



iPexpert
NEXT GENERATION

CCIE

ROUTING & SWITCHING

WORKBOOK

(Vol 2)

iPexpert's Cisco CCIE R&S (v5) 8-Hour Mock Lab Workbook (Vol. 2)

Congratulations! You now possess one of the ULTIMATE CCIE Lab preparation resources available! This resource was produced by senior engineers, technical instructors, and authors boasting decades of internetworking experience. Although there is no way to guarantee a 100% success rate on the CCIE Lab exam, we feel confident that your chances of passing the Lab will improve dramatically after completing this industry-recognized Workbook!

iPexpert is proud to lead the industry with multiple support options at your disposal free of charge. Our online communities have attracted a membership of your peers from around the world! At blog.ipexpert.com you can keep up to date with everything iPexpert does. At community.ipexpert.com, you may collaborate with your CCIE and CCIE-candidate peers.

Feedback

Do you have a suggestion or other feedback regarding this book or other iPexpert products? At iPexpert, we look to you – our valued clients – for the real world, frontline evaluation that we believe is necessary to improve continually. Please send an email with your thoughts to feedback@ipexpert.com or call 1.866.225.8064 (international callers dial +1.810.326.1444).

In addition, when you pass the CCIE Lab exam, we want to hear about it! Email your CCIE number to success@ipexpert.com and let us know how iPexpert helped you succeed. We would like to send you a gift of thanks and congratulations.

Additional CCIE Preparation Material

iPexpert is committed to developing the most effective Cisco CCIE Collaboration, Data Center, Routing & Switching, Security, and Wireless certification preparation tools available. Our team of certified networking professionals develops the most up-to-date and comprehensive

materials for networking certification, including self-paced workbooks, online Cisco hardware rental, classroom training, online (distance learning) instructor-led training, audio products, and video training materials.

We employ an experienced and accomplished team of experts to create, maintain, and constantly update our products. At iPexpert, we are focused on making your CCIE Lab preparation more effective.

A message from the Author

The scenarios covered in this workbook were developed by CCIEs to help you prepare for the Cisco CCIE Lab. It is strongly recommended that you use other reading materials in addition to this workbook.

Training is not this CCIE workbook's objective. The intent of these labs is to test your knowledge and abilities in implementing Cisco Enterprise Security Solutions.

Time management is very important. If you get stuck on a lab scenario be sure to write it down. Formulate a Checklist for skipped sections and then return to those sections once you have gone through the entire lab. Be sure to revisit the questions that you do not understand.

For more information on the CCIE lab, please visit <http://www.cisco.com/go/ccie>.

Helpful Hints

- Keep It Simple; try to avoid any extra work (example: adding descriptions).
- Always reference everything from the Documentation Website:
<http://www.cisco.com/web/psa/products/index.html>
- Save your router configurations often (wr is the quickest command).

Troubleshooting Lab 1

Please look at the provided diagrams and read through the whole lab before you start. Read the directions very carefully to make sure you are doing what is being asked of you. This is very important when you take Cisco's CCIE lab.

Each incident contains a small diagram that is designed to show you the focus of where the issue is.

Multiple topology diagrams are available for this lab, including an IPv4 and a BGP diagram.

General Rules

- You may modify, but not delete or remove any prefix-lists, route-maps, or access-lists.
- Do not modify any IP addressing on any interfaces
- The BB routers are not accessible
- Static/default routes are NOT allowed unless otherwise stated in the task
- Save your configurations often

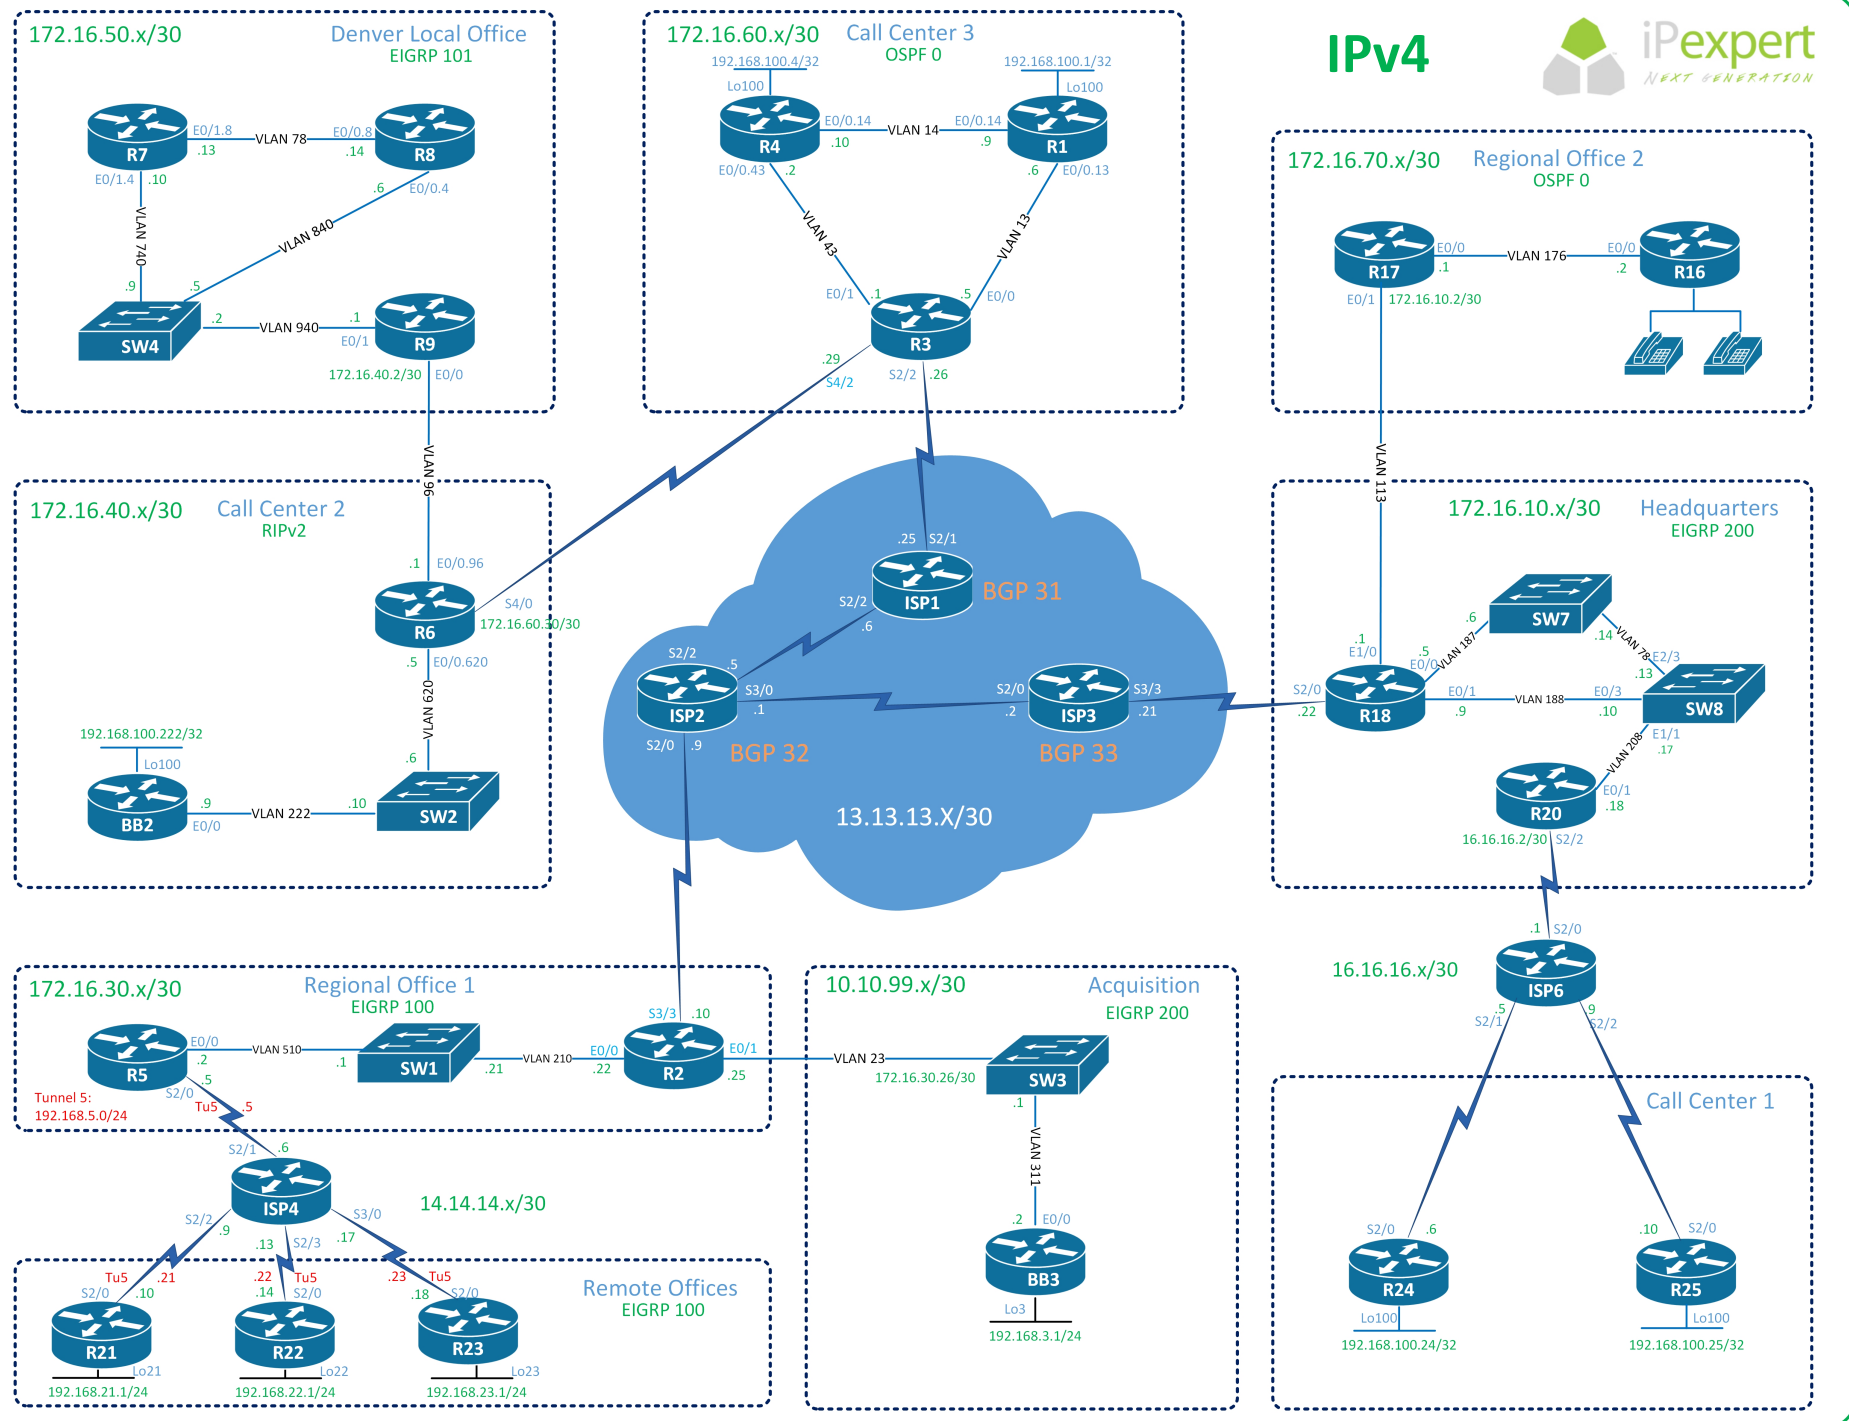
Estimated Time to Complete: **2-2.5 Hours**

