ms
r7#
```

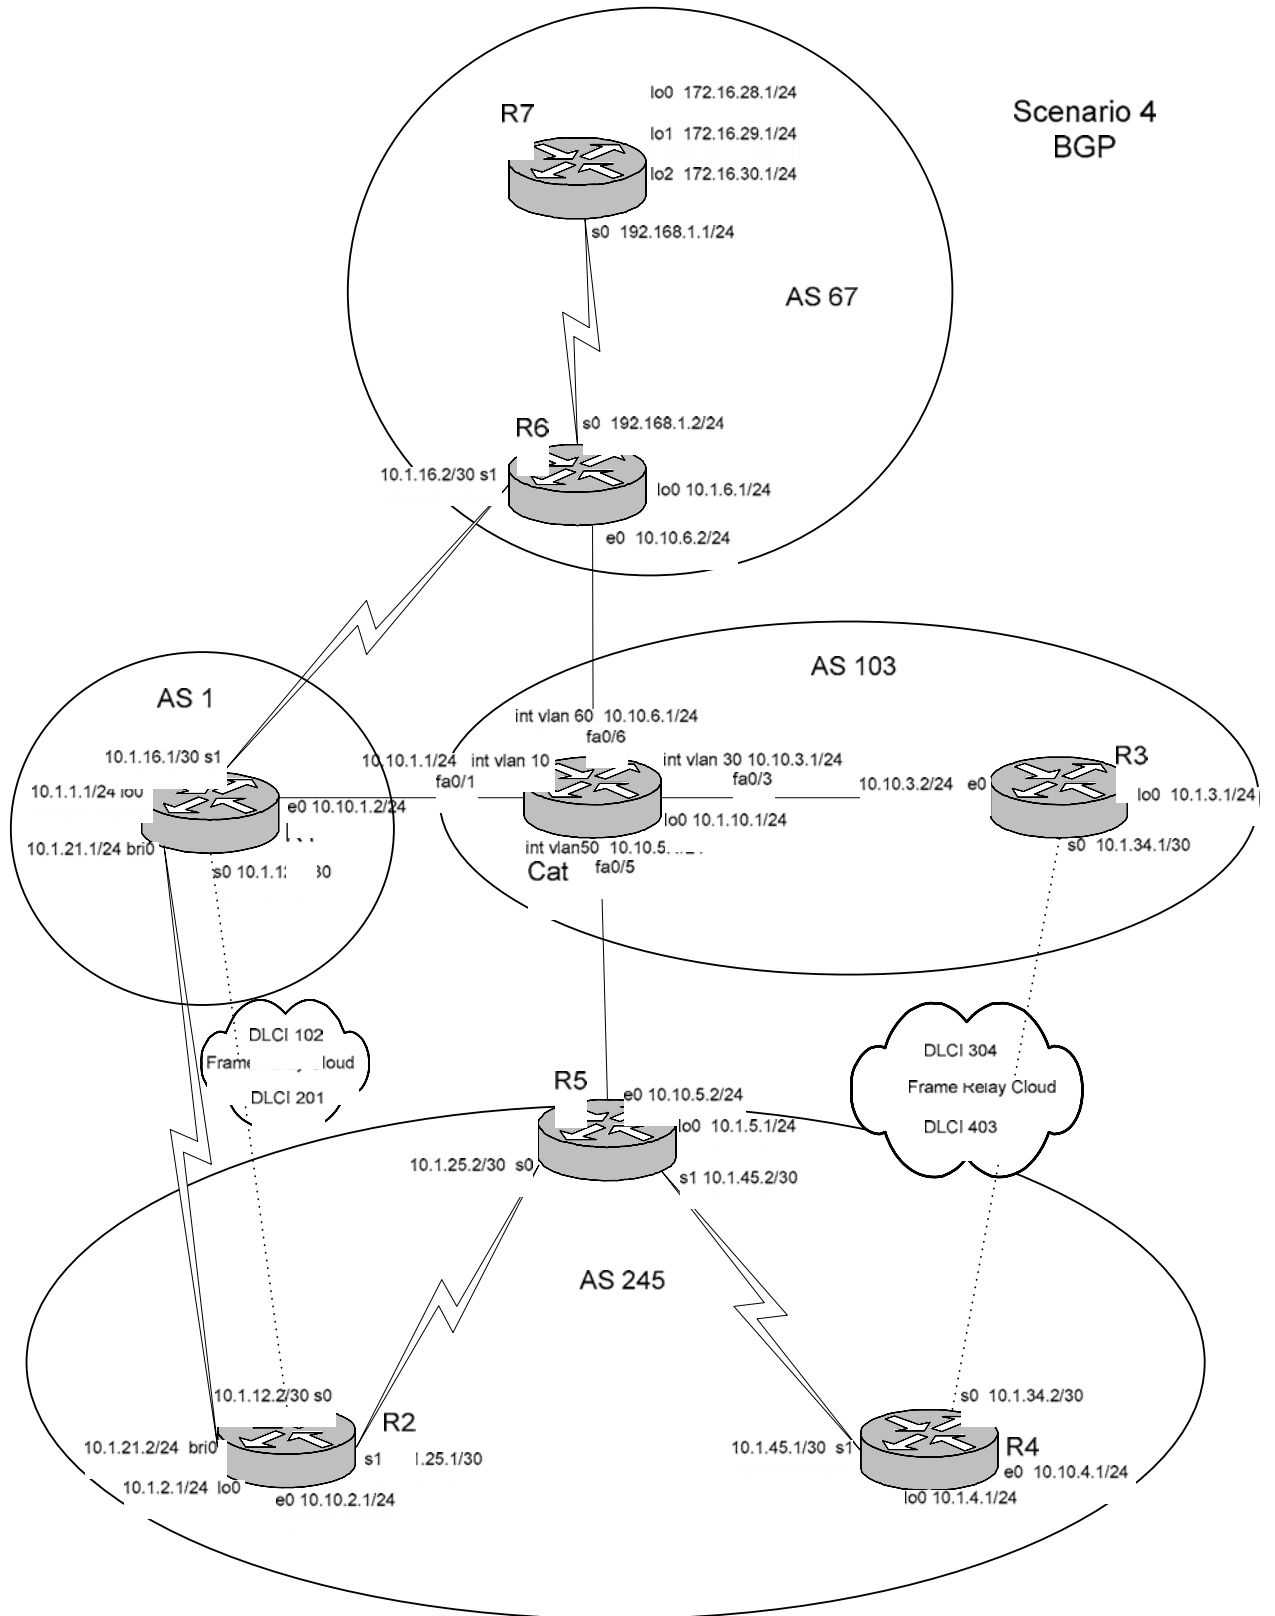
Chapter Four

Scenario Four

Scenario 4 IGPs



Scenario 4
BGP



Tasks

1. Configure Cat interfaces and VLANs and the router ethernet interfaces as shown in the table and the diagram.

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

2. Configure the frame relay links as shown in the diagram.
3. Configure the serial links as shown in the diagram. Use PPP as the encapsulation on each serial link.
4. Configure OSPF on R3 and Cat. Insure that all interfaces are in OSPF.
5. Configure IS-IS on R6 and R7. All interfaces will be in IS-IS, but interfaces which cannot form an IS-IS adjacency should be passive.
6. Configure EIGRP on R2, R5, and R4. R5 e0, R5 lo0, R4 lo0, R4 s0 and R4 e0 will be passive interfaces in EIGRP.
7. Configure MD5 route authentication in the EIGRP domain.
8. Configure RIP as shown in the diagram. Configure ISDN on R1 and R2 with a floating static route on each router pointing to the loopback0 network of the other router. Do not redistribute any routes among internal routing protocols. Make interfaces passive where appropriate.
9. Configure RIP MD5 route authentication.

10. Configure BGP on all routers and have each AS peer with every other AS with which it is directly connected. For EBGP peering between R1 and R2, use the loopbacks.

11. Advertise in BGP R2 e0, R3 lo0, R4 e0 and lo0, R5 lo0, R6 lo0, R7 lo0 and lo1 and lo2, and Cat lo0.

12. Configure BGP authentication between AS 67 and AS 103.

13. Configure BGP in AS 245 so that AS 245 will always prefer to send outbound traffic on the R5-Cat ethernet link.

14. Set the weight attribute on R6 so that it will always prefer sending outbound traffic on the R6-Cat ethernet link.

15. Configure R5 lo0 as a multicast source for testing purposes. Configure multicast using dense-mode PIM and insure that R4 e0 and R2 e0 can receive multicast traffic sourced from R5 lo0. Use the multicast group address 239.0.0.1.

16. Enable IP accounting on R6 s0 to keep track of all IP traffic (outbound on that interface), regardless of source or destination address. On R6 e0, enable accounting to show the source MAC-address of incoming frames.

Configuration Scripts

R1

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

Current configuration : 2285 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
!
key chain ripkey
  key 1
    key-string cisco
!
!
!
!
interface Loopback0
  ip address 10.1.1.1 255.255.255.0
!
interface Ethernet0
  ip address 10.10.1.2 255.255.255.0
!
interface Serial0
  ip address 10.1.12.1 255.255.255.252
  ip rip authentication mode md5
  ip rip authentication key-chain ripkey
  encapsulation frame-relay
  ip split-horizon
  frame-relay map ip 10.1.12.1 102
  frame-relay map ip 10.1.12.2 102 broadcast
  no frame-relay inverse-arp
!
interface Serial1
  ip address 10.1.16.1 255.255.255.252
  encapsulation ppp
```

```

!
interface BRI0
 ip address 10.1.21.1 255.255.255.0
 encapsulation ppp
 dialer idle-timeout 240
 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
 passive-interface default
 no passive-interface Serial0
 network 10.0.0.0
 no auto-summary
!
router bgp 1
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.1.2.1 remote-as 245
 neighbor 10.1.2.1 ebgp-multihop 255
 neighbor 10.1.2.1 update-source Loopback0
 neighbor 10.1.16.2 remote-as 67
 neighbor 10.10.1.1 remote-as 103
 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
!
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 : 2723 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
!
key chain ripkey
  key 1
    key-string cisco
key chain eigrpkey
  key 1
    key-string hongkong
!
!
!
!
interface Loopback0
  ip address 10.1.2.1 255.255.255.0
!
interface Ethernet0
  ip address 10.10.2.1 255.255.255.0
  ip pim dense-mode
  ip igmp join-group 239.0.0.1
!
interface Serial0
  ip address 10.1.12.2 255.255.255.252
  ip rip authentication mode md5
  ip rip authentication key-chain ripkey
  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
  no frame-relay inverse-arp
!
interface Serial1
  ip address 10.1.25.1 255.255.255.252
  ip authentication mode eigrp 20 md5
  ip authentication key-chain eigrp 20 eigrpkey
  ip pim dense-mode
  encapsulation ppp
  clockrate 1300000
```

```

!
interface BRI0
 ip address 10.1.21.2 255.255.255.0
 encapsulation ppp
 dialer map ip 10.1.21.1 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
 ppp authentication chap
!
router eigrp 20
 network 10.1.25.0 0.0.0.3
 no auto-summary
 no eigrp log-neighbor-changes
!
router rip
 version 2
 passive-interface default
 no passive-interface Serial0
 network 10.0.0.0
 no auto-summary
!
router bgp 245
 no synchronization
 bgp log-neighbor-changes
 network 10.10.2.0 mask 255.255.255.0
 neighbor 10.1.1.1 remote-as 1
 neighbor 10.1.1.1 ebgp-multihop 255
 neighbor 10.1.1.1 update-source Loopback0
 neighbor 10.1.1.1 route-map setas in
 neighbor 10.1.25.2 remote-as 245
 neighbor 10.1.25.2 next-hop-self
 no auto-summary
!
ip kerberos source-interface any
ip classless
ip route 10.1.1.0 255.255.255.0 10.1.21.1 240
ip http server
!
dialer-list 1 protocol ip permit
route-map setas permit 10
 set as-path prepend 1 1
!
!
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 : 1567 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
 ip ospf network point-to-point
!
interface Ethernet0
 ip address 10.10.3.2 255.255.255.0
!
interface Serial0
 ip address 10.1.34.1 255.255.255.252
 encapsulation frame-relay
 frame-relay map ip 10.1.34.1 304
 frame-relay map ip 10.1.34.2 304 broadcast
 no frame-relay inverse-arp
!
interface Serial1
 no ip address
 shutdown
!
router ospf 64
 log-adjacency-changes
 network 0.0.0.0 255.255.255.255 area 0
!
router bgp 103
 no synchronization
 bgp log-neighbor-changes
 network 10.1.3.0 mask 255.255.255.0
 neighbor 10.1.34.2 remote-as 245
```

```

neighbor 10.10.3.1 remote-as 103
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 : 2079 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
!
key chain eigrpkey
  key 1
    key-string hongkong
!
!
!
!
interface Loopback0
  ip address 10.1.4.1 255.255.255.0
!
interface Ethernet0
  ip address 10.10.4.1 255.255.255.0
  ip pim dense-mode

```

```
 ip igmp join-group 239.0.0.1
 !
interface Serial0
 ip address 10.1.34.2 255.255.255.252
 encapsulation frame-relay
 frame-relay map ip 10.1.34.1 403 broadcast
 frame-relay map ip 10.1.34.2 403
 no frame-relay inverse-arp
 !
interface Serial1
 ip address 10.1.45.1 255.255.255.252
 ip authentication mode eigrp 20 md5
 ip authentication key-chain eigrp 20 eigrpkey
 ip pim dense-mode
 encapsulation ppp
 clockrate 1300000
 !
interface TokenRing0
 no ip address
 shutdown
 !
router eigrp 20
 passive-interface default
 no passive-interface Serial1
 network 10.1.4.0 0.0.0.255
 network 10.1.34.0 0.0.0.3
 network 10.1.45.0 0.0.0.3
 network 10.10.4.0 0.0.0.255
 no auto-summary
 no eigrp log-neighbor-changes
 !
router bgp 245
 no synchronization
 bgp log-neighbor-changes
 network 10.1.4.0 mask 255.255.255.0
 network 10.10.4.0 mask 255.255.255.0
 neighbor 10.1.34.1 remote-as 103
 neighbor 10.1.45.2 remote-as 245
 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

r4#
```

R5

```

r5#sh runn
Building configuration...

Current configuration : 2136 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
!
key chain eigrpkey
  key 1
    key-string hongkong
!
!
!
!
interface Loopback0
  ip address 10.1.5.1 255.255.255.0
!
interface Ethernet0
  ip address 10.10.5.2 255.255.255.0
  ip pim dense-mode
!
interface Serial0
  ip address 10.1.25.2 255.255.255.252
  ip authentication mode eigrp 20 md5
  ip authentication key-chain eigrp 20 eigrpkey
  ip pim dense-mode
  encapsulation ppp
!
interface Serial1
  ip address 10.1.45.2 255.255.255.252
  ip authentication mode eigrp 20 md5
  ip authentication key-chain eigrp 20 eigrpkey
  ip pim dense-mode
  encapsulation ppp
!
router eigrp 20
  passive-interface default
  no passive-interface Serial0
  no passive-interface Serial1

```

```

network 10.1.5.0 0.0.0.255
network 10.1.25.0 0.0.0.3
network 10.1.45.0 0.0.0.3
network 10.10.5.0 0.0.0.255
no auto-summary
no eigrp log-neighbor-changes
!
router bgp 245
no synchronization
bgp log-neighbor-changes
network 10.1.5.0 mask 255.255.255.0
neighbor 10.1.25.1 remote-as 245
neighbor 10.1.25.1 route-reflector-client
neighbor 10.1.45.1 remote-as 245
neighbor 10.1.45.1 route-reflector-client
neighbor 10.10.5.1 remote-as 103
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

r5#

```

R6

```

r6#sh runn
Building configuration...

Current configuration : 1743 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 10.10.6.2 255.255.255.0  
  ip accounting mac-address input  
!  
interface Ethernet1  
  no ip address  
  shutdown  
!  
interface Serial0  
  ip address 192.168.1.2 255.255.255.0  
  ip accounting output-packets  
  ip router isis  
  encapsulation ppp  
  fair-queue  
  clockrate 1300000  
!  
interface Serial1  
  ip address 10.1.16.2 255.255.255.252  
  encapsulation ppp  
  clockrate 1300000  
!  
router isis  
  passive-interface default  
  no passive-interface Serial0  
  net 49.0067.6666.6666.6666.00  
!  
router bgp 67  
  no synchronization  
  bgp log-neighbor-changes  
  network 10.1.6.0 mask 255.255.255.0  
  neighbor 10.1.16.1 remote-as 1  
  neighbor 10.10.6.1 remote-as 103  
  neighbor 10.10.6.1 password cisco  
  neighbor 10.10.6.1 weight 65535  
  neighbor 192.168.1.1 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

r6#

```

R7

```

r7#sh runn
Building configuration...

Current configuration : 1674 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 172.16.28.1 255.255.255.0
!
interface Loopback1
 ip address 172.16.29.1 255.255.255.0
!
interface Loopback2
 ip address 172.16.30.1 255.255.255.0
!
interface Ethernet0
 no ip address
 shutdown
!
interface Ethernet1
 no ip address
 shutdown
!
interface Serial0
 ip address 192.168.1.1 255.255.255.0
 ip router isis
 encapsulation ppp
 no fair-queue

```

```

!
interface Serial11
  no ip address
  shutdown
!
router isis
  passive-interface default
  no passive-interface Serial0
  net 49.0067.7777.7777.7777.00
!
router bgp 67
  no synchronization
  bgp log-neighbor-changes
  network 172.16.28.0 mask 255.255.255.0
  network 172.16.29.0 mask 255.255.255.0
  network 172.16.30.0 mask 255.255.255.0
  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#sh runn
Building configuration...

Current configuration : 2958 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 extend system-id  
!  
!  
!  
interface Loopback0  
  ip address 10.1.10.1 255.255.255.0  
  ip ospf network point-to-point  
!  
interface FastEthernet0/1  
  switchport access vlan 10  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/2  
  switchport access vlan 20  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/3  
  switchport access vlan 30  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/4  
  switchport access vlan 40  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/5  
  switchport access vlan 50  
  switchport mode access  
  no ip address  
!  
interface FastEthernet0/6  
  switchport access vlan 60  
  switchport mode access  
  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  
  no ip address  
!  
interface FastEthernet0/12  
  no ip address  
!
```

```
interface FastEthernet0/13
  no ip address
!
interface FastEthernet0/14
  no ip address
!
interface FastEthernet0/15
  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
!
interface Vlan10
  ip address 10.10.1.1 255.255.255.0
!
interface Vlan30
  ip address 10.10.3.1 255.255.255.0
!
interface Vlan50
  ip address 10.10.5.1 255.255.255.0
!
interface Vlan60
  ip address 10.10.6.1 255.255.255.0
```

```
!  
router ospf 64  
  log-adjacency-changes  
  network 0.0.0.0 255.255.255.255 area 0  
!  
router bgp 103  
  no synchronization  
  bgp log-neighbor-changes  
  network 10.1.10.0 mask 255.255.255.0  
  neighbor 10.10.1.2 remote-as 1  
  neighbor 10.10.3.2 remote-as 103  
  neighbor 10.10.5.2 remote-as 245  
  neighbor 10.10.6.2 remote-as 67  
  neighbor 10.10.6.2 password cisco  
  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#
```

Routing Tables

R1

```
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

    172.16.0.0/24 is subnetted, 3 subnets
B       172.16.28.0 [20/0] via 10.1.16.2, 00:14:24
B       172.16.29.0 [20/0] via 10.1.16.2, 00:14:24
B       172.16.30.0 [20/0] via 10.1.16.2, 00:14:24
    10.0.0.0/8 is variably subnetted, 16 subnets, 3 masks
B       10.1.10.0/24 [20/0] via 10.10.1.1, 00:14:24
C       10.10.1.0/24 is directly connected, Ethernet0
B       10.10.2.0/24 [20/0] via 10.1.2.1, 00:14:24
B       10.10.4.0/24 [20/0] via 10.1.2.1, 00:14:24
C       10.1.12.0/30 is directly connected, Serial0
B       10.1.3.0/24 [20/0] via 10.10.1.1, 00:14:25
R       10.1.2.0/24 [120/1] via 10.1.12.2, 00:00:24, Serial0
C       10.1.1.0/24 is directly connected, Loopback0
B       10.1.6.0/24 [20/0] via 10.1.16.2, 00:14:25
B       10.1.5.0/24 [20/0] via 10.1.2.1, 00:14:25
B       10.1.4.0/24 [20/0] via 10.1.2.1, 00:14:26
R       10.1.25.2/32 [120/1] via 10.1.12.2, 00:00:25, Serial0
R       10.1.25.0/30 [120/1] via 10.1.12.2, 00:00:25, Serial0
C       10.1.16.2/32 is directly connected, Serial1
C       10.1.16.0/30 is directly connected, Serial1
C       10.1.21.0/24 is directly connected, BRI0
r1#
```