Total Possible Points: **24**

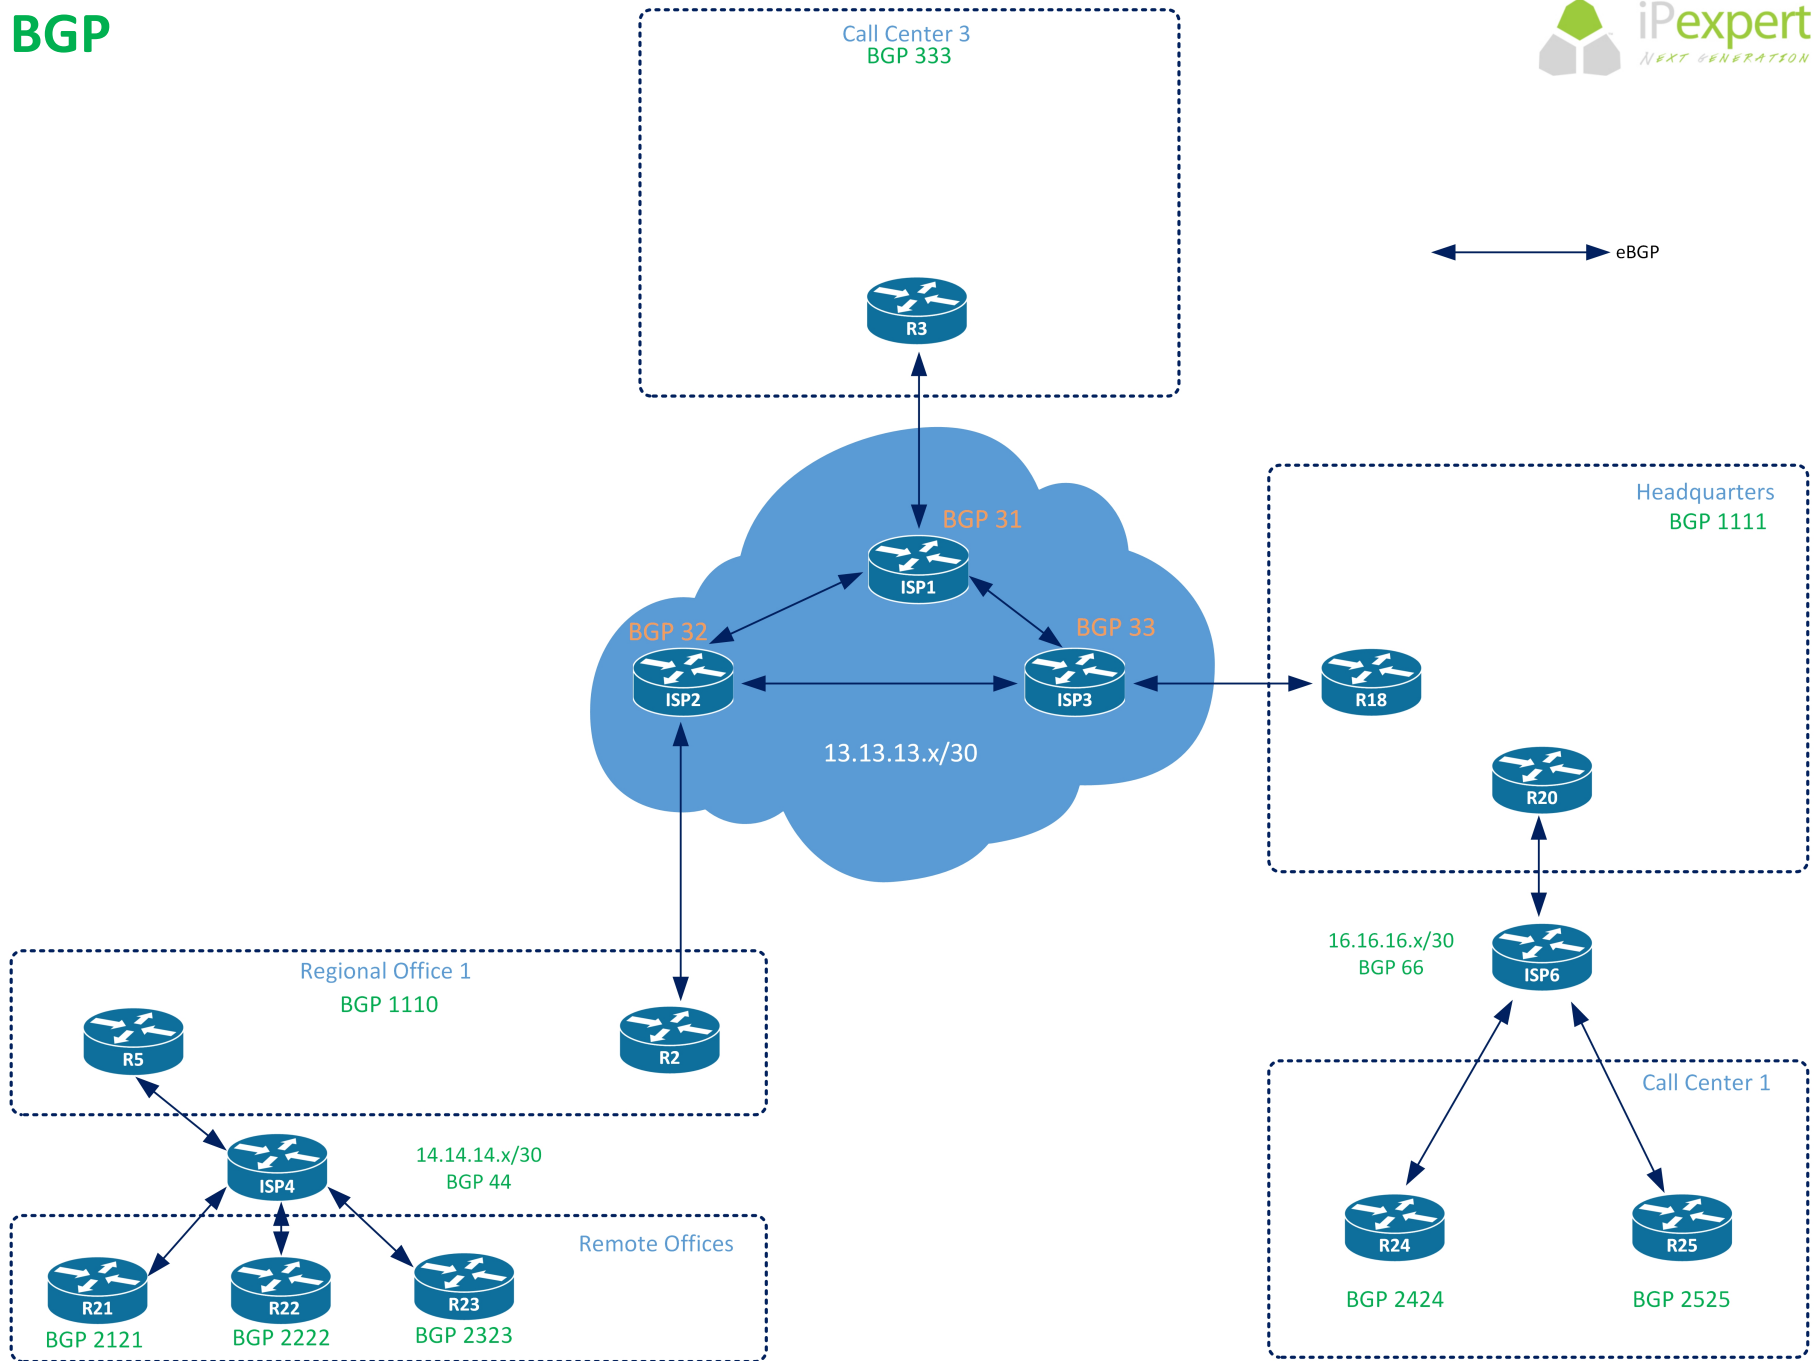
Pre-setup

Please login to your vRack at ProctorLabs.com and load the initial Configuration. This lab is intended to be used with online rack access provided by our partner Proctor Labs (www.proctorlabs.com). Connect to the terminal server and complete the troubleshooting tasks as detailed below.

IPv4

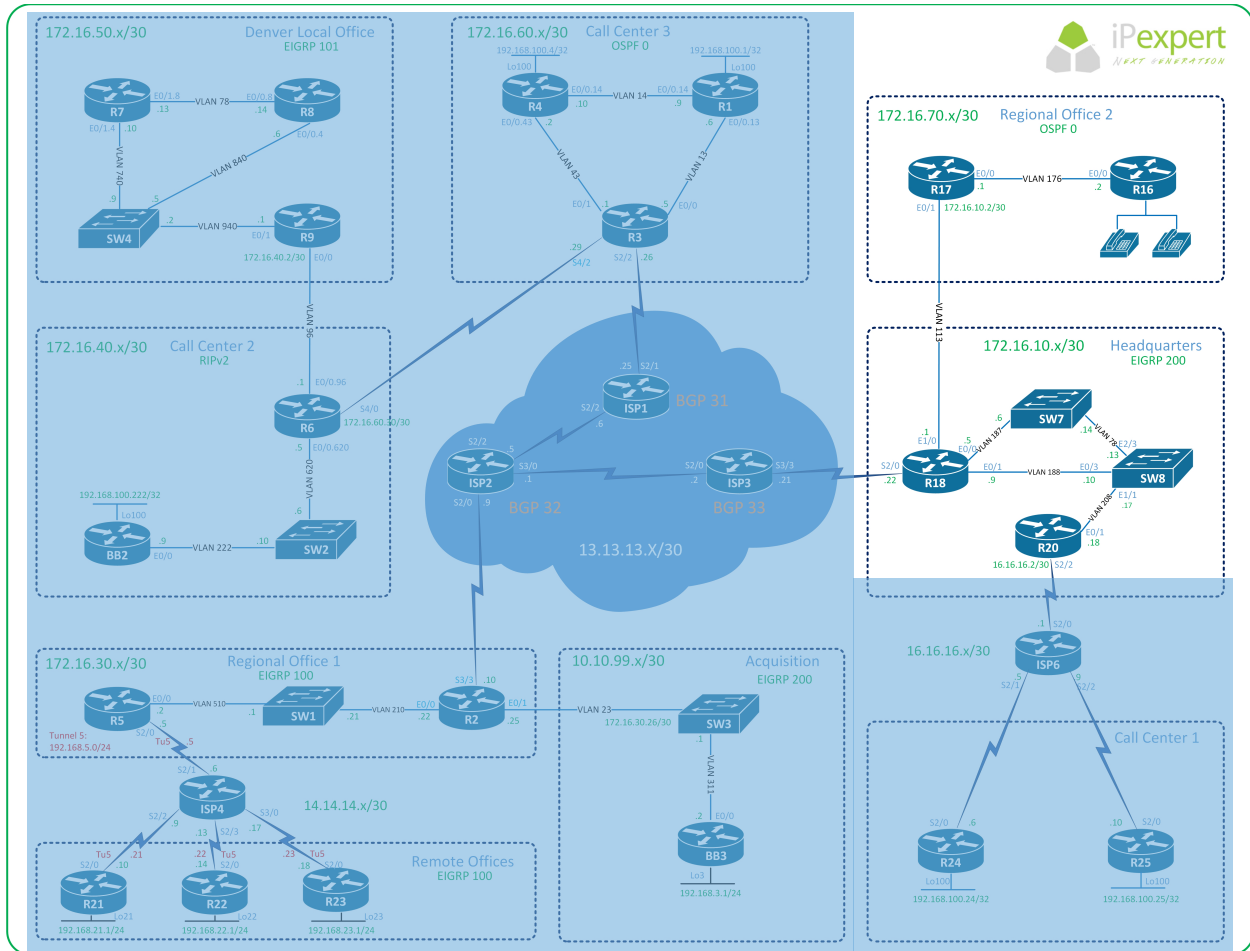


BGP



Incident 1

(2 points)



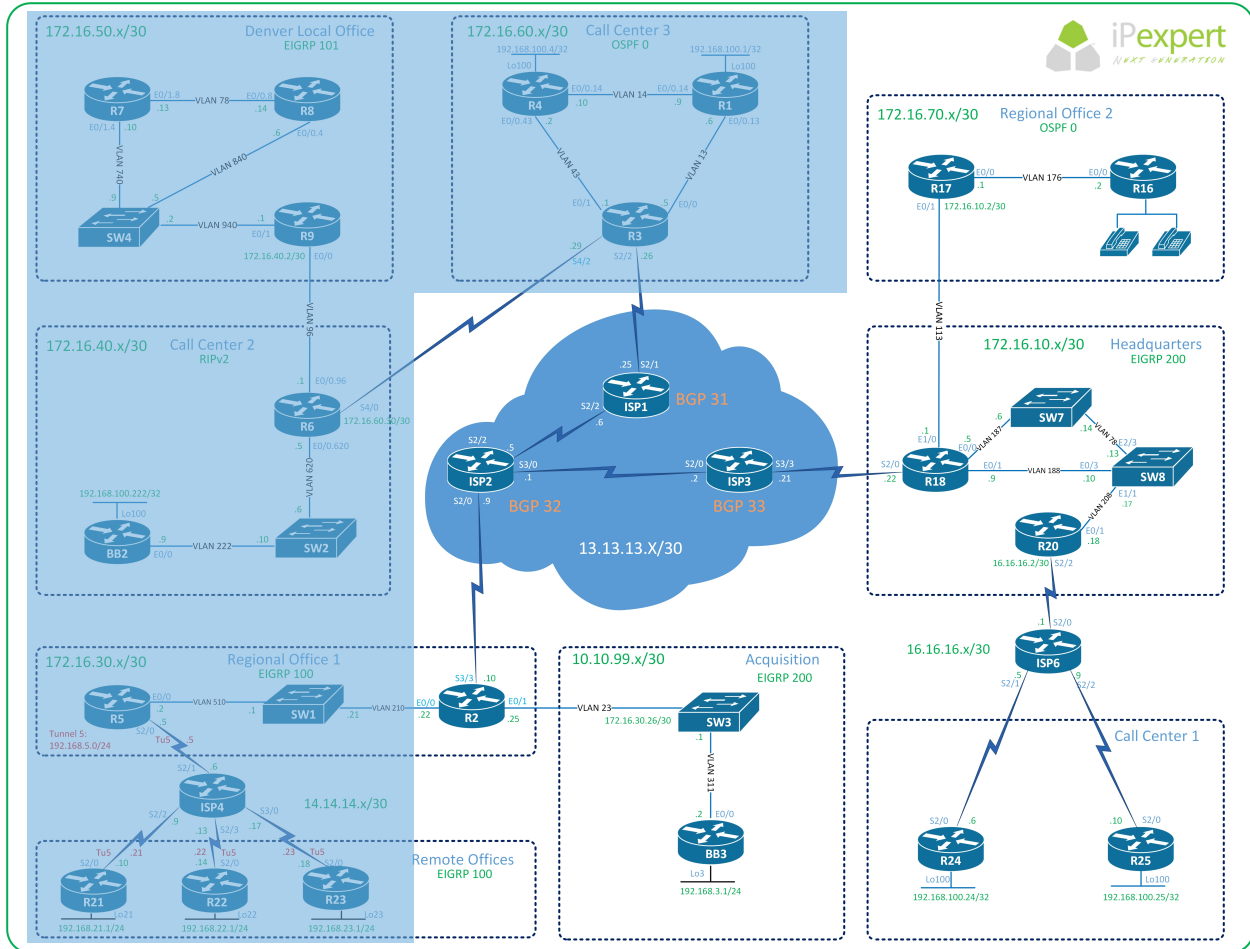
- R20 cannot copy the startup-config from R17 using TFTP. Isolate and fix the issues so that the output matches below:

```
R20#copy tftp://172.16.10.2/startup-config null:
Accessing tftp://172.16.10.2/startup-config . . .
Loading startup-config from 172.16.10.2 (via Ethernet1/0) : !
[OK - 2449 bytes]
```

```
2449 bytes copied in 0.81 secs (21876 bytes/sec)
R20#
```