R2

```
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, 3 subnets
B    172.16.28.0 [200/0] via 10.10.5.1, 01:09:09
B    172.16.29.0 [200/0] via 10.10.5.1, 01:09:09
B    172.16.30.0 [200/0] via 10.10.5.1, 01:09:09
10.0.0.0/8 is variably subnetted, 20 subnets, 3 masks
B    10.1.10.0/24 [200/0] via 10.10.5.1, 01:09:13
R    10.10.1.0/24 [120/1] via 10.1.12.1, 00:00:09, Serial0
C    10.10.2.0/24 is directly connected, Ethernet0
D    10.10.4.0/24 [90/2707456] via 10.1.25.2, 01:13:48, Serial1
D    10.10.5.0/24 [90/2195456] via 10.1.25.2, 01:13:49, Serial1
C    10.1.12.0/30 is directly connected, Serial0
B    10.1.3.0/24 [200/0] via 10.10.5.1, 01:09:10
C    10.1.2.0/24 is directly connected, Loopback0
R    10.1.1.0/24 [120/1] via 10.1.12.1, 00:00:10, Serial0
B    10.1.6.0/24 [200/0] via 10.10.5.1, 01:09:10
D    10.1.5.0/24 [90/2297856] via 10.1.25.2, 01:13:49, Serial1
D    10.1.4.0/24 [90/2809856] via 10.1.25.2, 01:13:49, Serial1
C    10.1.25.2/32 is directly connected, Serial1
C    10.1.25.0/30 is directly connected, Serial1
R    10.1.16.2/32 [120/1] via 10.1.12.1, 00:00:10, Serial0
R    10.1.16.0/30 [120/1] via 10.1.12.1, 00:00:10, Serial0
C    10.1.21.0/24 is directly connected, BRI0
D    10.1.45.0/30 [90/2681856] via 10.1.25.2, 01:13:49, Serial1
D    10.1.45.1/32 [90/2681856] via 10.1.25.2, 01:13:49, Serial1
D    10.1.34.0/30 [90/3193856] via 10.1.25.2, 01:13:49, Serial1
r2#

```

R3

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

```

172.16.0.0/24 is subnetted, 3 subnets
B    172.16.28.0 [200/0] via 10.10.6.2, 23:13:40
B    172.16.29.0 [200/0] via 10.10.6.2, 23:13:40
B    172.16.30.0 [200/0] via 10.10.6.2, 23:13:41
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
O    10.10.1.0/24 [110/11] via 10.10.3.1, 23:14:29, Ethernet0
O    10.1.10.0/24 [110/11] via 10.10.3.1, 23:14:29, Ethernet0
B    10.10.2.0/24 [20/0] via 10.1.34.2, 01:40:29

```

```

C      10.10.3.0/24 is directly connected, Ethernet0
B      10.10.4.0/24 [20/0] via 10.1.34.2, 01:40:02
O      10.10.5.0/24 [110/11] via 10.10.3.1, 23:14:30, Ethernet0
O      10.10.6.0/24 [110/11] via 10.10.3.1, 23:14:30, Ethernet0
C      10.1.3.0/24 is directly connected, Loopback0
B      10.1.6.0/24 [200/0] via 10.10.6.2, 23:13:41
B      10.1.5.0/24 [20/0] via 10.1.34.2, 01:40:02
B      10.1.4.0/24 [20/0] via 10.1.34.2, 01:40:02
C      10.1.34.0/30 is directly connected, Serial0
r3#

```

R4

```

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 not set

```

       172.16.0.0/24 is subnetted, 3 subnets
B      172.16.28.0 [200/0] via 10.10.5.1, 01:41:32
B      172.16.29.0 [200/0] via 10.10.5.1, 01:41:32
B      172.16.30.0 [200/0] via 10.10.5.1, 01:41:32
10.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
B      10.1.10.0/24 [200/0] via 10.10.5.1, 01:41:32
B      10.10.2.0/24 [200/0] via 10.1.25.1, 01:41:27
C      10.10.4.0/24 is directly connected, Ethernet0
D      10.10.5.0/24 [90/2195456] via 10.1.45.2, 01:42:10, Serial11
B      10.1.3.0/24 [200/0] via 10.10.5.1, 01:41:33
B      10.1.6.0/24 [200/0] via 10.10.5.1, 01:41:33
D      10.1.5.0/24 [90/2297856] via 10.1.45.2, 23:15:26, Serial11
C      10.1.4.0/24 is directly connected, Loopback0
D      10.1.25.0/30 [90/2681856] via 10.1.45.2, 23:15:26, Serial11
D      10.1.25.1/32 [90/2681856] via 10.1.45.2, 23:15:26, Serial11
C      10.1.45.2/32 is directly connected, Serial11
C      10.1.45.0/30 is directly connected, Serial11
C      10.1.34.0/30 is directly connected, Serial0
r4#

```

R5

```

r5#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, 3 subnets
B       172.16.28.0 [20/0] via 10.10.5.1, 01:42:41
B       172.16.29.0 [20/0] via 10.10.5.1, 01:42:41
B       172.16.30.0 [20/0] via 10.10.5.1, 01:42:41
    10.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
B       10.1.10.0/24 [20/0] via 10.10.5.1, 01:42:41
B       10.10.2.0/24 [200/0] via 10.1.25.1, 01:42:16
D       10.10.4.0/24 [90/2195456] via 10.1.45.1, 23:16:11, Serial1
C       10.10.5.0/24 is directly connected, Ethernet0
B       10.1.3.0/24 [20/0] via 10.10.5.1, 01:42:42
B       10.1.6.0/24 [20/0] via 10.10.5.1, 01:42:42
C       10.1.5.0/24 is directly connected, Loopback0
D       10.1.4.0/24 [90/2297856] via 10.1.45.1, 23:16:12, Serial1
C       10.1.25.0/30 is directly connected, Serial0
C       10.1.25.1/32 is directly connected, Serial0
C       10.1.45.0/30 is directly connected, Serial1
C       10.1.45.1/32 is directly connected, Serial1
D       10.1.34.0/30 [90/2681856] via 10.1.45.1, 23:16:12, Serial1
r5#

```

R6

```

r6#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, 3 subnets
i L1    172.16.28.0 [115/10] via 192.168.1.1, Serial0
i L1    172.16.29.0 [115/10] via 192.168.1.1, Serial0
i L1    172.16.30.0 [115/10] via 192.168.1.1, Serial0
    10.0.0.0/8 is variably subnetted, 10 subnets, 3 masks
B       10.1.10.0/24 [20/0] via 10.10.6.1, 00:44:40
B       10.10.2.0/24 [20/0] via 10.10.6.1, 00:33:40
B       10.10.4.0/24 [20/0] via 10.10.6.1, 00:33:40
C       10.10.6.0/24 is directly connected, Ethernet0
B       10.1.3.0/24 [20/0] via 10.10.6.1, 00:44:41
C       10.1.6.0/24 is directly connected, Loopback0
B       10.1.5.0/24 [20/0] via 10.10.6.1, 00:33:41
B       10.1.4.0/24 [20/0] via 10.10.6.1, 00:33:41

```

```

C      10.1.16.1/32 is directly connected, Serial1
C      10.1.16.0/30 is directly connected, Serial1
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.1.1/32 is directly connected, Serial0
C      192.168.1.0/24 is directly connected, Serial0
r6#

```

R7

```

r7#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, 3 subnets
C      172.16.28.0 is directly connected, Loopback0
C      172.16.29.0 is directly connected, Loopback1
C      172.16.30.0 is directly connected, Loopback2
      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
B      10.1.10.0/24 [200/0] via 10.10.6.1, 00:45:42
B      10.10.2.0/24 [200/0] via 10.10.6.1, 00:34:42
B      10.10.4.0/24 [200/0] via 10.10.6.1, 00:34:42
i L1   10.10.6.0/24 [115/10] via 192.168.1.2, Serial0
B      10.1.3.0/24 [200/0] via 10.10.6.1, 00:45:39
i L1   10.1.6.0/24 [115/10] via 192.168.1.2, Serial0
B      10.1.5.0/24 [200/0] via 10.10.6.1, 00:34:38
B      10.1.4.0/24 [200/0] via 10.10.6.1, 00:34:43
i L1   10.1.16.0/30 [115/10] via 192.168.1.2, Serial0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.1.0/24 is directly connected, Serial0
C      192.168.1.2/32 is directly connected, Serial0
r7#

```

Cat

```

cat#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, 3 subnets
B    172.16.28.0 [20/0] via 10.10.6.2, 00:45:47
B    172.16.29.0 [20/0] via 10.10.6.2, 00:45:47
B    172.16.30.0 [20/0] via 10.10.6.2, 00:45:47
10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
C    10.10.1.0/24 is directly connected, Vlan10
C    10.1.10.0/24 is directly connected, Loopback0
B    10.10.2.0/24 [20/0] via 10.10.5.2, 00:45:36
C    10.10.3.0/24 is directly connected, Vlan30
B    10.10.4.0/24 [20/0] via 10.10.5.2, 00:45:36
C    10.10.5.0/24 is directly connected, Vlan50
C    10.10.6.0/24 is directly connected, Vlan60
O    10.1.3.0/24 [110/2] via 10.10.3.2, 00:46:31, Vlan30
B    10.1.6.0/24 [20/0] via 10.10.6.2, 00:45:48
B    10.1.5.0/24 [20/0] via 10.10.5.2, 00:45:36
B    10.1.4.0/24 [20/0] via 10.10.5.2, 00:45:36
O    10.1.34.0/30 [110/65] via 10.10.3.2, 00:46:31, Vlan30
cat#
```

Explanation

1. Configure Cat interfaces and VLANs and the router ethernet interfaces as shown in the table and the diagram.

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

No explanation needed.

2. Configure the frame relay links as shown in the diagram.

You can see that there are no hub-and-spoke arrangements, and no subinterfaces. Therefore, a reliable way to configure frame relay will be to disable inverse-arp and use static frame relay maps for IP. There are no frame-relay links in the IS-IS domain, so you will not need to worry about frame-relay mapping for CLNS.

In addition, you should enable split horizon under R1 s0 and R2 s0, since split horizon is disabled by default on physical frame relay interfaces.

3. Configure the serial links as shown in the diagram. Use PPP as the encapsulation on each serial link.

PPP is easy to configure, compatible with equipment from many vendors, and makes available additional features like PPP authentication and multilink fragmentation and interleaving.

4. Configure OSPF on R3 and Cat. Insure that all interfaces are in OSPF.

You can put an interface into the routing protocol, or you can redistribute it into the routing protocol as a connected network. The problem with redistributing a connected

network into a routing protocol like OSPF or EIGRP is that it will be treated as an external route, which creates the potential for administrative distance issues.

All in all, I would say that it is simpler to go ahead and make an interface part of the routing protocol than to **"redistribute connected"** with a route-map. You need to be ready to do it both ways, and you need to be wary of subtle issues like the way IS-IS routes that are also "connected" routes do not automatically get redistributed from IS-IS to the target routing protocol (covered elsewhere in this volume). You can make an interface passive to prevent useless hellos from being generated, but this was not required in this task.

5. Configure IS-IS on R6 and R7. All interfaces will be in IS-IS, but interfaces which cannot form an IS-IS adjacency should be passive.

This is about as simple as an IS-IS task can be.

6. Configure EIGRP on R2, R5, and R4. R5 e0, R5 lo0, R4 lo0, R4 s0 and R4 e0 will be passive interfaces in EIGRP.

This is a simple EIGRP arrangement. There are no frame interfaces actively running EIGRP, so there are no split horizon issues.

It is a good idea to go ahead and disable automatic route summarization by typing "no auto-summary" under EIGRP. This goes for EIGRP, RIP version 2, and BGP. If you are smart enough to know the default behavior for auto-summarization for each of these routing protocols for whatever IOS revision you are using, you can insist on configuring only what has to be configured, but I am not so confident about these matters. I have experimented with leaving out "no auto-summary" where I believed that the default behavior would be not to summarize, and the output from "show ip protocols" indicated that auto-summarization remained enabled. Whether automatic summarization was actually occurring in every instance is another matter, but I do not like to take chances, so I like to manually disable automatic summarization.

7. Configure MD5 route authentication in the EIGRP domain.

Authentication is conceptually simple, but the syntax can be difficult to memorize and tedious to execute. See the explanation for Task 9, below.

8. Configure RIP as shown in the diagram. Configure ISDN on R1 and R2 with a floating static route on each router pointing to the loopback0 network of the other router. Do not redistribute any routes among internal routing protocols. Make interfaces passive where appropriate.

I designed this scenario so that the static route will not simply sit there unused in the absence of ping testing. BGP traffic between R1 and R2 will bring up the link. I set the idle-timeout so that the BGP session will usually stay up, even though the ISDN link drops.

9. Configure RIP MD5 route authentication.

The syntax for RIP authentication is straightforward, but tedious, and differs somewhat from the syntax for EIGRP authentication. For example, the first line under the interface is "ip rip authentication mode md5" versus "ip authentication mode eigrp 20 md5" for EIGRP. That is radically different syntax.

The only link involved is the serial link between R1 and R2.

When testing, remember that RIP routes should be "forever young." If you see a RIP route that is a minute and a half old, your authentication configuration may not be right.

10. Configure BGP on all routers and have each AS peer with every other AS with which it is directly connected. For EBGP peering between R1 and R2, use the loopbacks.

We are peering using the loopbacks in order to have some traffic that will automatically bring up the ISDN link between R1 and R2 in the event of a failure of the frame-relay link between them.

11. Advertise in BGP R2 e0, R3 lo0, R4 e0 and lo0, R5 lo0, R6 lo0, R7 lo0 and lo1 and lo2, and Cat lo0.

It has not been flagged as an issue, but AS 245 requires a route reflector. R5 should serve as the route reflector, because it is the BGP speaker in AS 245 that is directly connected to both R2 and R4. You could, alternatively, use a confederation in AS 245 and have a separate sub-AS for each BGP speaker in AS 245, but that is a slightly more complicated solution.

You will note that we are not advertising the R1 and R2 loopbacks in BGP. This is to prevent recursive routing issues. You can use the **backdoor** command on R1 to resolve a recursive routing issue with regard to R2 e0.

12. Configure BGP authentication between AS 67 and AS 103.

The syntax is extremely simple. Personally, I love the simplicity of BGP configuration syntax. In general, you nail everything down on a per-neighbor basis, so you have a lot of control. If you have a lot of routers to configure, you can use peer-groups.

13. Configure BGP in AS 245 so that AS 245 will always prefer to send outbound traffic on the R5-Cat ethernet link.

A simple way to accomplish this task is to use "as-path prepend" to lengthen the AS-path for all prefixes coming in from R1 to R2, and from R3 to R4. I chose to lengthen the AS-path by two. You could also use local-preference, but that is slightly more complicated, in my opinion. Here we can see the result on R2's BGP table:

```
r2#sib
```

BGP table version is 11, local router ID is 10.1.2.1
 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
* 10.1.3.0/24	10.1.1.1			0	1 1 1 103 i
*>i	10.10.5.1		100	0	103 i
*>i10.1.4.0/24	10.1.45.1	0	100	0	i
*>i10.1.5.0/24	10.1.25.2	0	100	0	i
* 10.1.6.0/24	10.1.1.1			0	1 1 1 67 i
*>i	10.10.5.1		100	0	103 67 i
* 10.1.10.0/24	10.1.1.1			0	1 1 1 103 i
*>i	10.10.5.1	0	100	0	103 i
*> 10.10.2.0/24	0.0.0.0	0		32768	i
*>i10.10.4.0/24	10.1.45.1	0	100	0	i
* 172.16.28.0/24	10.1.1.1			0	1 1 1 67 i
*>i	10.10.5.1		100	0	103 67 i
* 172.16.29.0/24	10.1.1.1			0	1 1 1 67 i
*>i	10.10.5.1		100	0	103 67 i
* 172.16.30.0/24	10.1.1.1			0	1 1 1 67 i
*>i	10.10.5.1		100	0	103 67 i

r2#

Of course, the prefixes with the lengthened AS-paths will not be preferred. However, they will be used in the event that the preferred prefix becomes unavailable.

There is a potential next-hop reachability problem for prefixes with the next-hop of 10.1.1.1, because R5 and R4 do not know how to reach 10.1.1.1. I configured next-hop-self on R2 so that it would advertise itself as the next-hop for all prefixes it advertises to R5.

14. Set the weight attribute on R6 so that it will always prefer sending outbound traffic on the R6-Cat ethernet link.

I set the weight to a high value (65535) for all prefixes received on the ethernet link. Here you can see the result in R6's BGP table:

```
r6#sh ip bgp
BGP table version is 18, local router ID is 10.1.6.1
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
* 10.1.3.0/24	10.1.16.1			0	1 103 i
*>	10.10.6.1			65535	103 i
* 10.1.4.0/24	10.1.16.1			0	1 245 i
*>	10.10.6.1			65535	103 245 i
* 10.1.5.0/24	10.1.16.1			0	1 245 i
*>	10.10.6.1			65535	103 245 i
*> 10.1.6.0/24	0.0.0.0	0		32768	i
* 10.1.10.0/24	10.1.16.1			0	1 103 i
*>	10.10.6.1		0	65535	103 i
* 10.10.2.0/24	10.1.16.1			0	1 245 i
*>	10.10.6.1			65535	103 245 i

```

* 10.10.4.0/24      10.1.16.1                0 1 245 i
*>                  10.10.6.1                65535 103 245 i
*>i172.16.28.0/24  192.168.1.1             0 100 0 i
*>i172.16.29.0/24  192.168.1.1             0 100 0 i
*>i172.16.30.0/24  192.168.1.1             0 100 0 i
r6#

```

15. Configure R5 lo0 as a multicast source for testing purposes. Configure multicast using dense-mode PIM and insure that R4 e0 and R2 e0 can receive multicast traffic sourced from R5 lo0. Use the multicast group address 239.0.0.1.

This is about as simple as a multicast task can be. Configuration of IP multicast is required on only three routers. If you wanted multicast traffic to be able to take a roundabout path from R5 to R4 or R2 in the event of a serial link failure in the direct path, then you would want to enable multicast routing on more routers and more interfaces. Remember to use extended ping for testing.

16. Enable IP accounting on R6 s0 to keep track of all IP traffic (outbound on that interface), regardless of source or destination address. On R6 e0, enable accounting to show the source MAC-address of incoming frames.

With "ip accounting" under the interface in question, the IOS will automatically specify "output-packets." Remember that accounting will not keep track of traffic generated by the router on which it is configured. So when you test, you should test from another router behind the router on which "ip accounting" is configured.

Always be aware of the direction in which traffic is moving, and what device is initiating a session, and what device is responding. Learning configuration syntax and understanding how routing protocols work is much of the battle, but if you do not gain a keen awareness of direction of traffic flows and location of various devices and networks relative to each other, you will be at a disadvantage. I say this as someone who took a long time to develop an adequate degree of this kind of awareness.

To test IP accounting, I started a very long series of pings from R5 lo0 to R7 lo2, and then went over to R6 and watched the number of packets increase quickly in a second or so:

```

r6#sh ip accounting
  Source          Destination          Packets          Bytes
  10.1.5.1         172.16.30.1         3823             382300

Accounting data age is 0
r6#sh ip accounting
  Source          Destination          Packets          Bytes
  10.1.5.1         172.16.30.1         4394             439400

Accounting data age is 0
r6#

```


You can clear the data collected by typing "**clear ip accounting.**" To test the accounting based on MAC address, I did a simple ping from Cat to R6 e0 (simple in that I did not need to specify source address or interface), and then went over to R6 and typed "**show interface mac**":

```
cat#ping 10.10.6.2
```

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

```
Sending 5, 100-byte ICMP Echos to 10.10.6.2, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

```
cat#
```

```
termserver>6
```

```
[Resuming connection 6 to r6 ... ]
```

```
r6#sh interface mac
```

```
Ethernet0
```

```
Input (511 free)
```

```
000b.bee6.3e80(237): 7 packets, 756 bytes, last: 1620ms ago
```

```
Total: 7 packets, 756 bytes
```

```
r6#
```

That MAC address is the MAC address of Cat's faste0/6 which connects to R6 e0. This is the only MAC address that will be the source of any frame that R6 e0 sees in this scenario.

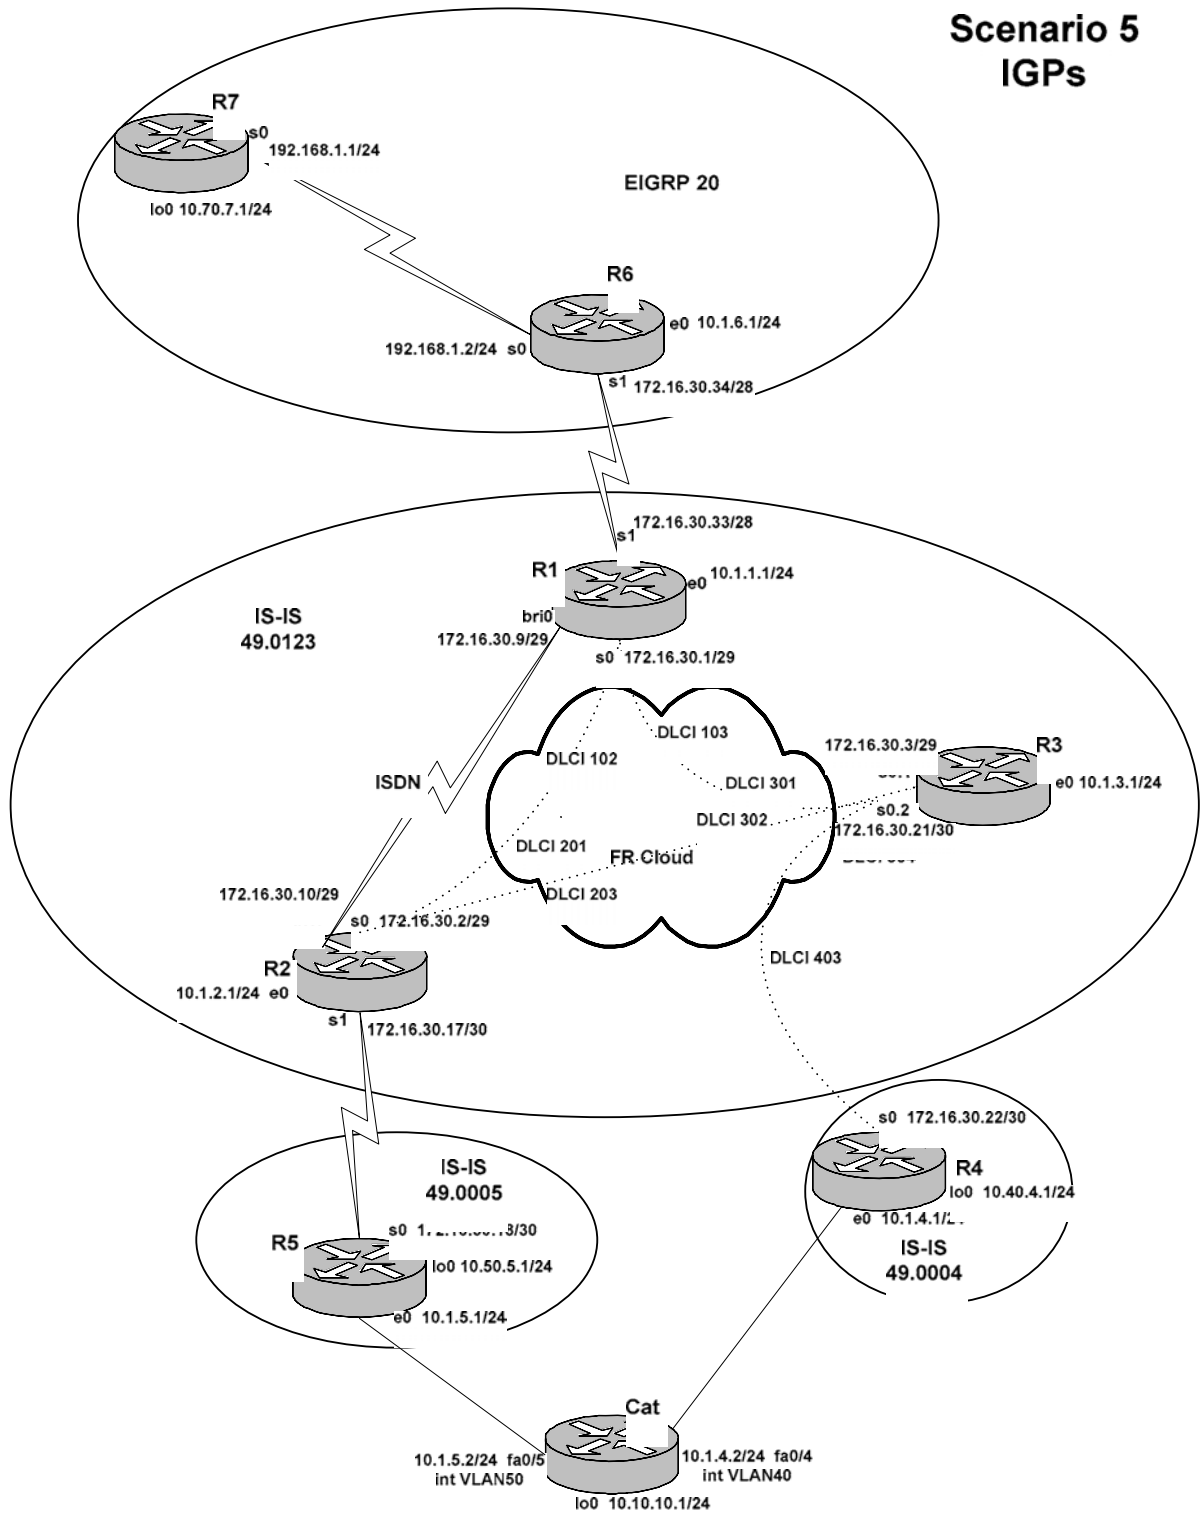
You may want to explore further the different kinds of IP accounting that can be configured under an interface. You can do IP accounting based on access-list violations, and based on IP precedence value.

Before configuring any sort of accounting on an interface in a production network, make sure you research the performance implications associated with IP accounting. In some cases, accounting may require that your router use a slower form of switching than it would normally use, while in other cases it does not. This may not matter in a lab setting, but it could matter a lot in a production network.

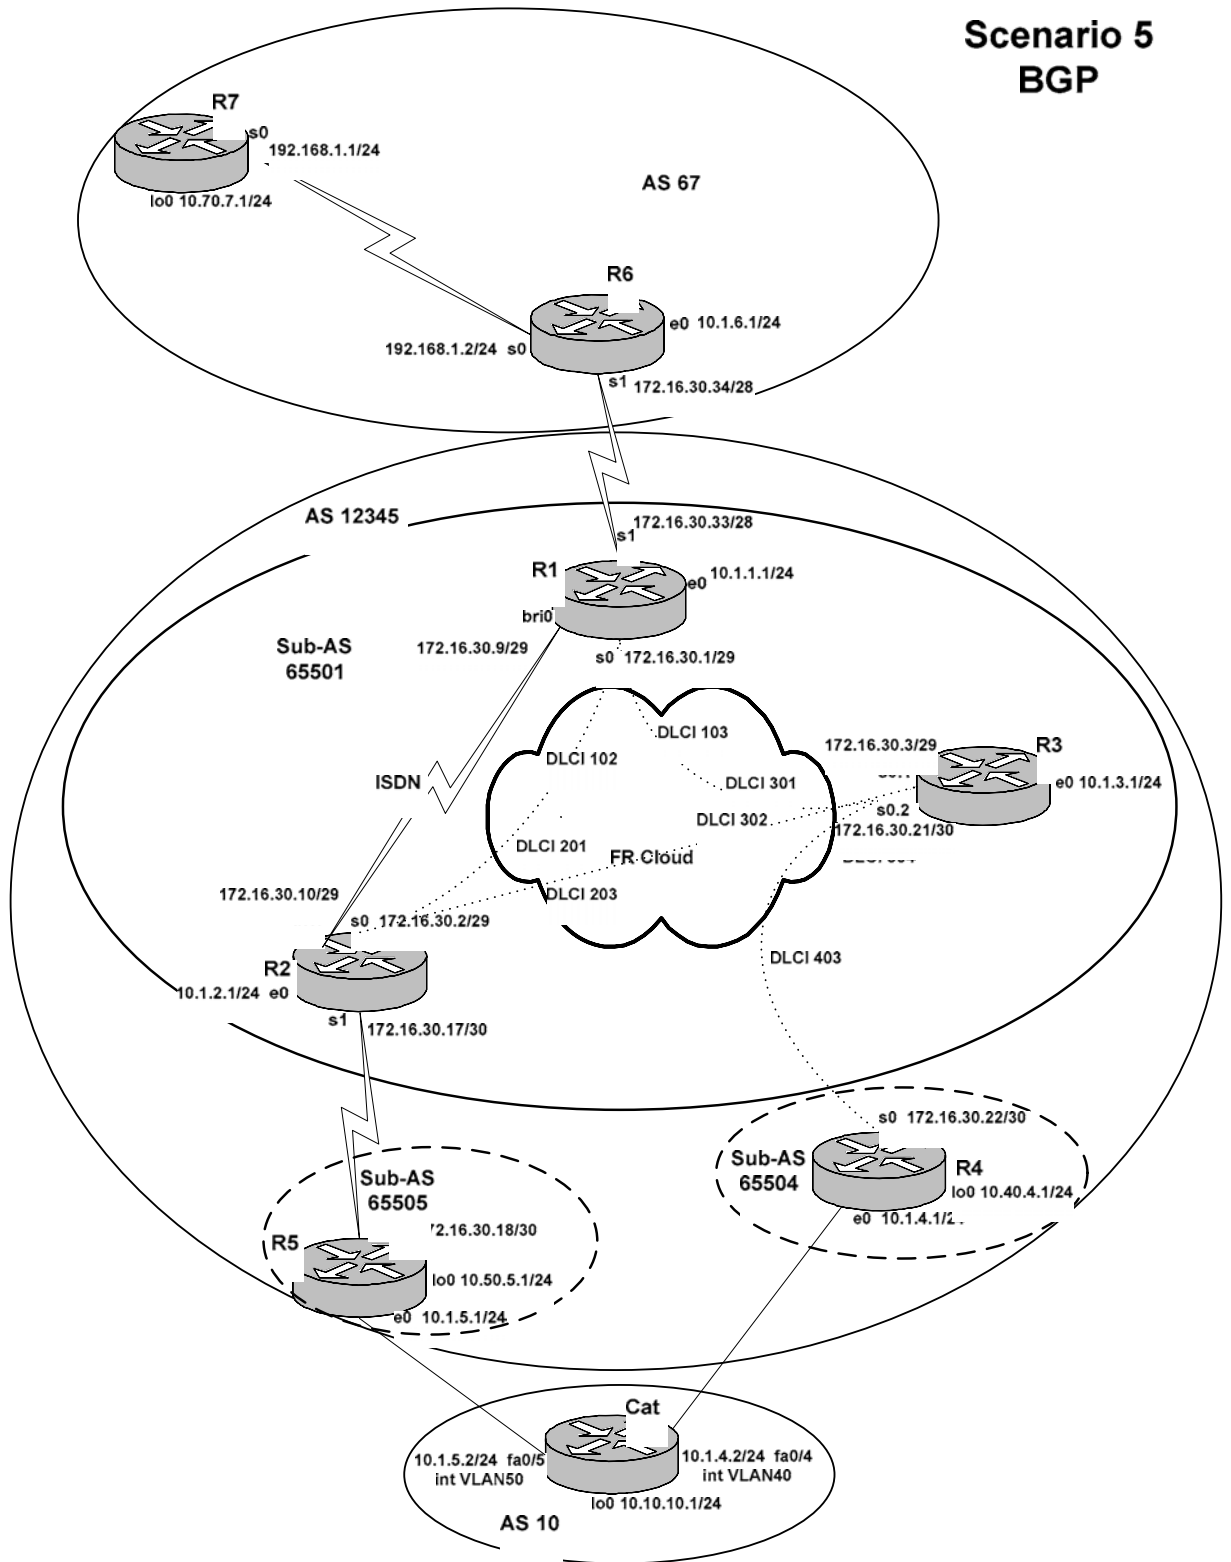
Chapter Five

Scenario Five

Scenario 5 IGPs



Scenario 5 BGP



Tasks

1. Configure the switch ports, VLANs, and router ethernet interfaces according to the diagram and table. You will need to create on Cat the Switched Virtual Interfaces (SVIs) VLAN40 and VLAN50.

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

2. Configure the frame relay links as shown in the diagram. Note that there is a full-mesh among R1, R2 and R3.

3. Configure the serial links as shown in the diagram using PPP encapsulation.

4. Configure IS-IS as shown in the diagram. Use passive interface on interfaces on which it does no good to send out hellos. Use L1, L2, or L1-L2 as you like.

5. Configure dialer watch on R1 so that R1 will dial R2 over the ISDN link when there is no longer a route to R2 e0 in R1's routing table. Configure ISDN so that bri0:2 will be used whenever bri0:1 is even minimally loaded.

6. Configure IS-IS interface authentication, area authentication, and domain authentication, using the passwords "interface," "area," and "domain," respectively.

7. Configure EIGRP on R6 and R7.

8. Enable IP routing on Cat, but do not run any internal routing protocol on it.

9. Do not configure any internal route redistribution.

10. Configure R1, R2, R3, R4, and R5 as BGP speakers in AS 12345. Do not peer from loopback to loopback in this scenario.

11. Configure R6 and R7 as BGP speakers in AS 67.

12. Configure Cat as a BGP speaker in AS 10.

13. Configure peering between AS 67 and AS 12345, and between AS 12345 and AS 10.

14. Advertise loopbacks in BGP, and Ethernets except for R2 e0. Do not peer from loopback to loopback.
15. Configure local-preference in AS 12345 so that the R5-Cat ethernet link will be preferred for any traffic destined to AS 10. We are not concerned about the path taken by traffic going in the opposite direction.

Configuration Scripts

R1

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

Current configuration : 2675 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
 ip router isis
 encapsulation frame-relay
 no arp frame-relay
 frame-relay map clns 103 broadcast
 frame-relay map clns 102 broadcast
 frame-relay map ip 172.16.30.1 102
 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
 isis password interface level-1
 isis password interface level-2
!
interface Serial1
 ip address 172.16.30.33 255.255.255.240
 encapsulation ppp
!
```

```

interface BRI0
 ip address 172.16.30.9 255.255.255.248
 ip router isis
 encapsulation ppp
 dialer map clns 49.0123.2222.2222.2222.00 name r2 broadcast 4082222222
 dialer map ip 10.1.2.0 name r2 broadcast 4082222222
 dialer map ip 172.16.30.10 name r2 broadcast 4082222222
 dialer load-threshold 5 either
 dialer watch-group 1
 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
 ppp multilink
 isis metric 30
 isis password interface level-1
 !
router isis
 passive-interface Ethernet0
 passive-interface Serial1
 net 49.0123.1111.1111.1111.00
 domain-password domain
 area-password area
 !
router bgp 65501
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 12345
 bgp confederation peers 65504 65505
 network 10.1.1.0 mask 255.255.255.0
 neighbor 172.16.30.2 remote-as 65501
 neighbor 172.16.30.3 remote-as 65501
 neighbor 172.16.30.34 remote-as 67
 no auto-summary
 !
ip kerberos source-interface any
ip classless
ip http server
 !
dialer watch-list 1 ip 10.1.2.0 255.255.255.0
dialer-list 1 protocol ip permit
 !
line con 0
 transport input none
line aux 0
line vty 0 4
 login
 !
end

r1#