Incident 2

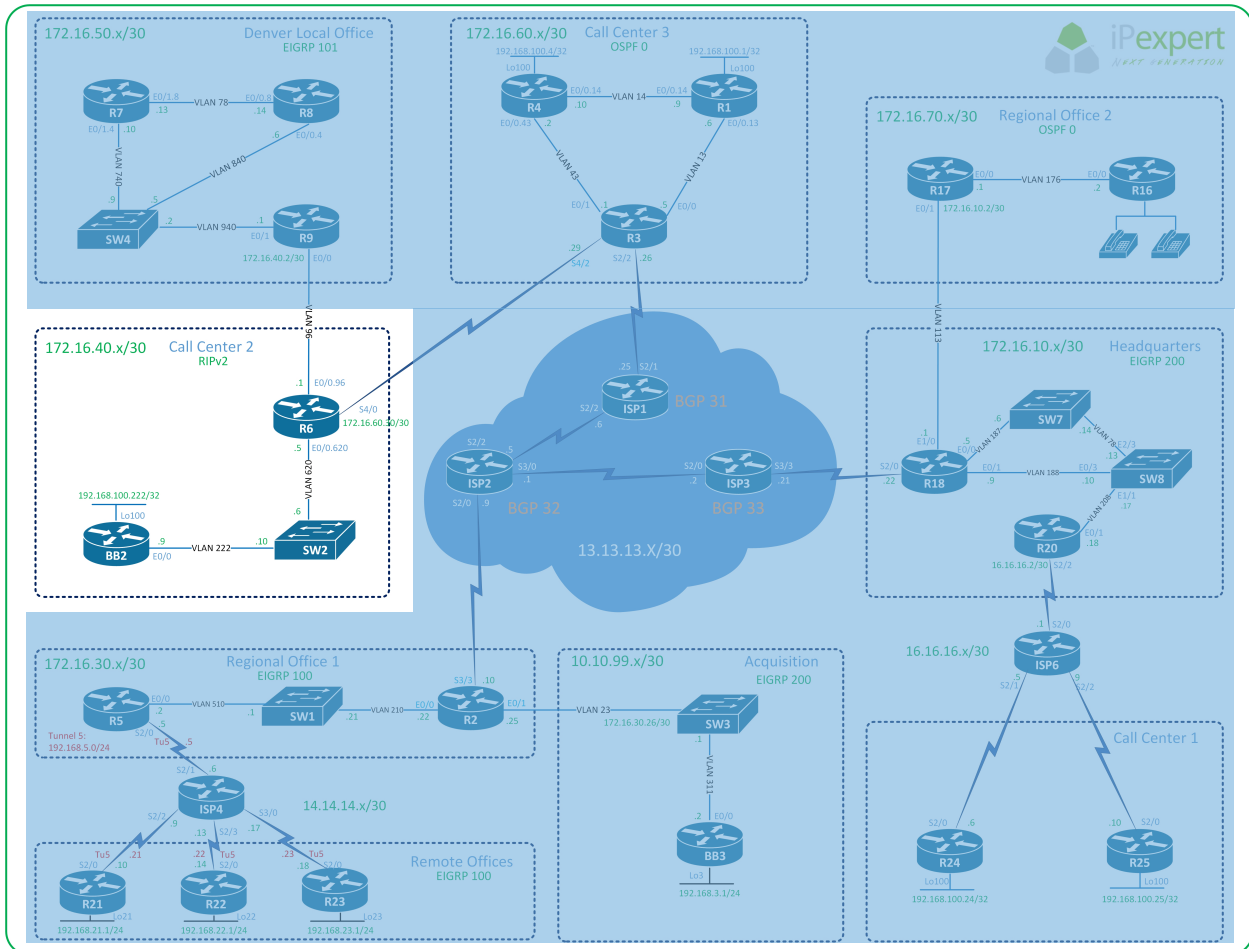
(3 points)



- Your company recently acquired a new office, called Acquisition. It is connected to Headquarters through the BGP connection at Regional Office 1. After a configuration change in WAN Cloud, the acquisition site lost connectivity to HQ.
- Troubleshoot and restore connectivity back to HQ.
- You may login to the ISP Routers in the cloud and make changes as needed.
- This incident contains multiple faults.

Incident 3

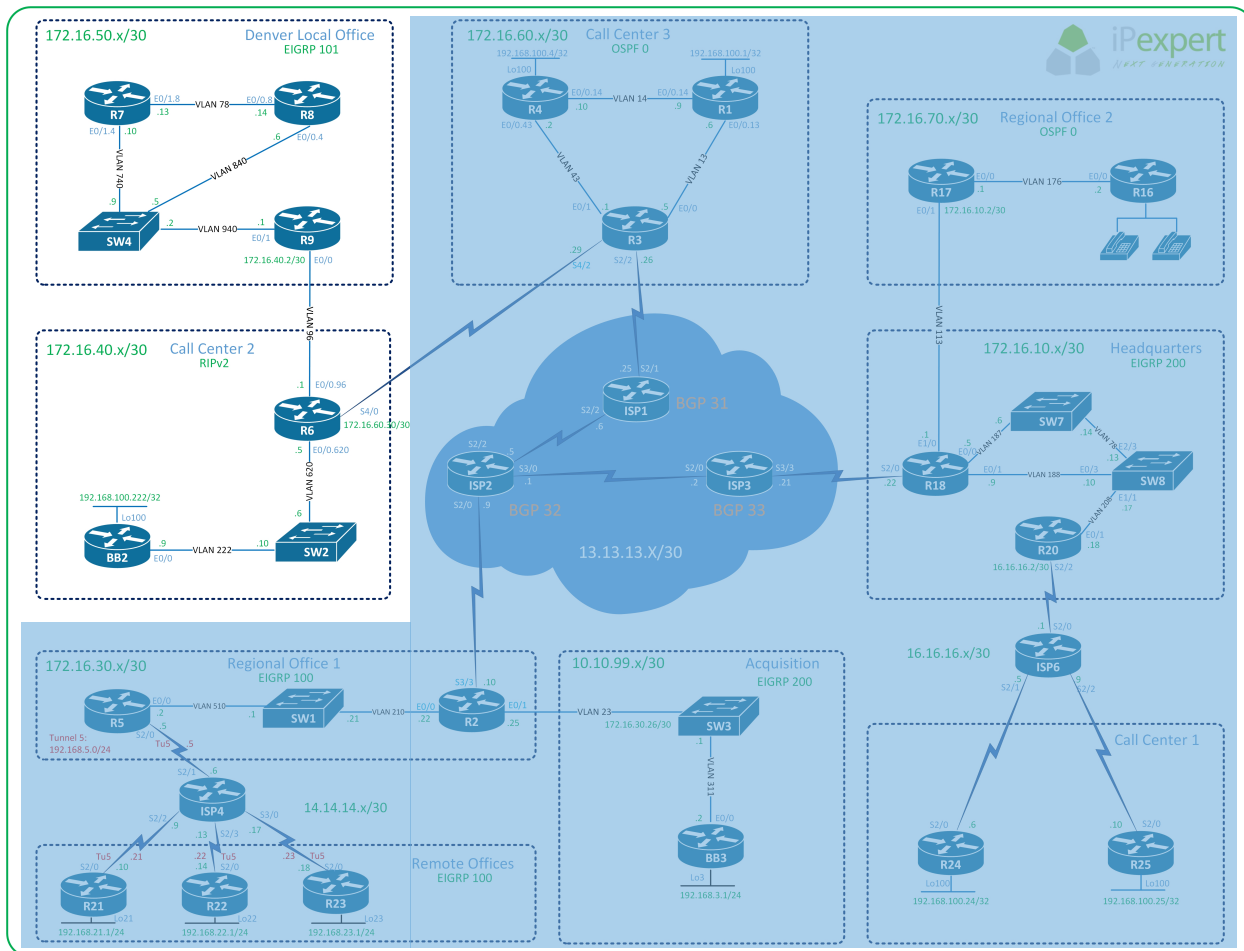
(2 points)



- Users at Call Center 2 are reporting that they cannot reach the loopback 100 interface of BB2.
- Isolate and fix the issue so that Loopback 100 of BB2 is reachable from all interfaces in the RIP routing domain.

Incident 4

(2 points)



- R6 has been configured to mark packets incoming from the Denver Local Office, destined for 192.168.100.222, with IP precedence of 5. SW2 fails to indicate the marked packets on the link with R6.
- Troubleshoot and fix the issue so that SW2 indicates packets marked with IP Precedence coming from the Denver Local Office and destined for 192.168.100.222.
- The following command should show the highlighted counter incrementing when traffic is generated from the 172.16.50.0/24 network. Generate a ping from 172.16.50.2 to 192.168.100.222 and verify that the highlighted counter increments.

```
SW2#sh interface vlan 620 precedence
```

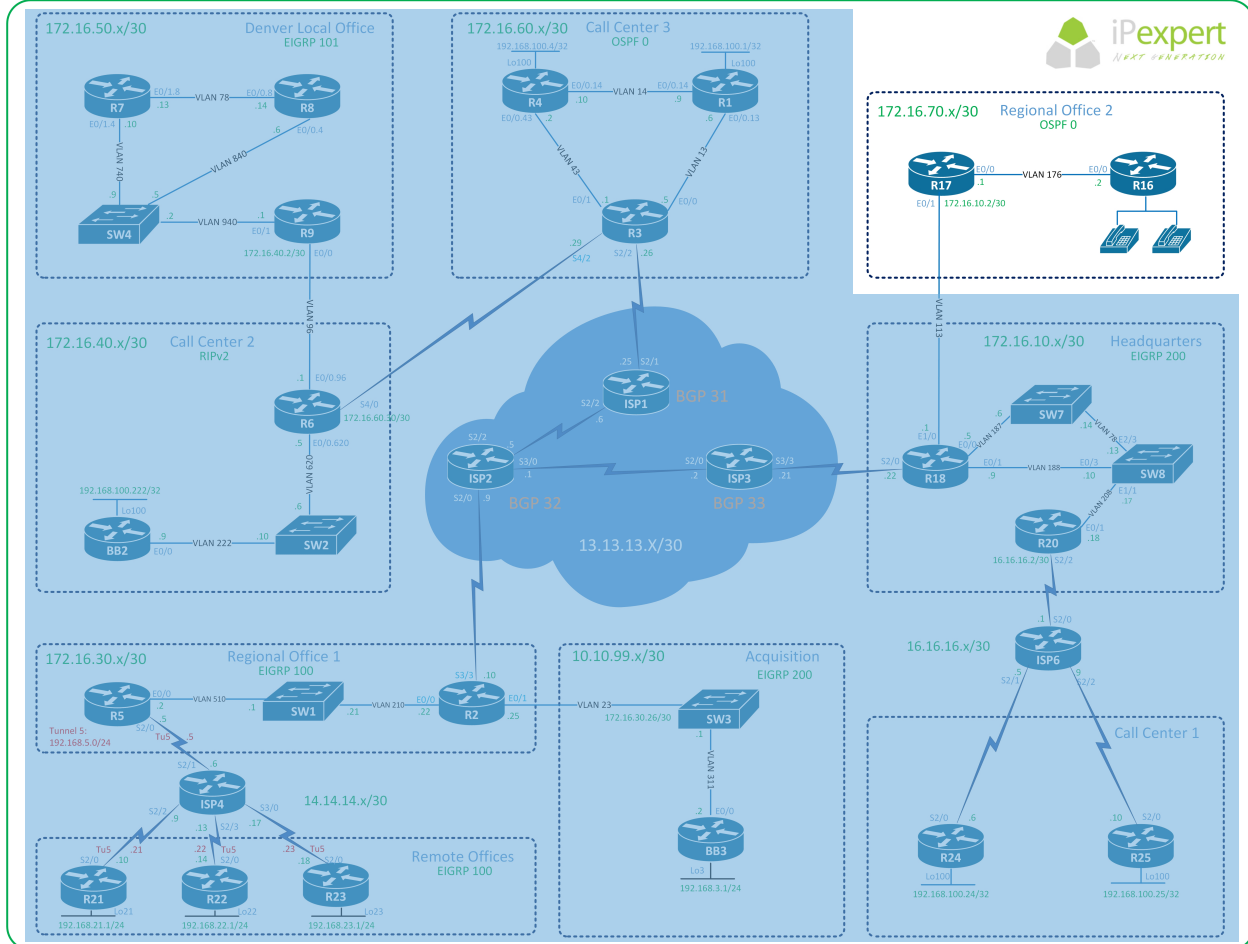
```
Vlan620
```

```
Input
```

```
Precedence 5: 25 packets, 2850 bytes
```

```
Precedence 6: 13 packets, 3978 bytes
```

- This incident contains multiple faults and is dependent on Incident 3

Incident 5**(2 points)**

- Voice users located at Regional Office 2 have reported that they are not able to make VoIP calls to HQ.
- Fix the issue so that connectivity is fully restored and the following ping command has a successful result:

```
R17#ping 172.16.16.16
```

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

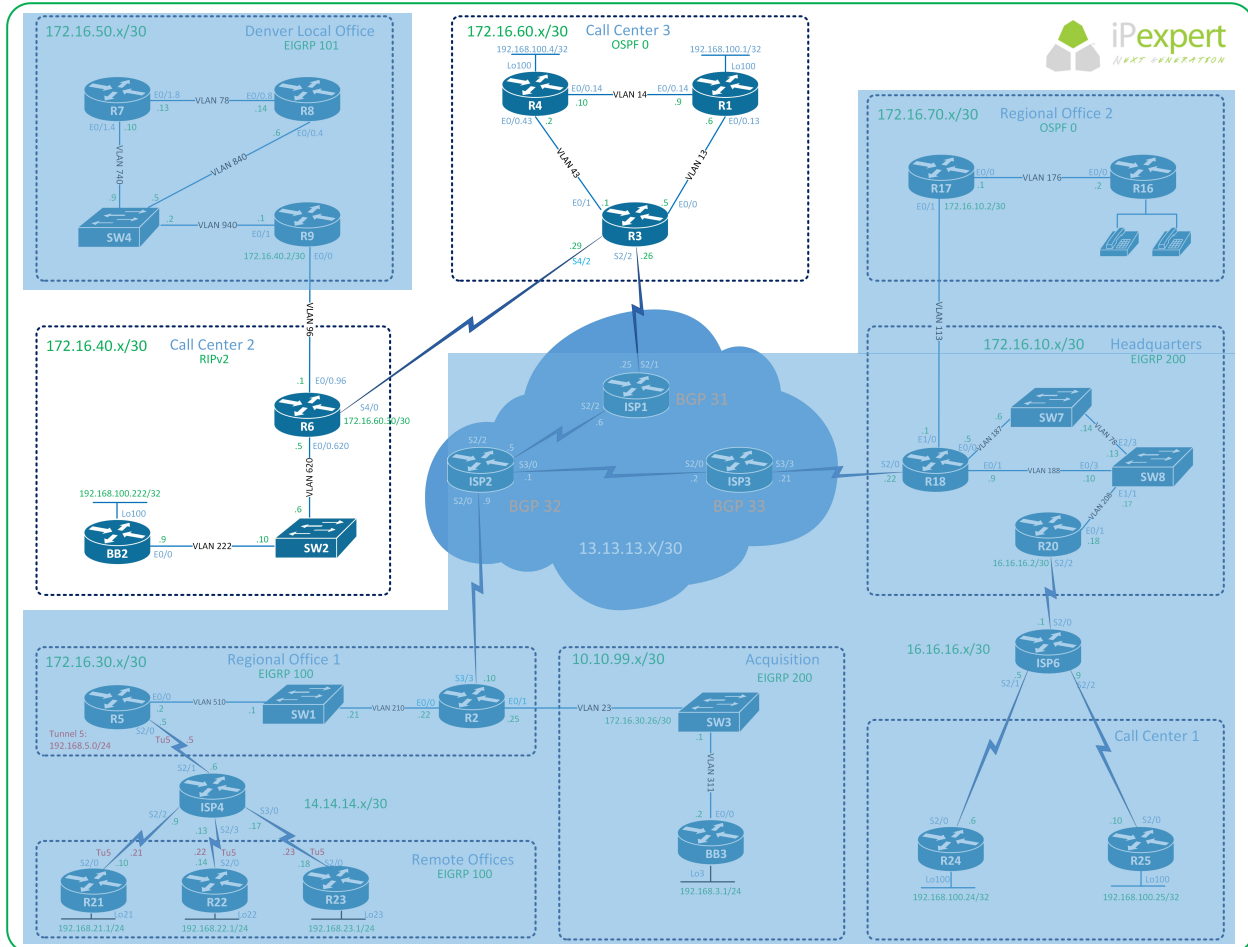
```
Sending 5, 100-byte ICMP Echos to 172.16.16.16, timeout is 2 seconds:
```

```
!!!!!!
```

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

Incident 6

(3 points)



- Users in Call Center 2 have reported that they have lost connectivity to the rest of the network. Isolate and fix the issues so that the Loopback 100 interface on R1 and R4 can ping 172.16.40.5.