```


R2

```

r2#sh runn
Building configuration...

Current configuration : 2579 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
 ip router isis
 encapsulation frame-relay
 no arp frame-relay
 frame-relay map clns 201 broadcast
 frame-relay map clns 203 broadcast
 frame-relay map ip 172.16.30.1 201 broadcast
 frame-relay map ip 172.16.30.2 201
 frame-relay map ip 172.16.30.3 203 broadcast
 no frame-relay inverse-arp
 isis password interface level-1
 isis password interface level-2
!
interface Serial1
 ip address 172.16.30.17 255.255.255.252
 ip router isis
 encapsulation ppp
 clockrate 1300000
 isis password interface level-2
!
interface BRI0
 ip address 172.16.30.10 255.255.255.248
 ip router isis

```

```

encapsulation ppp
dialer map clns 49.0123.1111.1111.1111.00 name r1 broadcast 4081111111
dialer map ip 172.16.30.9 name r1 broadcast 4081111111
dialer load-threshold 5 either
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
ppp authentication chap
ppp multilink
isis metric 30
isis password interface level-1
isis password interface level-2
!
router isis
  passive-interface Ethernet0
  net 49.0123.2222.2222.2222.00
  domain-password domain
  area-password area
!
router bgp 65501
  no synchronization
  bgp log-neighbor-changes
  bgp confederation identifier 12345
  bgp confederation peers 65504 65505
  neighbor 172.16.30.1 remote-as 65501
  neighbor 172.16.30.3 remote-as 65501
  neighbor 172.16.30.18 remote-as 65505
  no auto-summary
!
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#sh runn
Building configuration...

```

```
Current configuration : 2168 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
!
interface Serial0
 no ip address
 encapsulation frame-relay
 no arp frame-relay
 no frame-relay inverse-arp
!
interface Serial0.1 multipoint
 ip address 172.16.30.3 255.255.255.248
 ip router isis
 frame-relay map clns 302 broadcast
 frame-relay map clns 301 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 172.16.30.3 301
 isis password interface level-1
 isis password interface level-2
!
interface Serial0.2 multipoint
 ip address 172.16.30.21 255.255.255.252
 ip router isis
 frame-relay map clns 304 broadcast
 frame-relay map ip 172.16.30.21 304
 frame-relay map ip 172.16.30.22 304 broadcast
 isis password interface level-2
!
interface Serial1
 no ip address
 shutdown
!
router isis
 passive-interface Ethernet0
 net 49.0123.3333.3333.3333.00
```

```

domain-password domain
area-password area
!
router bgp 65501
  no synchronization
  bgp log-neighbor-changes
  bgp confederation identifier 12345
  bgp confederation peers 65504 65505
  network 10.1.3.0 mask 255.255.255.0
  neighbor 172.16.30.1 remote-as 65501
  neighbor 172.16.30.2 remote-as 65501
  neighbor 172.16.30.22 remote-as 65504
  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 : 1804 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 10.40.4.1 255.255.255.0
!
interface Ethernet0
 ip address 10.1.4.1 255.255.255.0
!
interface Serial0
 ip address 172.16.30.22 255.255.255.252
 ip router isis
 encapsulation frame-relay
 frame-relay map clns 403 broadcast
 frame-relay map ip 172.16.30.21 403 broadcast
 frame-relay map ip 172.16.30.22 403
 isis password interface level-2
!
interface Serial1
 no ip address
 shutdown
!
interface TokenRing0
 no ip address
 shutdown
!
router isis
 passive-interface Ethernet0
 passive-interface Loopback0
 net 49.0004.4444.4444.4444.00
 domain-password domain
 area-password area
!
router bgp 65504
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 12345
 bgp confederation peers 65501 65504 65505
 network 10.40.4.0 mask 255.255.255.0
 neighbor 10.1.4.2 remote-as 10
 neighbor 172.16.30.21 remote-as 65501
 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

r4#

```

R5

```

r5#sh runn
Building configuration...

Current configuration : 1710 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 10.50.5.1 255.255.255.0
!
interface Ethernet0
 ip address 10.1.5.1 255.255.255.0
!
interface Serial0
 ip address 172.16.30.18 255.255.255.252
 ip router isis
 encapsulation ppp
 isis password interface level-2
!
interface Serial1
 no ip address
 shutdown
!
router isis
 passive-interface Ethernet0
 passive-interface Loopback0
 net 49.0005.5555.5555.5555.00
 domain-password domain
 area-password area
!
router bgp 65505
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 12345
 bgp confederation peers 65501 65504
 network 10.50.5.0 mask 255.255.255.0
 neighbor 10.1.5.2 remote-as 10

```

```

neighbor 10.1.5.2 route-map setpref in
neighbor 172.16.30.17 remote-as 65501
no auto-summary
!
ip kerberos source-interface any
ip classless
ip http server
!
route-map setpref permit 10
  set local-preference 10000
!
!
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 : 1544 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 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
 clockrate 1300000
!
interface Serial1
 ip address 172.16.30.34 255.255.255.240
 encapsulation ppp
 clockrate 1300000
!
router eigrp 20
 network 10.1.6.0 0.0.0.255
 network 172.16.30.32 0.0.0.15
 network 192.168.1.0
 no auto-summary
 no eigrp log-neighbor-changes
!
router bgp 67
 no synchronization
 bgp log-neighbor-changes
 network 10.1.6.0 mask 255.255.255.0
 neighbor 172.16.30.33 remote-as 12345
 neighbor 192.168.1.1 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

r6#

```

R7

```

r7#sh runn
Building configuration...

Current configuration : 1480 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 10.70.7.1 255.255.255.0  
!  
interface Ethernet0  
 no ip address  
 shutdown  
!  
interface Ethernet1  
 no ip address  
 shutdown  
!  
interface Serial0  
 ip address 192.168.1.1 255.255.255.0  
 encapsulation ppp  
!  
interface Serial1  
 no ip address  
 shutdown  
!  
router eigrp 20  
 passive-interface Loopback0  
 network 10.70.7.0 0.0.0.255  
 network 192.168.1.0  
 no auto-summary  
 no eigrp log-neighbor-changes  
!  
router bgp 67  
 no synchronization  
 bgp log-neighbor-changes  
 network 10.70.7.0 mask 255.255.255.0  
 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#sh runn
Building configuration...

Current configuration : 2627 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 extend system-id
!
!
!
interface Loopback0
 ip address 10.10.10.1 255.255.255.0
!
interface FastEthernet0/1
 switchport access vlan 10
 switchport mode access
 no ip address
!
interface FastEthernet0/2
 switchport access vlan 20
 switchport mode access
 no ip address
!
interface FastEthernet0/3
 switchport access vlan 30
 switchport mode access
 no ip address
!
interface FastEthernet0/4
 switchport access vlan 40
 switchport mode access
 no ip address
!
interface FastEthernet0/5
 switchport access vlan 50
 switchport mode access
 no ip address
```

```
!  
interface FastEthernet0/6  
  switchport access vlan 60  
  switchport mode access  
  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  
  no ip address  
!  
interface FastEthernet0/12  
  no ip address  
!  
interface FastEthernet0/13  
  no ip address  
!  
interface FastEthernet0/14  
  no ip address  
!  
interface FastEthernet0/15  
  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
!
interface Vlan40
  ip address 10.1.4.2 255.255.255.0
!
interface Vlan50
  ip address 10.1.5.2 255.255.255.0
!
router bgp 10
  no synchronization
  bgp log-neighbor-changes
  network 10.10.10.0 mask 255.255.255.0
  neighbor 10.1.4.1 remote-as 12345
  neighbor 10.1.5.1 remote-as 12345
  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#
```

Routing Tables

R1

```
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, 8 subnets, 4 masks
C       172.16.30.34/32 is directly connected, Serial1
C       172.16.30.32/28 is directly connected, Serial1
C       172.16.30.0/29 is directly connected, Serial0
C       172.16.30.8/29 is directly connected, BRI0
i L2    172.16.30.18/32 [115/20] via 172.16.30.2, Serial0
i L1    172.16.30.16/30 [115/20] via 172.16.30.2, Serial0
i L2    172.16.30.17/32 [115/30] via 172.16.30.2, Serial0
i L1    172.16.30.20/30 [115/20] via 172.16.30.3, Serial0
10.0.0.0/24 is subnetted, 10 subnets
i L1    10.1.3.0 [115/10] via 172.16.30.3, Serial0
i L1    10.1.2.0 [115/10] via 172.16.30.2, Serial0
B       10.10.10.0 [200/0] via 10.1.5.2, 01:10:56
C       10.1.1.0 is directly connected, Ethernet0
B       10.1.6.0 [20/0] via 172.16.30.34, 3d02h
i L2    10.1.5.0 [115/20] via 172.16.30.2, Serial0
i L2    10.1.4.0 [115/20] via 172.16.30.3, Serial0
i L2    10.40.4.0 [115/20] via 172.16.30.3, Serial0
i L2    10.50.5.0 [115/20] via 172.16.30.2, Serial0
B       10.70.7.0 [20/0] via 172.16.30.34, 3d02h
r1#
```

R2

```
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
```

```

172.16.0.0/16 is variably subnetted, 7 subnets, 4 masks
i L2 172.16.30.34/32 [115/10] via 172.16.30.1, Serial0
i L1 172.16.30.32/28 [115/10] via 172.16.30.1, Serial0
C 172.16.30.0/29 is directly connected, Serial0
C 172.16.30.8/29 is directly connected, BRI0
C 172.16.30.18/32 is directly connected, Serial1
C 172.16.30.16/30 is directly connected, Serial1
i L1 172.16.30.20/30 [115/20] via 172.16.30.3, Serial0
10.0.0.0/24 is subnetted, 10 subnets
i L1 10.1.3.0 [115/10] via 172.16.30.3, Serial0
C 10.1.2.0 is directly connected, Ethernet0
i L1 10.1.1.0 [115/10] via 172.16.30.1, Serial0
B 10.10.10.0 [200/0] via 10.1.5.2, 3d02h
B 10.1.6.0 [200/0] via 172.16.30.34, 01:12:07
i L2 10.1.5.0 [115/10] via 172.16.30.18, Serial1
i L2 10.1.4.0 [115/20] via 172.16.30.3, Serial0
i L2 10.40.4.0 [115/20] via 172.16.30.3, Serial0
i L2 10.50.5.0 [115/10] via 172.16.30.18, Serial1
B 10.70.7.0 [200/0] via 172.16.30.34, 01:12:07
r2#

```

R3

```

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, 8 subnets, 4 masks
i L2 172.16.30.34/32 [115/10] via 172.16.30.1, Serial0.1
i L1 172.16.30.32/28 [115/10] via 172.16.30.1, Serial0.1
C 172.16.30.0/29 is directly connected, Serial0.1
i L1 172.16.30.8/29 [115/40] via 172.16.30.2, Serial0.1
        [115/40] via 172.16.30.1, Serial0.1
i L2 172.16.30.18/32 [115/20] via 172.16.30.2, Serial0.1
i L1 172.16.30.16/30 [115/20] via 172.16.30.2, Serial0.1
i L2 172.16.30.17/32 [115/30] via 172.16.30.2, Serial0.1
C 172.16.30.20/30 is directly connected, Serial0.2
10.0.0.0/24 is subnetted, 10 subnets
C 10.1.3.0 is directly connected, Ethernet0
i L1 10.1.2.0 [115/10] via 172.16.30.2, Serial0.1
i L1 10.1.1.0 [115/10] via 172.16.30.1, Serial0.1
B 10.10.10.0 [200/0] via 10.1.5.2, 3d02h
B 10.1.6.0 [200/0] via 172.16.30.34, 01:11:28
i L2 10.1.5.0 [115/20] via 172.16.30.2, Serial0.1

```

```

i L2    10.1.4.0 [115/10] via 172.16.30.22, Serial0.2
i L2    10.40.4.0 [115/10] via 172.16.30.22, Serial0.2
i L2    10.50.5.0 [115/20] via 172.16.30.2, Serial0.1
B       10.70.7.0 [200/0] via 172.16.30.34, 01:11:28
r3#

```

R4

```

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 not set

```

        172.16.0.0/16 is variably subnetted, 8 subnets, 4 masks
i L2    172.16.30.34/32 [115/20] via 172.16.30.21, Serial0
i L2    172.16.30.32/28 [115/20] via 172.16.30.21, Serial0
i L2    172.16.30.0/29 [115/20] via 172.16.30.21, Serial0
i L2    172.16.30.8/29 [115/50] via 172.16.30.21, Serial0
i L2    172.16.30.18/32 [115/30] via 172.16.30.21, Serial0
i L2    172.16.30.16/30 [115/30] via 172.16.30.21, Serial0
i L2    172.16.30.17/32 [115/40] via 172.16.30.21, Serial0
C       172.16.30.20/30 is directly connected, Serial0
        10.0.0.0/24 is subnetted, 10 subnets
i L2    10.1.3.0 [115/10] via 172.16.30.21, Serial0
i L2    10.1.2.0 [115/20] via 172.16.30.21, Serial0
i L2    10.1.1.0 [115/20] via 172.16.30.21, Serial0
B       10.10.10.0 [200/0] via 10.1.5.2, 3d02h
B       10.1.6.0 [200/0] via 172.16.30.34, 01:11:56
i L2    10.1.5.0 [115/30] via 172.16.30.21, Serial0
C       10.1.4.0 is directly connected, Ethernet0
C       10.40.4.0 is directly connected, Loopback0
i L2    10.50.5.0 [115/30] via 172.16.30.21, Serial0
B       10.70.7.0 [200/0] via 172.16.30.34, 01:11:56
r4#

```

R5

```

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

```

```

    i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS
inter area
    P - periodic downloaded static route

```

Gateway of last resort is not set

```

    172.16.0.0/16 is variably subnetted, 7 subnets, 4 masks
i L2   172.16.30.34/32 [115/20] via 172.16.30.17, Serial0
i L2   172.16.30.32/28 [115/20] via 172.16.30.17, Serial0
i L2   172.16.30.0/29 [115/20] via 172.16.30.17, Serial0
i L2   172.16.30.8/29 [115/40] via 172.16.30.17, Serial0
C      172.16.30.16/30 is directly connected, Serial0
C      172.16.30.17/32 is directly connected, Serial0
i L2   172.16.30.20/30 [115/30] via 172.16.30.17, Serial0
10.0.0.0/24 is subnetted, 10 subnets
i L2   10.1.3.0 [115/20] via 172.16.30.17, Serial0
i L2   10.1.2.0 [115/10] via 172.16.30.17, Serial0
i L2   10.1.1.0 [115/20] via 172.16.30.17, Serial0
B      10.10.10.0 [20/0] via 10.1.5.2, 3d03h
B      10.1.6.0 [200/0] via 172.16.30.34, 01:13:30
C      10.1.5.0 is directly connected, Ethernet0
i L2   10.1.4.0 [115/30] via 172.16.30.17, Serial0
i L2   10.40.4.0 [115/30] via 172.16.30.17, Serial0
C      10.50.5.0 is directly connected, Loopback0
B      10.70.7.0 [200/0] via 172.16.30.34, 01:13:29
r5#

```

R6

```
r6#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, 2 subnets, 2 masks
C      172.16.30.32/28 is directly connected, Serial1
C      172.16.30.33/32 is directly connected, Serial1
10.0.0.0/24 is subnetted, 8 subnets
B      10.1.3.0 [20/0] via 172.16.30.33, 01:12:40
B      10.10.10.0 [20/0] via 172.16.30.33, 01:13:48
B      10.1.1.0 [20/0] via 172.16.30.33, 3d03h
C      10.1.6.0 is directly connected, Ethernet0
B      10.40.4.0 [20/0] via 172.16.30.33, 01:12:42
B      10.50.5.0 [20/0] via 172.16.30.33, 01:13:49
D      10.70.7.0 [90/2297856] via 192.168.1.1, 4d00h, Serial0
192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.1.1/32 is directly connected, Serial0

```



```
C      192.168.1.0/24 is directly connected, Serial0
r6#
```

R7

```
r7#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, 2 subnets, 2 masks
D      172.16.30.32/28 [90/2681856] via 192.168.1.2, 01:37:18, Serial0
D      172.16.30.33/32 [90/2681856] via 192.168.1.2, 01:37:18, Serial0
       10.0.0.0/24 is subnetted, 8 subnets
B      10.1.3.0 [200/0] via 172.16.30.33, 01:13:06
B      10.10.10.0 [200/0] via 172.16.30.33, 01:14:18
B      10.1.1.0 [200/0] via 172.16.30.33, 01:33:57
D      10.1.6.0 [90/2195456] via 192.168.1.2, 01:34:39, Serial0
B      10.40.4.0 [200/0] via 172.16.30.33, 01:13:12
B      10.50.5.0 [200/0] via 172.16.30.33, 01:14:14
C      10.70.7.0 is directly connected, Loopback0
       192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.1.0/24 is directly connected, Serial0
C      192.168.1.2/32 is directly connected, Serial0
r7#
```

Cat

```
cat#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
```

```

       10.0.0.0/24 is subnetted, 10 subnets
B      10.1.3.0 [20/0] via 10.1.4.1, 3d03h
B      10.1.1.0 [20/0] via 10.1.5.1, 01:14:15
C      10.10.10.0 is directly connected, Loopback0
B      10.1.6.0 [20/0] via 10.1.5.1, 01:14:46
```

```
C      10.1.5.0 is directly connected, Vlan50
C      10.1.4.0 is directly connected, Vlan40
B      10.40.4.0 [20/0] via 10.1.4.1, 3d03h
B      10.50.5.0 [20/0] via 10.1.5.1, 3d03h
B      10.70.7.0 [20/0] via 10.1.5.1, 01:14:16
cat#
```

Explanation

1. Configure the switch ports, VLANs, and router ethernet interfaces according to the diagram and table. You will need to create on Cat the Switched Virtual Interfaces (SVIs) VLAN40 and VLAN50.

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

No explanation needed.

2. Configure the frame relay links as shown in the diagram. Note that there is a full-mesh among R1, R2 and R3.

It is very important that there is a full-mesh of frame-relay links here. A full-mesh lends itself to IS-IS, and also means that we will not need to configure R1 as a BGP route reflector. Be careful to configure the full-mesh, rather than following habit and configuring the hub-and-spoke arrangement so often used in practice scenarios. Here is output from "sh frame map" on R1, R2, and R3:

```
r1#sh frame map
Serial0 (up): CLNS dlci 102(0x66,0x1860), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): CLNS dlci 103(0x67,0x1870), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): ip 172.16.30.1 dlci 102(0x66,0x1860), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): ip 172.16.30.2 dlci 102(0x66,0x1860), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): ip 172.16.30.3 dlci 103(0x67,0x1870), static,
              broadcast,
              CISCO, status defined, active
r1#
```

```

r2#sh frame map
Serial0 (up): CLNS dlci 201(0xC9,0x3090), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): CLNS dlci 203(0xCB,0x30B0), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): ip 172.16.30.1 dlci 201(0xC9,0x3090), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): ip 172.16.30.2 dlci 201(0xC9,0x3090), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): ip 172.16.30.3 dlci 203(0xCB,0x30B0), static,
              broadcast,
              CISCO, status defined, active
r2#

r3#sh frame map
Serial0.2 (up): CLNS dlci 304(0x130,0x4C00), static,
               broadcast,
               CISCO, status defined, active
Serial0.1 (up): CLNS dlci 301(0x12D,0x48D0), static,
               broadcast,
               CISCO, status defined, active
Serial0.1 (up): CLNS dlci 302(0x12E,0x48E0), static,
               broadcast,
               CISCO, status defined, active
Serial0.1 (up): ip 172.16.30.1 dlci 301(0x12D,0x48D0), static,
               broadcast,
               CISCO, status defined, active
Serial0.1 (up): ip 172.16.30.2 dlci 302(0x12E,0x48E0), static,
               broadcast,
               CISCO, status defined, active
Serial0.1 (up): ip 172.16.30.3 dlci 301(0x12D,0x48D0), static,
               broadcast,
               CISCO, status defined, active
Serial0.2 (up): ip 172.16.30.21 dlci 304(0x130,0x4C00), static,
               broadcast,
               CISCO, status defined, active
Serial0.2 (up): ip 172.16.30.22 dlci 304(0x130,0x4C00), static,
               broadcast,
               CISCO, status defined, active
r3#

```

Whenever you configure frame relay, you should check to make sure that you do not have any DLCIs available that should not be available. Even after passing the CCIE Lab and developing numerous practice labs, I still find that I sometimes forget to disable inverse-arp somewhere in a scenario.

3. Configure the serial links as shown in the diagram using PPP encapsulation.

No explanation needed.

4. Configure IS-IS as shown in the diagram. Use passive interface on interfaces on which it does no good to send out hellos. Use L1, L2, or L1-L2 as you like.

We can see from the following output that all IS-IS routers are participating in IS-IS. You may also note that R2 and R3 are both L1 and L2 routers in relation to R1.