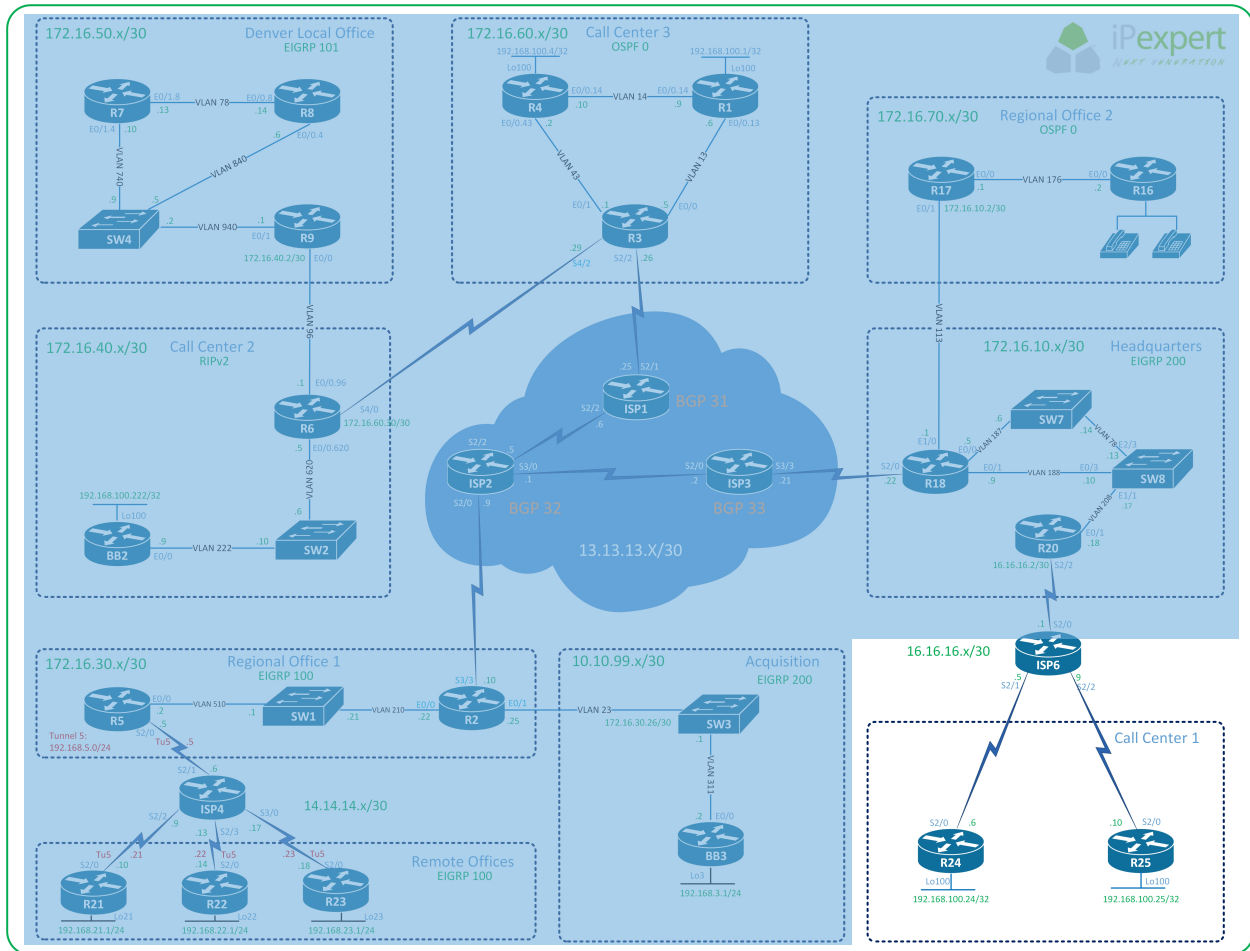
```
R1#ping 172.16.40.5 source loopback100
```

```
R4#ping 172.16.40.5 source loopback100
```

- Refer to the Troubleshooting BGP Diagram.
- There are multiple faults in this incident.

Incident 7

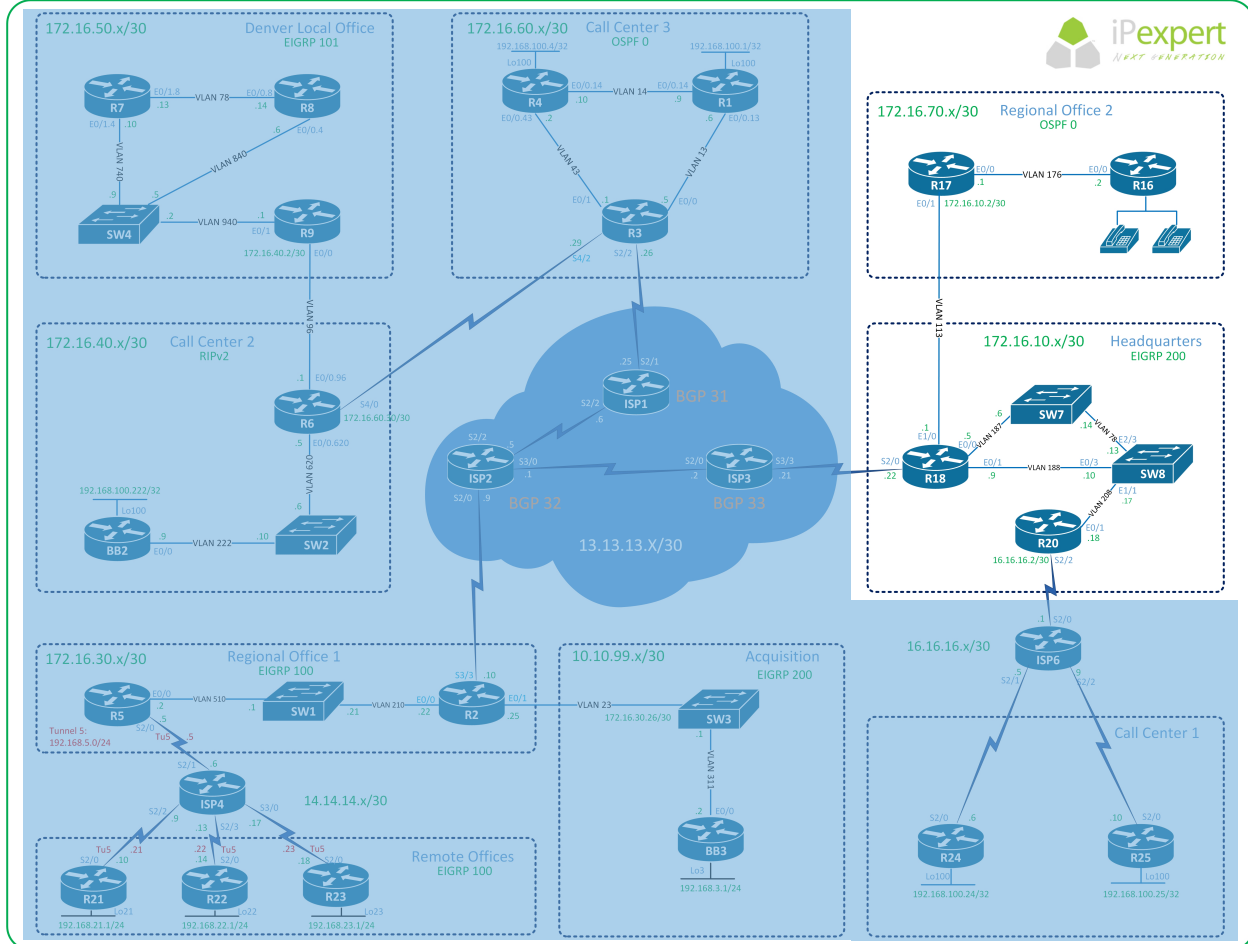
(2 points)



- Loopback100 on R24 has lost connectivity to Loopback100 on R25. Troubleshoot and restore connectivity so that the following pings are successful:

```
R24#ping 192.168.100.25 source loopback100
R25#ping 192.168.100.24 source loopback100
```

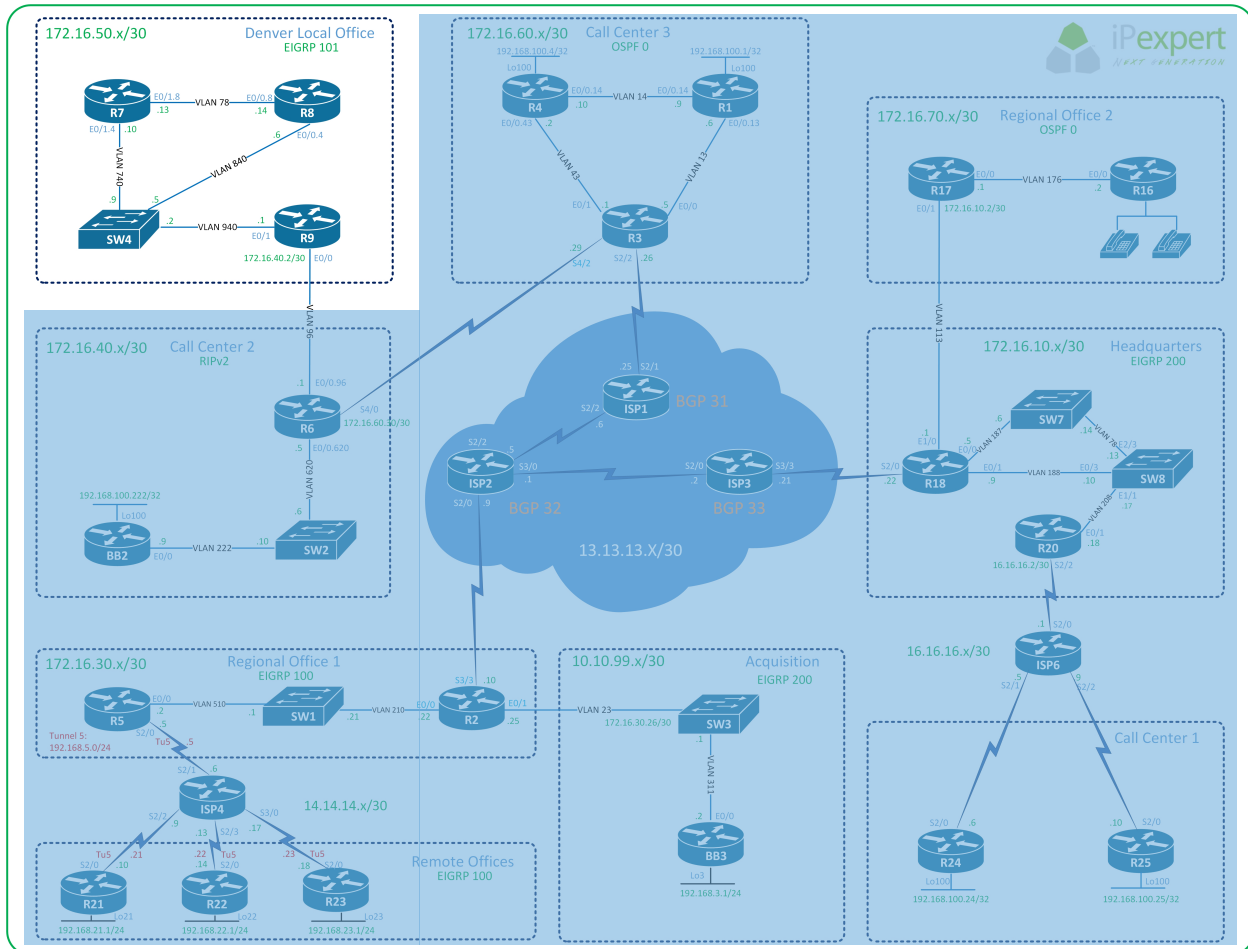
- Refer to the Troubleshooting BGP Diagram.
- This incident contains multiple faults.

Incident 8**(3 points)**

- SW8 is preferring the link through SW7 to get to Regional Office 2. Optimize the path by preferring the direct link between SW8 and R18 instead.
- All links between R18, SW7, and SW8 are 10Mbps.
- The path from SW8 to Regional Office 2 should go directly through R18 and use the path through SW7 as a backup.
- Verify that both routes are in the EIGRP topology table, but only the preferred route is in the routing table and is NOT load balancing.
- This incident contains multiple faults.

Incident 9

(2 points)



- Users at the Denver Local Office complain that their multicast applications are not working properly. R9 cannot ping all members of the 226.8.8.8 group.
- This ticket has multiple faults.
- Fix the issue so that R9 can ping the multicast group and gets a response from SW4, R7, and R8.
- Do not activate multicast on any interfaces that are currently not running multicast.
- The following ping needs to reflect the output below:

```
R9#ping 226.8.8.8
```

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

```
Sending 1, 100-byte ICMP Echos to 226.8.8.8, timeout is 2 seconds:
```

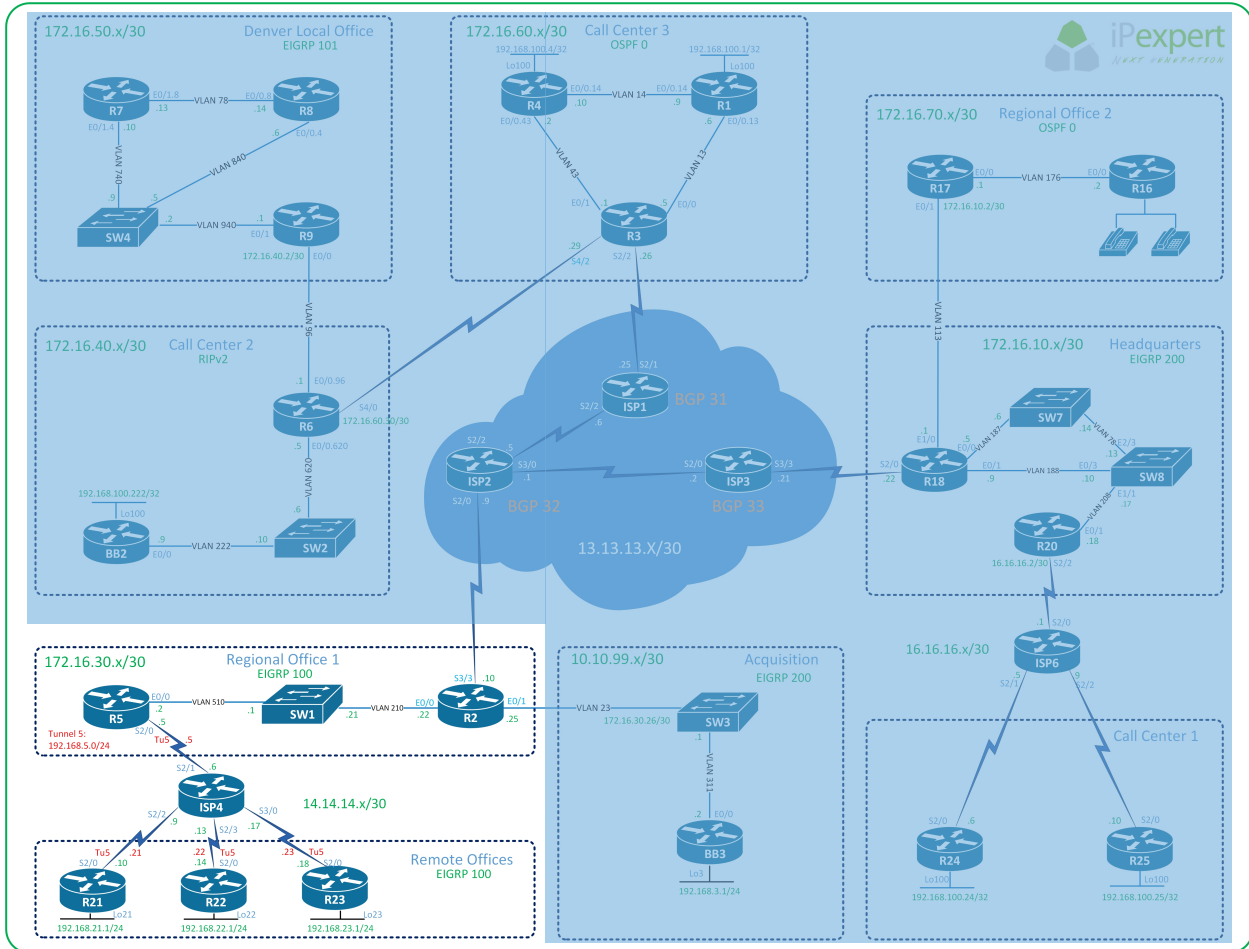
```
Reply to request 0 from 172.16.50.2, 2 ms
```

```
Reply to request 0 from 172.16.50.6, 2 ms
```

```
Reply to request 0 from 172.16.50.10, 2 ms
```

Incident 10

(3 points)



- Users at the Remote Offices connected to Regional Office 1 are complaining of various issues.
- Troubleshoot and fix any issues preventing the Loopback interfaces of R21, R22, and R23 from talking to all devices in the EIGRP 100 domain.
- EIGRP 100 should have full reachability.
- This incident contains multiple faults.

Diagnostic Lab 1

Please look at the provided information and read through this entire lab before you start. Read the directions very carefully to make sure you are doing what is being asked of you. This is very important when you take Cisco's CCIE lab.

Each task may contain a large amount of information including diagrams, email chains, trouble tickets, device configs, and Wireshark captures. It is extremely important that you read through each piece of information before answering the task.

Each task will require you to provide an answer to the issues provided, based off of the information that is presented.

General Rules

- You do not have access to any equipment
- You are not required to configure any equipment
- Questions may be best selection, fill in the blank, multiple choice, order of operations, or best match.

Estimated Time to Complete: 30 Minutes

Total Possible Points: N/A

Pre-setup

- There is no pre-setup for this lab.
- No rack access is required for this lab.

Task #1

(3 points):

A new trouble ticket has been escalated to you. The following information has been provided to help with understanding the issue. Diagnose and identify the issue using the information provided.

Email chain between customer and helpdesk

From: Jiminy Cricket
Sent: Wednesday, August 13, 2014 9:17 AM
To: iPexpert Helpdesk
Subject: EIGRP Config Tuning – HELP!

Hi,

We recently attempted to tune our EIGRP config to have faster convergence times and better route summarization. After completing the configs, our Iowa offices have very intermittent connectivity. We do not have backup configs and need help to figure out what is causing this issue.

Jiminy Cricket
IT Manager, Acme Corp.
Direct: 111-111-1111
E-mail: jiminy.cricket@acme.com

From: iPexpert Helpdesk
Sent: Wednesday, August 13, 2014 9:23 AM
To: Jiminy Cricket
Subject: RE: EIGRP Config Tuning – HELP!

Mr. Cricket,

We would love to assist with this issue. We have opened up a ticket named Incident 1 for internal tracking. In order to better help, please provide the following:

1. A network diagram that shows the EIGRP connectivity
2. The router configs of the devices having issues
3. Perform and capture a “show log” on each router

Once we have the above information, we will review, assign an engineer, and get back to you.

Tom Clancy
HelpDesk Representative
Office: 999-999-9999 | helpdesk@ipexpert.com

From: Jiminy Cricket
Sent: Wednesday, August 13, 2014 9:35 AM
To: iPexpert Helpdesk
Subject: EIGRP Config Tuning – HELP!

The information requested has been attached. I am unable to connect steadily to the Iowa routers due to this issue, so I am not able to get the debugs from R2 or R3. I attached the debug from R1. Please understand that this is an urgent matter and if there is any way we can expedite this request, I would really appreciate it.

Jiminy Cricket
IT Manager, Acme Corp.

Direct: 111-111-1111

E-mail: jiminy.cricket@acme.com

From: iPexpert Helpdesk

Sent: Wednesday, August 13, 2014 9:55 AM

To: Jiminy Cricket

Subject: RE: EIGRP Config Tuning – HELP!

Mr. Cricket,

This incident has been assigned to our top tier Network Engineer for review. You should hear something back very soon. Thank you for your patience.

Tom Clancy

HelpDesk Representative

Office: 999-999-9999 | helpdesk@ipexpert.com

Router Configuration

- **R1**

```
R1#sh run
Building configuration...

Current configuration : 2963 bytes
!
version 15.4
no service timestamps debug uptime
no service timestamps log uptime
no service password-encryption
!
hostname R1
!
!
no aaa new-model
clock timezone CET 1 0
!
!
no ip domain lookup
ip cef
no ipv6 cef
!
multilink bundle-name authenticated
!
!
redundancy
!
interface Loopback111
 ip address 192.168.1.1 255.255.255.0
!
interface Ethernet0/0
 ip address 10.10.1.1 255.255.255.0
```

```
ip summary-address eigrp 2001 192.168.0.0 255.255.254.0
ip hello-interval eigrp 2001 120
!
interface Ethernet0/1
no ip address
!
router eigrp 2001
variance 10
network 10.10.1.1 0.0.0.0
network 192.168.1.0
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
control-plane
!
line con 0
exec-timeout 0 0
logging synchronous
line aux 0
line vty 0 4
login
transport input none
```

▪ **R2**

```
R2#sh run
Building configuration...

Current configuration : 2926 bytes
!
!
version 15.4
no service timestamps debug uptime
no service timestamps log uptime
no service password-encryption
!
hostname R2
!
!
no aaa new-model
clock timezone CET 1 0
!
no ip domain lookup
ip cef
no ipv6 cef
!
multilink bundle-name authenticated
```

```
!  
!  
interface Loopback222  
  ip address 192.168.2.2 255.255.255.0  
!  
interface Ethernet0/0  
  no ip address  
!  
interface Ethernet0/1  
  no ip address  
!  
interface Ethernet0/1.1  
  encapsulation dot1Q 12  
  ip address 10.10.1.2 255.255.255.0  
  ip hold-time eigrp 2001 60  
  
!  
interface Ethernet0/1.2  
  encapsulation dot1Q 23  
  ip address 10.10.2.2 255.255.255.0  
!  
router eigrp 2001  
  variance 12  
  network 10.10.1.2 0.0.0.0  
  network 10.10.2.2 0.0.0.0  
  redistribute connected  
!  
ip forward-protocol nd  
!  
!  
no ip http server  
no ip http secure-server  
!  
ip access-list extended EIGRP  
  deny ip 172.16.0.0 0.0.255.255 any  
  deny ip 10.10.0.0 0.0.255.255 any  
  permit ip any any  
!  
control-plane  
!  
line con 0  
  exec-timeout 0 0  
  logging synchronous  
line aux 0  
line vty 0 4  
  login  
  transport input none
```

▪ **R3**

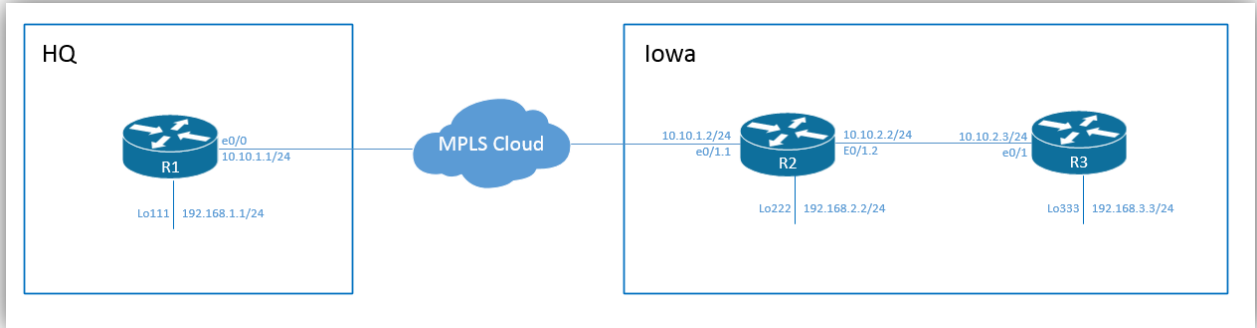
```
R3#sh run
```

Building configuration...

```
Current configuration : 2926 bytes
!
!
version 15.4
no service timestamps debug uptime
no service timestamps log uptime
no service password-encryption
!
hostname R3
!
boot-start-marker
boot-end-marker
!
!
no aaa new-model
clock timezone CET 1 0
!
!
no ip domain lookup
ip cef
no ipv6 cef
!
multilink bundle-name authenticated
!
interface Loopback333
 ip address 192.168.3.3 255.255.255.0
!
interface Ethernet0/0
 no ip address
!
interface Ethernet0/1
 ip address 10.10.2.3 255.255.255.0
!
!
router eigrp 2001
 network 10.0.0.0
 redistribute connected
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
control-plane
!
line con 0
 exec-timeout 0 0
 logging synchronous
line aux 0
```

```
line vty 0 4
login
transport input none
```

EIGRP Topology



Router Debugs

```
R1#
06:42:15: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is down: holding time expired
06:42:17: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is up
06:43:17: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is down: holding time expired
06:43:19: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is up
06:44:19: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is down: holding time expired
06:44:21: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is up
06:45:21: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is down: holding time expired
06:45:23: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is up
06:46:23: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is down: holding time expired
06:46:25: %DUAL-5-NBRCHANGE: IP-EIGRP 2001: Neighbor 10.10.1.2 (Ethernet0/0)
is up
```

Using the information provided, select the most logical cause of the issue from the list below:

- The network statements are not correctly configured on R2
- The network summary from R1 is configured with the wrong subnet mask.

- The EIGRP hello-interval on R1 is set smaller than the hold-interval on R2, causing the neighbors to flap continuously.
- R1 has an ACL that filters and drops EIGRP packets.
- The EIGRP hold-interval on R2 is set smaller than the hello-interval on R1, causing the neighbors to flap continuously.
- R2 has an ACL that filters and drops EIGRP packets.
- R1 int e0/0 has the wrong subnet mask.
- The EIGRP hold-interval on R2 is set smaller than the hello-interval on R2, causing the neighbors to flap continuously.
- The variance is incorrectly set on R2.
- The network summary on R1 is configured with the wrong subnet mask.
- R2 int e0/0 has an incorrectly configured subnet mask.
- The EIGRP hello-interval on R2 is set smaller than the hold-interval on R1, causing the neighbors to flap continuously.
- The variance is incorrectly set on R1.

Task #2

(3 points):

You have been out on vacation for the last 2 weeks. While you were out, your company added an ISP using BGP. Your co-worker is having an issue trying to load balance across the 2 providers and does not understand the BGP path selection process. Review the information provided for a better understanding of the issue.

ISP Information

- ISP1
 - Public Subnet: 13.13.13.0/29
 - Gateway: 13.13.13.1
 - BGP AS 1313
 - Bandwidth: 100Mbps
 - Extended Communities: Yes
- ISP2
 - Public Subnet: 14.14.14.0/29
 - Gateway: 14.14.14.1
 - BGP AS 1414
 - Bandwidth: 100Mbps
 - Extended Communities: Yes

Using the information provided, explain to your co-worker the process that BGP takes to select the best path, in order from the top down (Order from 1-12).

- Put the following BGP route selection criteria in order from the top down.
 - Prefer eBGP over iBGP paths.
 - Prefer the path with the lowest origin type.
 - Prefer the route that comes from the BGP router with the lowest router ID.
 - Prefer the path with the highest WEIGHT.
 - Prefer the path that was locally originated via a network or aggregate BGP subcommand or through redistribution from an IGP.
 - Prefer the path that comes from the lowest neighbor address.
 - When both paths are external, prefer the path that was received first (the oldest one).
 - Prefer the path with the highest LOCAL_PREF.
 - Prefer the path with the minimum cluster list length.
 - Prefer the path with the lowest IGP metric to the BGP next hop.
 - Prefer the path with the shortest AS_PATH.

- Prefer the path with the lowest MED.