```
r1#sh isis topology
```

```
IS-IS paths to level-1 routers
System Id      Metric  Next-Hop      Interface      SNPA
r1             --
r2           10    r2           Se0           DLCI 102
r3           10    r3           Se0           DLCI 103
```

```
IS-IS paths to level-2 routers
System Id      Metric  Next-Hop      Interface      SNPA
r1             --
r2           10    r2           Se0           DLCI 102
r3           10    r3           Se0           DLCI 103
r4           20    r3           Se0           DLCI 103
r5           20    r2           Se0           DLCI 102
r1#
```

"Sh isis topo" is the first-thing I would do to check my IS-IS configuration, but it is by no means the final step. You might want to "show isis database" and also check each router's routing table to insure that all IS-IS routes are shared among the IS-IS routers.

```
r1#sh isis database
```

```
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
r1.00-00      * 0x0000038B  0xE729        1101           1/0/0
r2.00-00      0x000002D8  0x527A        773            1/0/0
r3.00-00      0x00000237  0x1AB0        834            1/0/0
r3.01-00      0x00000238  0x420A        1027           0/0/0
r3.02-00      0x00000157  0x5821        1172           0/0/0
IS-IS Level-2 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
r1.00-00      * 0x00000398  0x82BB        1085           0/0/0
r2.00-00      0x0000038D  0xE354        1059           0/0/0
r3.00-00      0x00000241  0xDABE        744            0/0/0
r3.01-00      0x00000234  0x9E58        1030           0/0/0
r4.00-00      0x0000031B  0xF196        762            0/0/0
r4.02-00      0x000002A6  0x281B        1169           0/0/0
r5.00-00      0x000001BF  0xDBD9        651            0/0/0
r1#
```

```
r1#sh ip rou isis
```

```
172.16.0.0/16 is variably subnetted, 8 subnets, 4 masks
i L2 172.16.30.18/32 [115/20] via 172.16.30.2, Serial0
i L1 172.16.30.16/30 [115/20] via 172.16.30.2, Serial0
```

```

i L2    172.16.30.17/32 [115/30] via 172.16.30.2, Serial0
i L1    172.16.30.20/30 [115/20] via 172.16.30.3, Serial0
        10.0.0.0/24 is subnetted, 10 subnets
i L1    10.1.3.0 [115/10] via 172.16.30.3, Serial0
i L1    10.1.2.0 [115/10] via 172.16.30.2, Serial0
i L2    10.1.5.0 [115/20] via 172.16.30.2, Serial0
i L2    10.1.4.0 [115/20] via 172.16.30.3, Serial0
i L2    10.40.4.0 [115/20] via 172.16.30.3, Serial0
i L2    10.50.5.0 [115/20] via 172.16.30.2, Serial0
r1#

```

The routes to networks in other IS-IS areas are level-2 routes, as they should be.

5. Configure dialer watch on R1 so that R1 will dial R2 over the ISDN link when there is no longer a route to R2 e0 in R1 s routing table. Configure ISDN so that bri0:2 will be used whenever bri0:1 is even minimally loaded.

Since IS-IS is designed to be used in carrier networks, it is fair to say that it was not designed with customers' ISDN dial-on-demand backup links in mind. Nevertheless, dialer watch would seem to offer a way to use dial-on-demand ISDN links in an IP network using integrated IS-IS for IP routing. I cannot say it is easy to configure, and I had enough problems that I would personally hesitate to implement dialer watch with IS-IS in a production network.

Just as you need a frame relay mapping for CLNS, you need a dialer map for CLNS on an ISDN interface. This frame relay map statement does not require the dial number, a fact one of my technical reviewers pointed out and tested. Once dialer watch brings up the ISDN interface, an adjacency should form as long as you have enabled IS-IS under the ISDN interface.

If you have trouble getting dialer-watch to work, remove the dialer watch-group under the interface and remove the dialer watch-list in global configuration mode. You may also want to reset the ISDN link.

With dialer watch, the subnet mask used in the dialer watch list must match the subnet mask of the watched network precisely. It is good to enable ip routing debugging to see exactly what networks are added and taken away when you shut down the primary link.

I configured "isis metric 30" under R1 interface bri0, to make sure that the bri0 interface would not be favored under normal conditions.

You may want to tinker with the configuration to fine-tune it. Note that I used a dialer-load threshold of 5. Cisco warns against using a small value like 1 or 2 because it can cause both bri0:1 and bri0:2 to come up at the same time and contend for the line. I recently saw an example on CCO with a dialer-threshold set to 1, so do not expect this advice about avoiding very small dialer-threshold values to be universally implemented.

6. Configure IS-IS interface authentication, area authentication, and domain authentication, using the passwords "interface," "area," and "domain," respectively.

There are three levels of IS-IS authentication:

- 1) interface authentication
- 2) area authentication
- 3) domain authentication

The password for interface authentication is specific for L1 or L2, so be sure you understand what kind of adjacency exists on that interface, which you can ascertain by "show clns neighbor." There is a useful Cisco article entitled, "Configuring IS-IS Authentication."

A point to remember about authentication of routes in IS-IS is that you could easily have all your IS-IS adjacencies up but still not have routes advertised throughout an area or between areas in a domain. You need to carefully examine what is happening before assuming that IS-IS is configured correctly and moving on to other tasks.

It could be considered unnecessary "over-configuration," but I went ahead and configured both level-1 and level-2 passwords on interfaces on which both level-1 and level-2 adjacencies could form. You will only get L2 adjacencies between routers located in two different IS-IS areas, so I did not specify L1 passwords on those links. If you want to nail things down more precisely, you could specify that only a level-1 or level-2 adjacency can form on an interface, and then specify a password for only the appropriate level.

7. Configure EIGRP on R6 and R7.

Make sure that you include R6 s1 in EIGRP (or redistribute it into EIGRP) so that R7 will be able to reach the next hop for BGP prefixes received from R1. Alternatively, you could use "next-hop-self" on R6.

8. Enable IP routing on Cat, but do not run any internal routing protocol on it.

No explanation needed.

9. Do not configure any internal route redistribution.

This means that for routing among autonomous systems we will be relying on BGP. Obviously, this also means that connectivity will be not be any-to-any.

10. Configure R1, R2, R3, R4, and R5 as BGP speakers in AS 12345. Do not peer from loopback to loopback in this scenario.

There is no need to configure R1 as a route-reflector because there is a full-mesh among R1, R2, and R3. We will need to use a confederation, because R4 and R5 are not in a full-mesh with all AS 12345 BGP speakers.

You may have noticed that I am guilty of slight "overconfiguration" of this task, in that I manually configured the BGP confederation peers on R1, where this was not necessary. You do not need to specify the confederation peers on R1, since it peers with BGP speakers within the same sub-AS, and with a BGP speaker in AS 67.

11. Configure R6 and R7 as BGP speakers in AS 67.

If you did not make sure that R7 knows how to reach the R6 s1 s network, you will have a next-hop reachability problem on R7. You can correct this by correcting the internal routing appropriately, or by using "next-hop-self" on R6.

If you can do so, you should always disable synchronization and auto-summarization under BGP. This is a persistent error for me, as I am tempted to go ahead and configure the BGP neighbors and advertise the networks. These tasks require some concentration, and before I know it I have forgotten to configure "no synch" and "no auto."

The problem with failing to configure "no synch" is that your BGP tables may look great, with a "best" route for every destination (">"). It is only when you check your routing table that you notice that some routes that seemed to be available through BGP were never put into the routing table.

12. Configure Cat as a BGP speaker in AS 10.

An examination of Cat's BGP table shows the BGP prefixes advertised from the other autonomous systems, and there is a valid and "best" route for every appropriate network:

```
cat#sh ip bgp
```

```
BGP table version is 22, local router ID is 10.10.10.1
```

```
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	10.1.4.1			0	12345 i
*>		10.1.5.1			0	12345 i
*	10.1.2.0/24	10.1.5.1			0	12345 i
*>		10.1.4.1			0	12345 i
*	10.1.3.0/24	10.1.5.1			0	12345 i
*>		10.1.4.1			0	12345 i
*	10.1.6.0/24	10.1.4.1			0	12345 67 i
*>		10.1.5.1			0	12345 67 i
*>	10.10.10.0/24	0.0.0.0	0		32768	i
*	10.40.4.0/24	10.1.5.1			0	12345 i
*>		10.1.4.1	0		0	12345 i
*	10.50.5.0/24	10.1.4.1			0	12345 i
*>		10.1.5.1	0		0	12345 i
*	10.70.7.0/24	10.1.4.1			0	12345 67 i
*>		10.1.5.1			0	12345 67 i

cat#

13. Configure peering between AS 67 and AS 12345, and between AS 12345 and AS 10.

No explanation needed.

14. Advertise loopbacks in BGP, and Ethernets except for R2 e0. Do not peer from loopback to loopback.

The problem with advertising R2 e0 s network in BGP would be that the BGP route would take a while to drop out of R1 s routing table, so dialer watch would not work well.

15. Configure local-preference in AS 12345 so that the R5-Cat ethernet link will be preferred for any traffic destined to AS 10. We are not concerned about the path taken by traffic going in the opposite direction.

On R5, we set local preference to a very high value for all prefixes received from Cat on the Cat-R5 ethernet link. The local-preference value should be shared throughout AS 12345. R1 s BGP table shows the result:

```
r1#sh ip bgp
BGP table version is 41, local router ID is 172.16.30.33
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      0.0.0.0           0      32768 i
*>i10.1.2.0/24      172.16.30.2       0      100    0 i
*>i10.1.3.0/24      172.16.30.3       0      100    0 i
*> 10.1.6.0/24      172.16.30.34      0
*>i10.10.10.0/24  10.1.5.2          0 10000  0 (65505) 10 i
*>i10.40.4.0/24     172.16.30.22      0      100    0 (65504) i
*>i10.50.5.0/24     172.16.30.18      0      100    0 (65505) i
*> 10.70.7.0/24     172.16.30.34      0      67 i
r1#
```

The next-hop is Cat VLAN50 s IP address, which is as it should be.