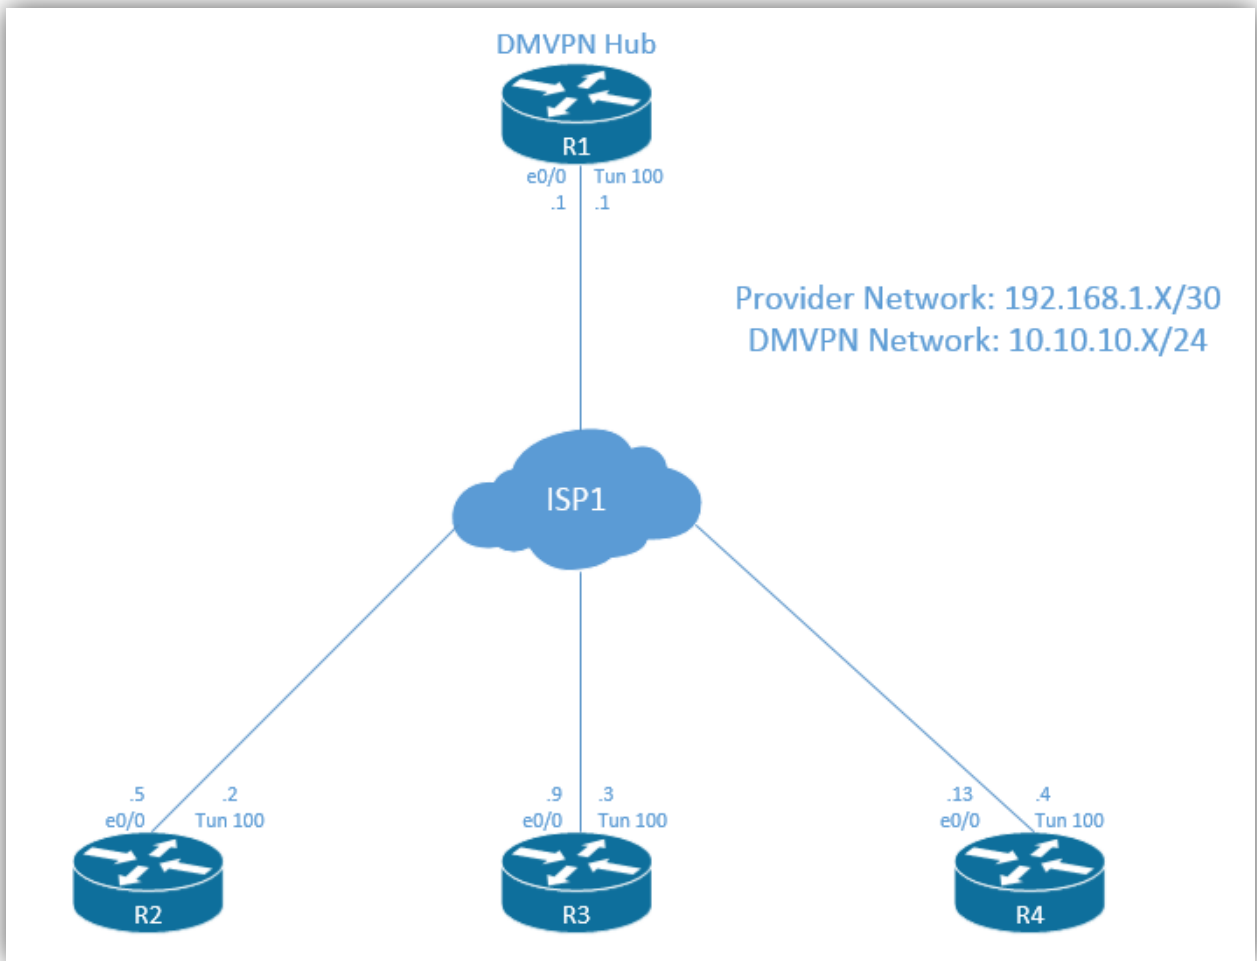
Task #3**(3 points):**

Users at the DMVPN spoke sites recently opened a trouble ticket that has been assigned to you. They are complaining that they cannot reach the other spokes sites, but can reach the DMVPN Hub.

Router Debugs

```
Router1#debug ip ospf events
OSPF events debugging is on
*Sep 8 02:58:36.069: OSPF: Interface FastEthernet0/0 going Up
*Sep 8 02:58:36.069: OSPF: Send hello to 224.0.0.5 area 0 on FastEthernet0/0
    from 10.10.10.1
*Sep 8 02:58:46.069: OSPF: Send hello to 224.0.0.5 area 0 on FastEthernet0/0
    from 10.10.10.1
*Sep 8 02:58:56.069: OSPF: Send hello to 224.0.0.5 area 0 on FastEthernet0/0
    from 10.10.10.1
*Sep 8 02:59:06.069: OSPF: Send hello to 224.0.0.5 area 0 on FastEthernet0/0
    from 10.10.10.1
*Sep 8 02:59:16.069: OSPF: end of Wait on interface FastEthernet0/0
*Sep 8 02:59:16.069: OSPF: DR/BDR election on FastEthernet0/0
*Sep 8 02:59:16.069: OSPF: Elect BDR 10.10.10.4
*Sep 8 02:59:16.069: OSPF: Elect DR 10.10.10.4
*Sep 8 02:59:16.069: OSPF: Elect BDR 0.0.0.0
*Sep 8 02:59:16.069: OSPF: Elect DR 10.10.10.4
*Sep 8 02:59:16.069:          DR: 10.10.10.4 (Id)   BDR: none
*Sep 8 02:59:16.069: OSPF: Send hello to 224.0.0.5 area 0 on FastEthernet0/0
    from 10.10.10.1
*Sep 8 03:25:46.084: %OSPF-5-ADJCHG: Process 100, Nbr 10.10.10.4 on
    FastEthernet0/0 from LOADING to FULL, Loading Done
Router1#no debug ip ospf events
```

OSPF Topology



Assuming R2 and R3 are configured correctly and using the information provided, choose the device that contains the configuration error and then choose the best method to fix the issue:

Device with Issue:

- R1
- R2
- R3
- R4
- ISP1

Area of Issue:

- Change OSPF network type
- Change OSPF priority to 1
- Change to Area 1
- Change OSPF cost to 100
- Change OSPF priority to 0

Configuration Lab 1

Please look at the provided diagrams and read through the whole lab before you start. Read the directions very carefully to make sure you are doing what is being asked of you. This is very important when you take Cisco's CCIE lab.

Multiple topology diagrams are available for this lab. Be sure to understand each diagram and the information being conveyed.

General Rules

- All IPv4 address are pre-configured except SVI, tunnel, and sub-interfaces, unless otherwise noted.
- All Service Provider routers are pre-configured and cannot be accessed during the lab.
- Do not modify any IP addressing on any interfaces.
- The BB routers are not accessible.
- Static/default routes are NOT allowed unless otherwise stated in the task.
- Save your configurations often.

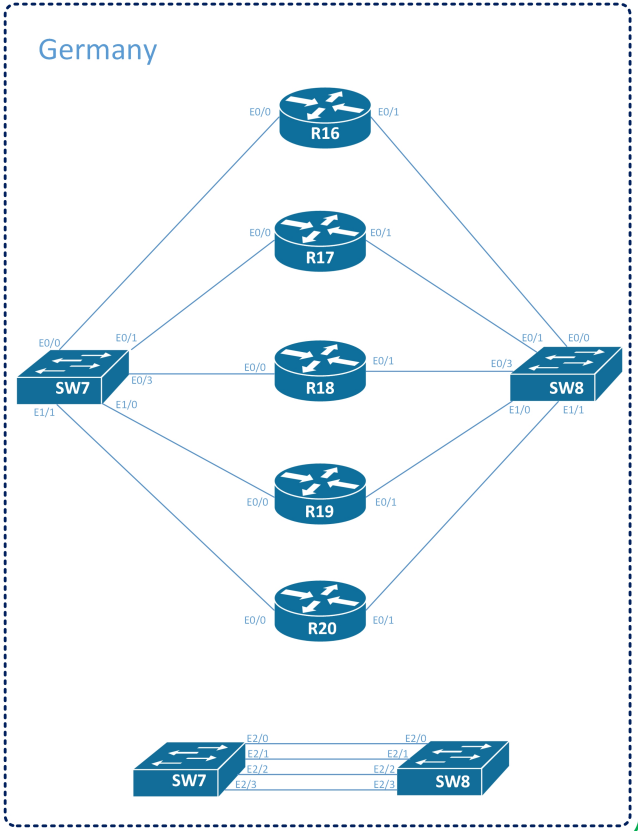
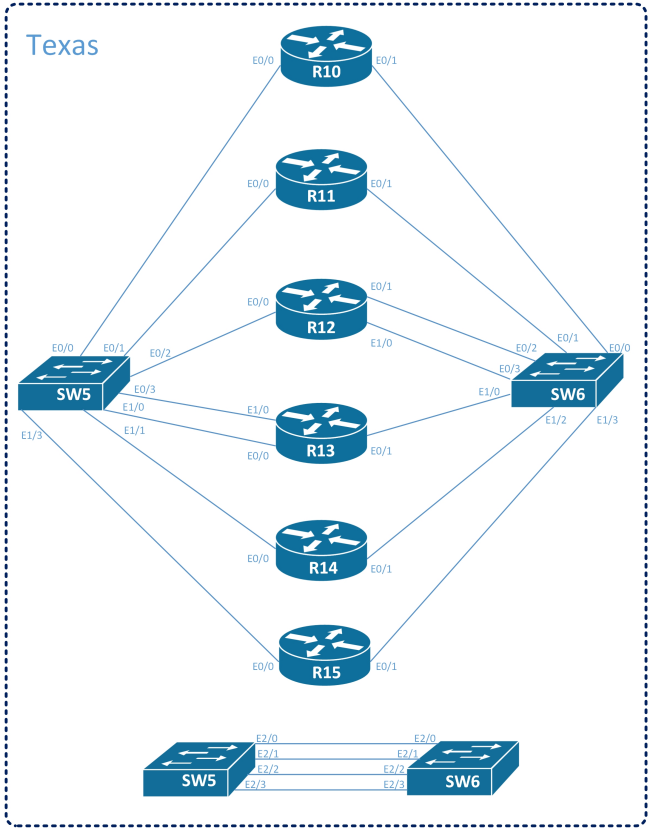

Estimated Time to Complete: **5.5-6 Hours**

Total Possible Points: **76**

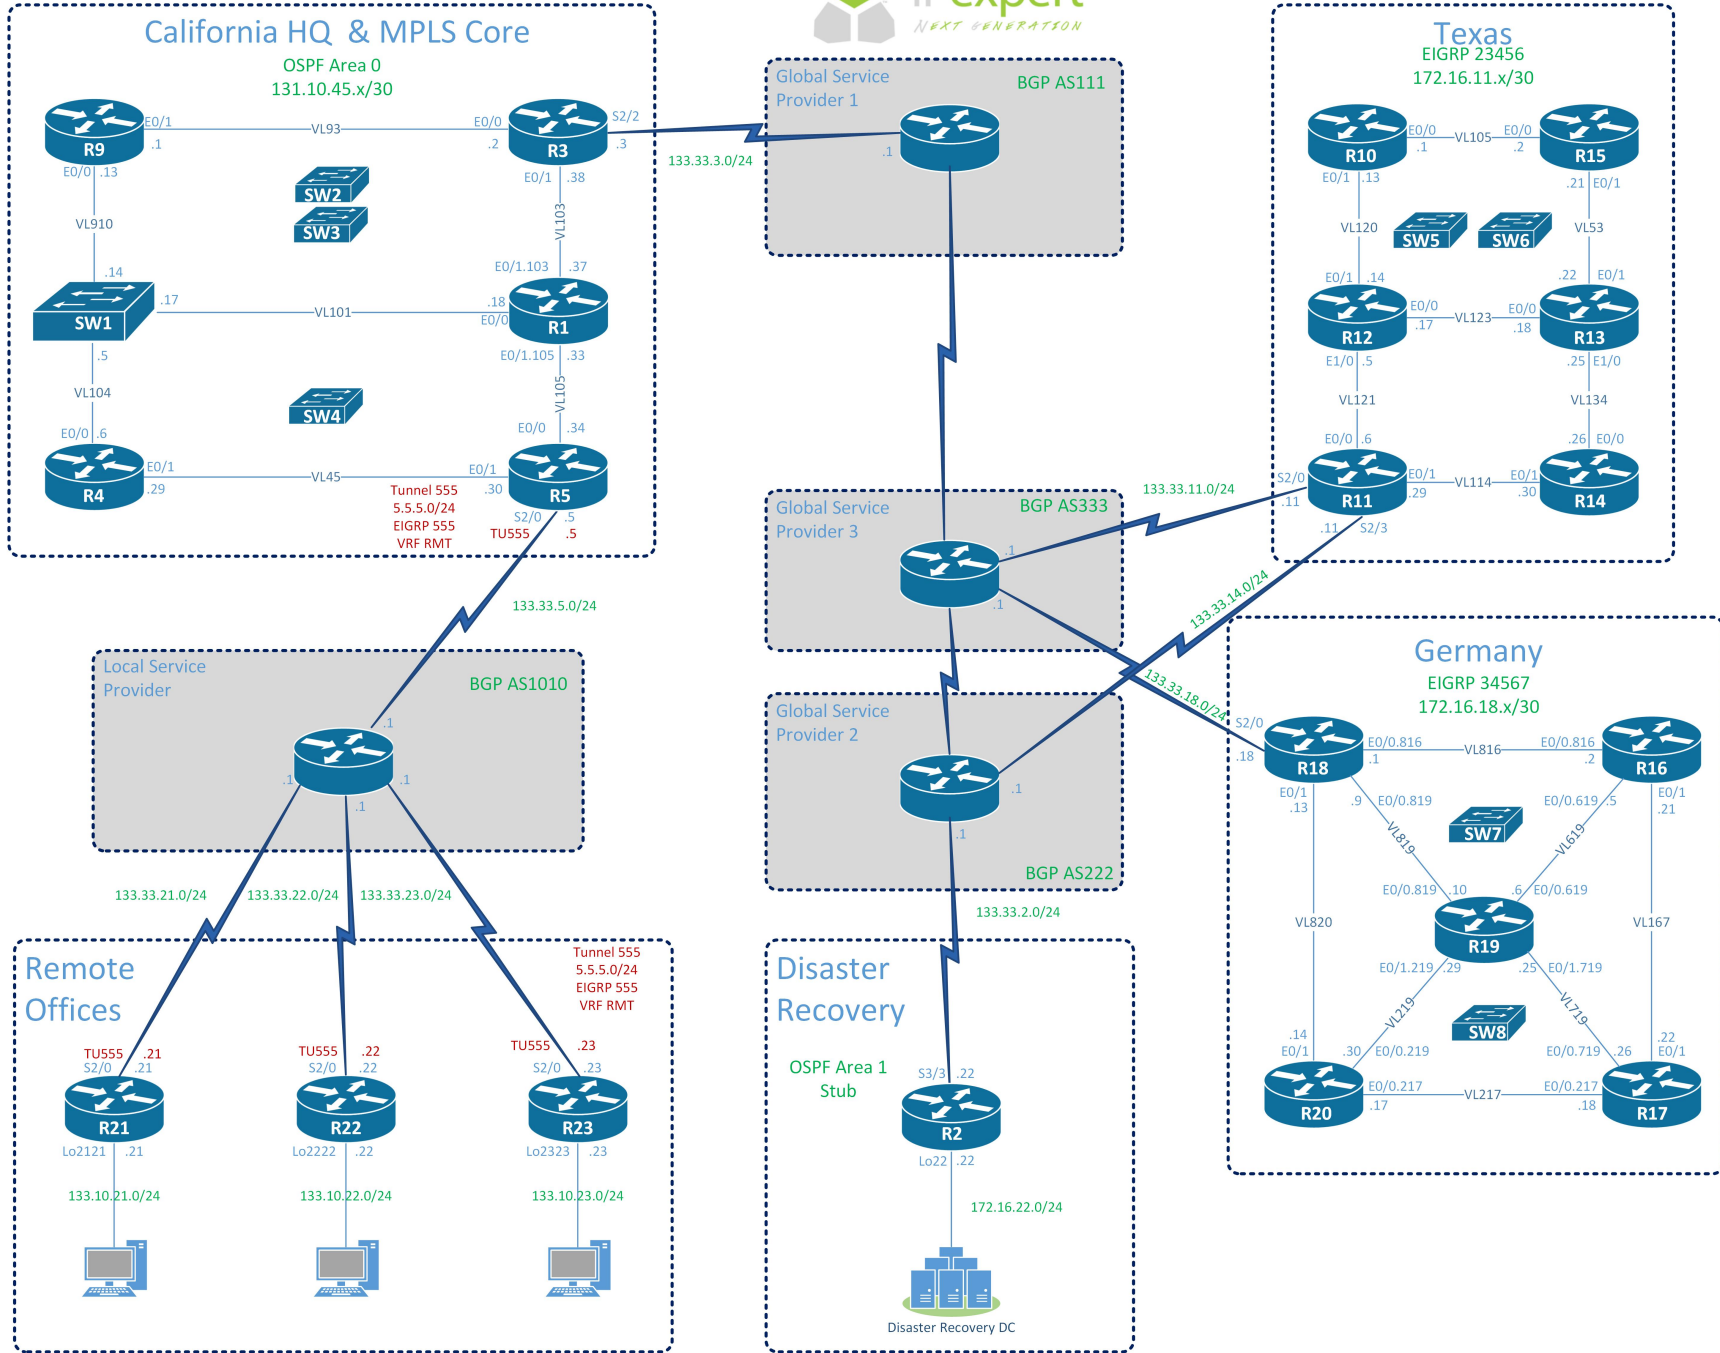
Pre-setup

- Please login to your vRack at ProctorLabs.com and load the initial Configuration.
- This lab is intended to be used with online rack access provided by our partner Proctor Labs (www.proctorlabs.com). Connect to the terminal server and complete the troubleshooting tasks as detailed below.

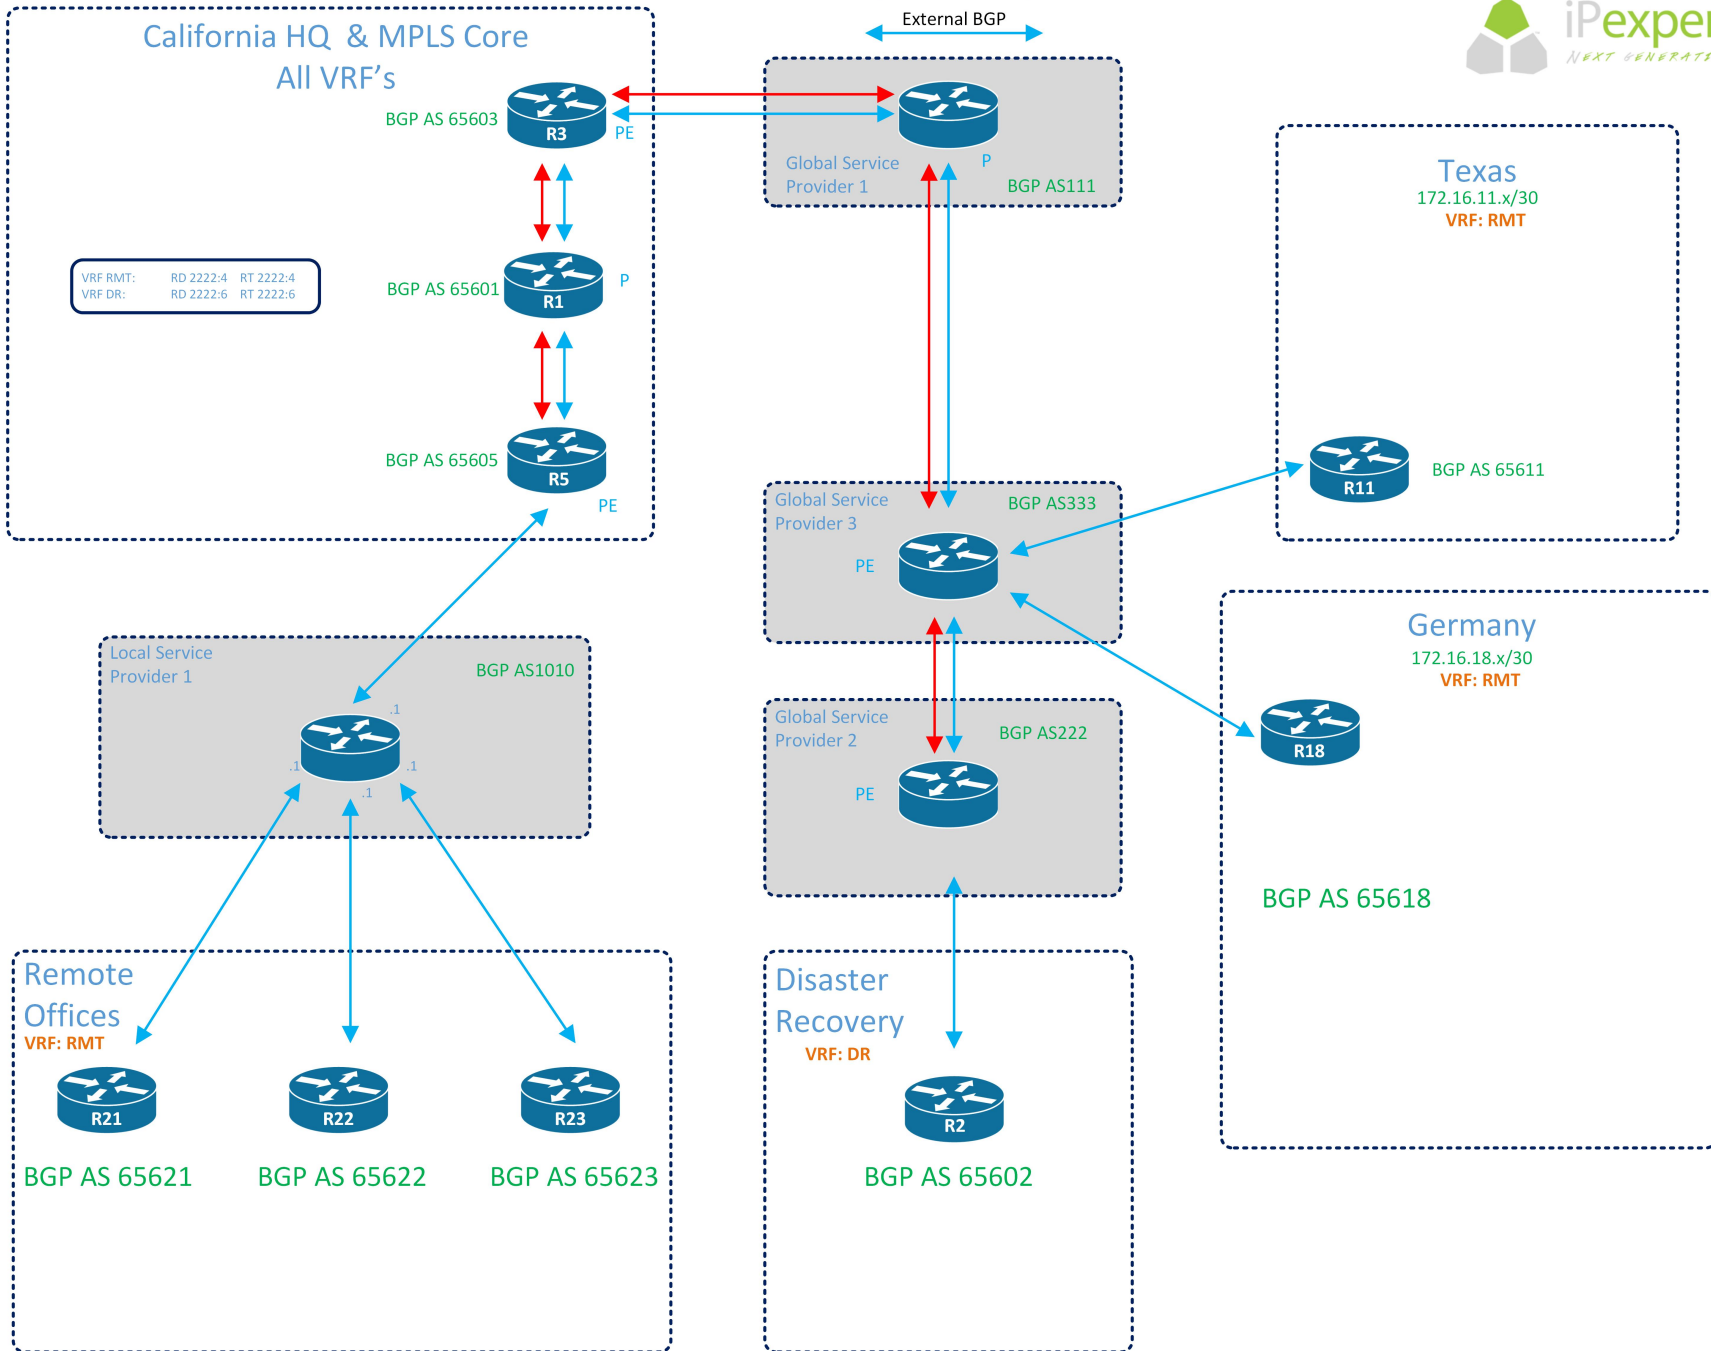
Layer 2



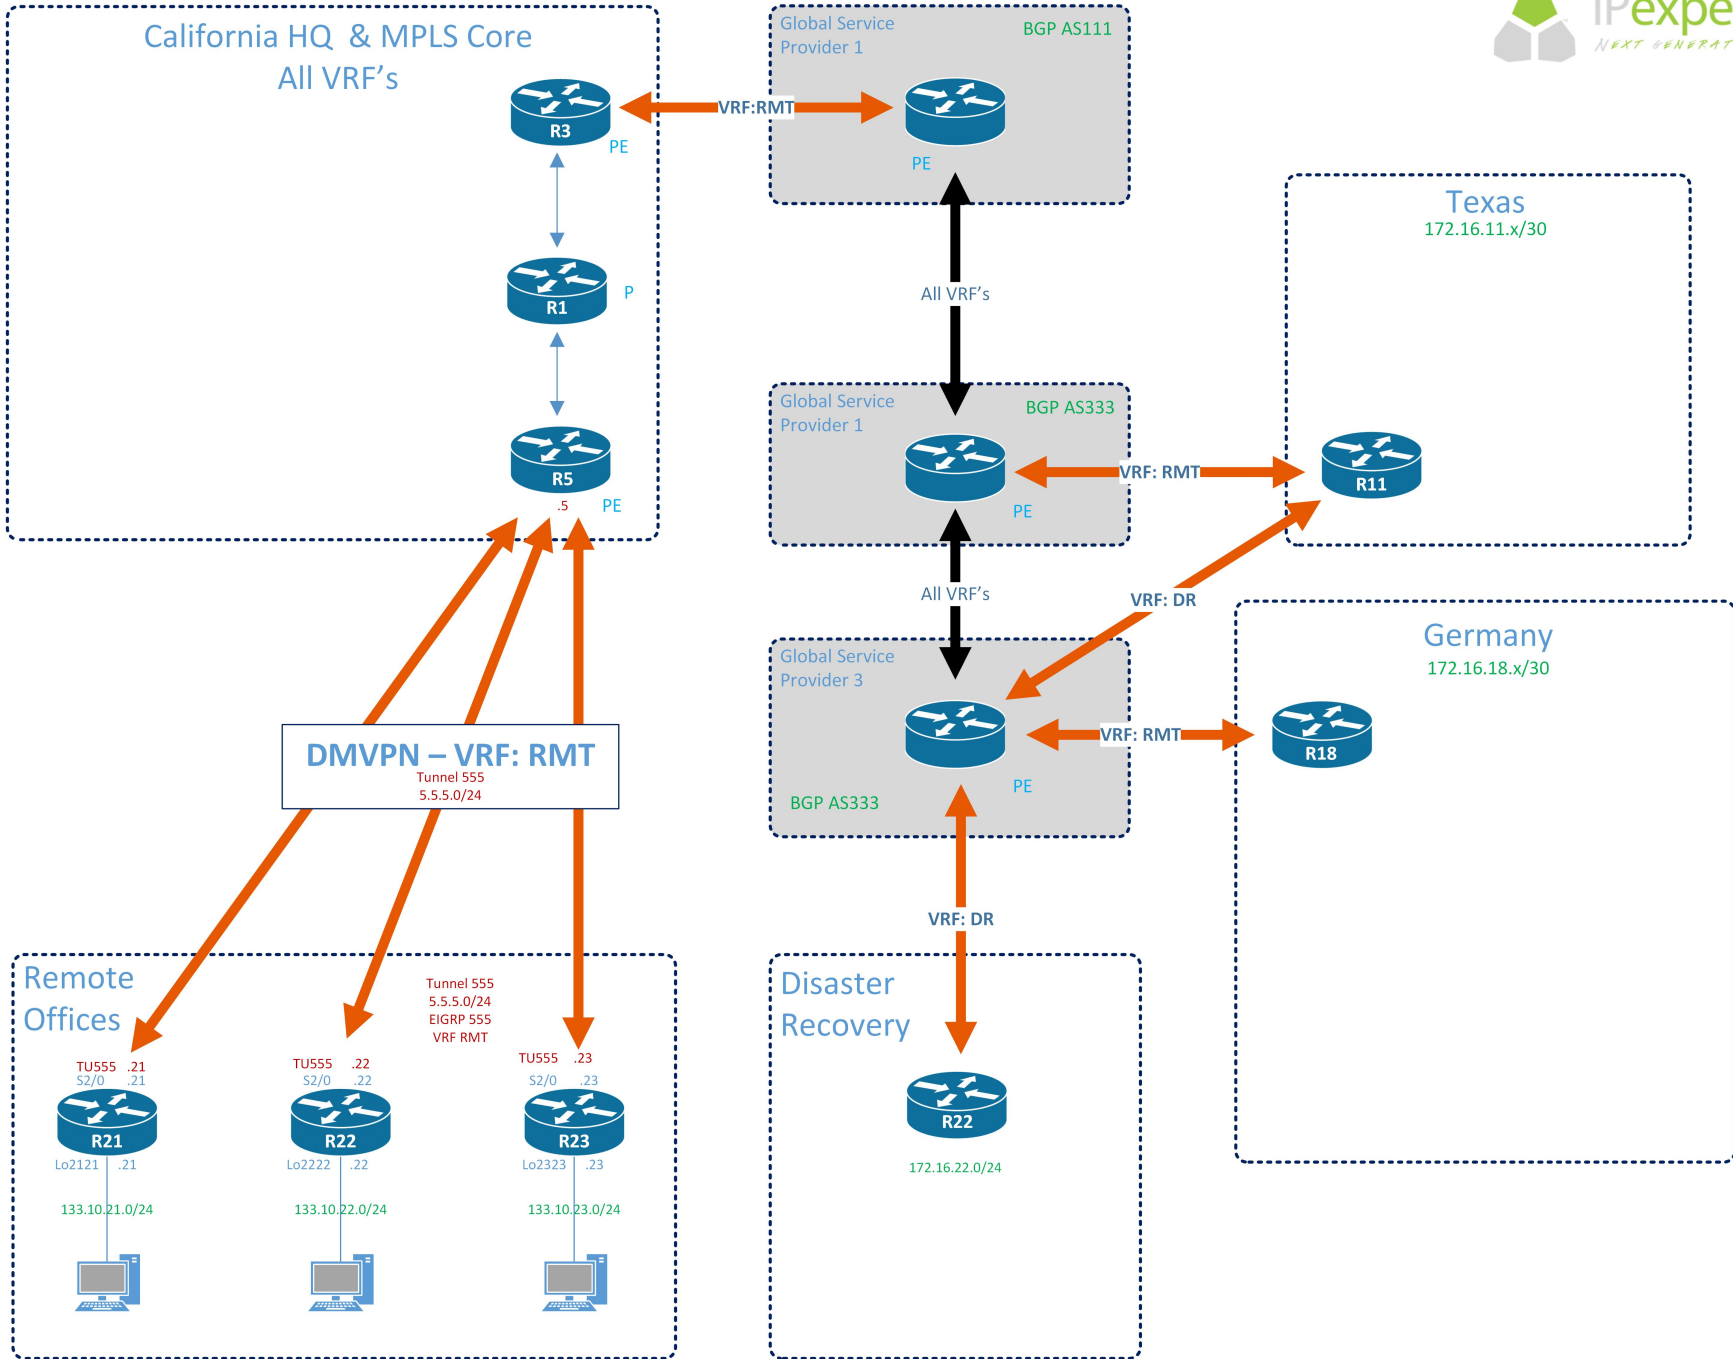
IPv4

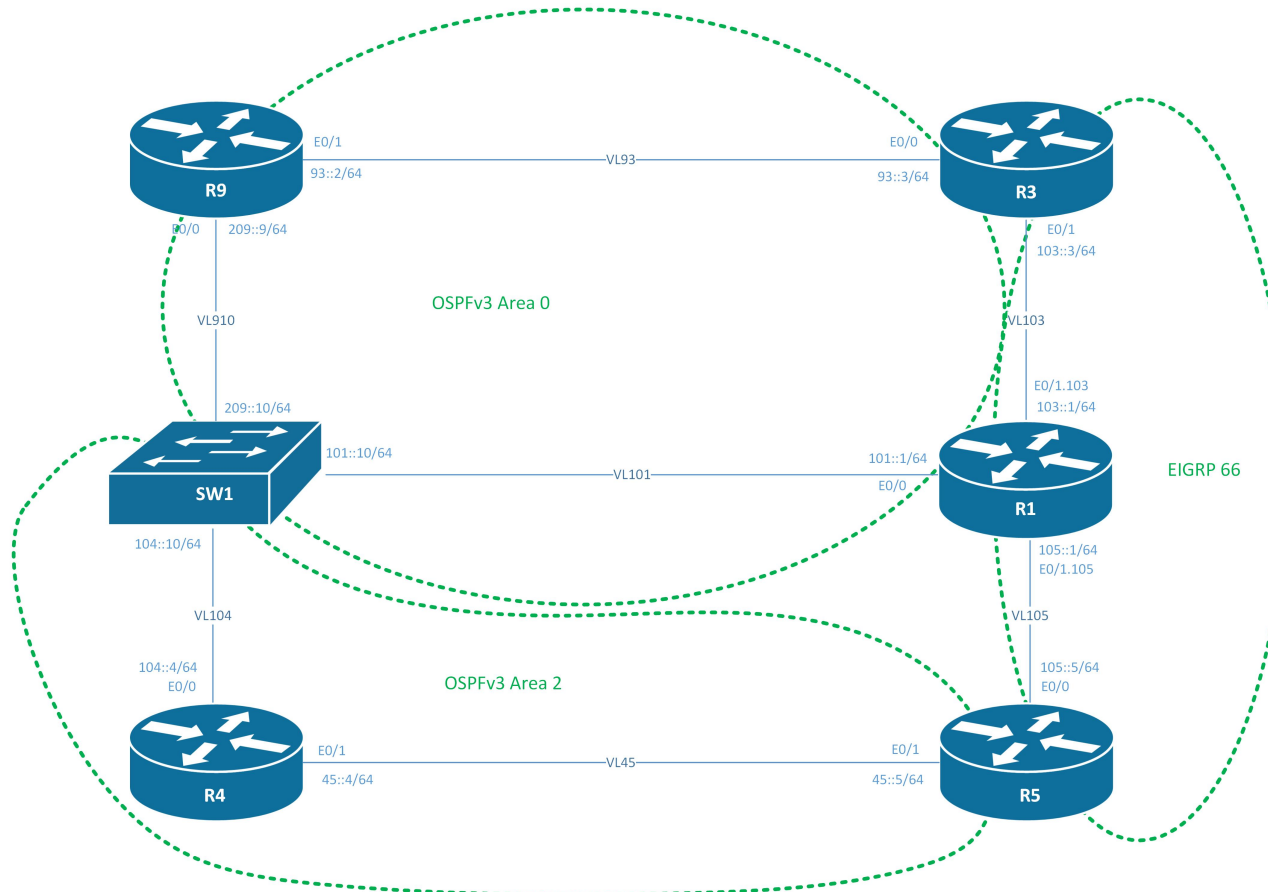


BGP



IPv4 VPN





Layer 2 Technologies

(14 points)

- California HQ Switch-to-Switch Links (2 points)
 - Using the Layer 2 diagram, configure the switch-to-switch links on SW1, SW2, SW3, and SW4 as dot1q trunks
 - Make sure that the trunk configuration is not negotiated.
 - Only allow the VLANs specified for California HQ across the trunks.
 - Create an ether-channel on interfaces 3/0 and 3/1, and 3/2 on both SW1 and SW2 using a Cisco proprietary technology.
 - Make sure the ether-channel is negotiated and that SW1 initiates the negotiation.
- California HQ VLANs (2 points)
 - Configure the necessary VLAN's at California HQ on SW1.
 - Use VTP domain name CCIE.
 - Use VTP password iPexpert!
 - SW1 should always be the VTP be the master. All other switches should be set to client.
 - Do not configure any VLAN's on SW2, SW3, or SW4. They should learn the VLAN's from the VTP server.
- California HQ Router Links (2 points)
 - Using the Layer 2 diagram, configure all interfaces connected to a router as an access port at California HQ. All Trunk links to routers should use 802.1Q and only allow the VLAN's needed.

Device	Interface	VLAN
R1	E0/0	VL101
	E0/1	Trunk
R3	E0/0	VL93
	E0/1	VL103
R4	E0/0	VL104
	E0/1	VL45
R5	E0/0	VL105
	E0/1	VL45
R9	E0/0	VL910
	E0/1	VL93

- Configure the SVI interfaces on both SW1 as referenced in the IPv4 diagram.
- California HQ Spanning-Tree (3 points)

- Use 802.1S.
- All even VLANs should belong to instance 10; all odd VLANs should belong to instance 11.
- SW1 should be the root bridge for all odd VLANs and the secondary root bridge for all even VLANs.
- SW2 should be the primary root bridge for all even VLANs and the secondary root bridge for all odd VLANs.
- The primary and secondary bridges should be statically set to protect against other switches becoming the root bridge.
- Texas VLANs (2 points)
 - Configure the necessary VLANs in Texas on SW5 and SW6.
 - Do not use VTP to accomplish this.
 - All switches should pass VTP information but not use it to configure VLANs.
- Texas Router Links (2 points)
 - Using the Layer 2 diagram, configure all interfaces connected to a router as an access port in Texas.
 - Assign the appropriate VLAN to each link.

Device	Interface	VLAN
R10	E0/0	VL105
	E0/1	VL120
R11	E0/0	VL121
	E0/1	VL114
R12	E0/0	VL123
	E0/1	VL120
	E1/0	VL121
R13	E0/0	VL123
	E0/1	VL53
	E1/1	VL134
R14	E0/0	VL134
	E0/1	VL114
R15	E0/0	VL105
	E0/1	VL53

- Texas Switch-to-Switch Links (1 points)
 - Configure the switch-to-switch links on SW5 and SW6 as dot1q trunks.
 - Make sure that the trunk links are not negotiated.
 - All trunks should use 802.1q as the encapsulation protocol.

- Only allow the VLAN's required for the Texas layer 2 topology across the switch-to-switch trunk links.
- Create an ether-channel, bonding E2/0-3 between the switches. Use LACP to accomplish this task.
- Verify Connectivity
 - SW7 and SW8 in Germany have been pre-configured. Verify the configuration and test connectivity between directly connected devices.
 - Verify all directly connected devices can ping each other in Germany, Texas, and California HQ.

IP Routing

(31 Points)

- OSPF at California HQ (4 points)
 - Create OSPF Area 0 in California HQ.
 - All devices should use the same OSPF process ID.
 - Add all interfaces to the OSPF process except the serial links that leave the Autonomous System.
 - Restrict OSPF to these interfaces without using the passive-interface feature.
 - All addresses in the OSPF domain should be reachable by all devices California HQ.
- EIGRP AS 23456 (4 points)
 - Create EIGRP AS 23456 in Texas.
 - Add all interfaces in Texas to the EIGRP process except those connected to other Autonomous Systems.
 - All EIGRP adjacencies should be authenticated using MD5 and the password CCIERock\$.
 - All subnets included in EIGRP 23456 should be reachable from every device in the AS, including the Loopback interface of each router.
 - Use EIGRP wide metrics.
- EIGRP AS 34567 (3 points)
 - Add all interfaces in Germany to the EIGRP process except those connected to other Autonomous Systems.
 - R17 and R16 should not form an adjacency with each other.
 - All subnets included in EIGRP 34567 should be reachable from every device in the AS.
 - Use EIGRP wide metrics.
- BGP AS 65611 (3 points)

- On R11, create the EBGP peering with BGP AS333 and AS222.
- Use directly connected interfaces for the peerings.
- On R11, mutually redistribute BGP and EIGRP.
- Do not advertise any routes learned from BGP back into BGP.
- BGP AS 65618 (2 points)
 - On R18, create the EBGP peering with BGP AS333.
 - Use directly connected interfaces for the peering addresses.
 - On R18, mutually redistribute BGP and EIGRP.
- BGP 65602 (2 points)
 - Create the EBGP peerings between AS65602 and AS222.
 - Use the directly connected interfaces to make these peerings.
 - On R2, redistribute Loopback 22 into BGP, but do not redistribute any other routes.
 - R2 should only be reachable by the 172.16.11.X networks and all loopback 0 router interfaces in Texas. Perform this filtering on R11.
- BGP AS 65621, 65622, 65623 (3 points)
 - Create the EBGP peerings from AS65621, 65622, and 65623 to AS1010.
 - Use the directly connected serial interfaces to make these peerings.
 - Do not perform any redistribution in these AS's
 - The local Service Provider is not advertising any routes except for those that are directly connected to it.
- IPv4 Multicast (4 points)
 - Create PIM Sparse-Dense-Mode neighbor relationships between R1, R3, R5, and SW1.
 - R3 should be used as the rendezvous point.
 - e0/0 of R3, E0/0 of R5, and e0/1 of R1 should all act as multicast clients of the group 224.5.5.5.
 - SW1 should be able to ping the multicast group and get a response from all 3 multicast clients.
 - Use the global routing table to accomplish this task.
- IPv6 (6 Points)
 - Assign the IPv6 addressing according to the following table:

Device	Interface	IPv6 Address
R1	E0/0	2001:101::1/64
	E0/1.103	2001:103::1/64

	E0/1.105	2001:105::1/64
R9	E0/0	2001:209::9/64
	E0/1	2001:93::9/64
R3	E0/0	2001:93::3/64
	E0/1	2001:103::3/64
R4	E0/0	2001:104::4/64
	E0/1	2001:45::4/64
R5	E0/0	2001:105::5/64
	E0/1	2001:45::5/64
SW1	VLAN 101	2001:101::10/64
	VLAN 910	2001:209::10/64
	VLAN 104	2001:104::10/64

- Configure EIGRP AS66.
 - Only add the interfaces that are in the IPv6 diagram.
- Configure OSPFv3 Area 0 and Area 2.
 - Only add the interfaces that are in the IPv6 diagram.
- Perform mutual redistribution on R1 between OSPF and EIGRP.
- Verify all devices can reach all assigned IPv6 address in the IPv6 diagram.

IPv4 VPN

(21 Points)

- MPLS VPN (5 points)
 - R5 is a PE router for Tunnel 555/VRF:RMT.
 - Enable MPLS on the necessary links in the MPLS backbone using the IPv4 VPN diagram.
 - Enable LDP routing and create the necessary LDP peerings.
 - Only enable MPLS on interfaces specified in the IPv4 VPN diagram.

- Enable VPNv4 BGP peerings between all PE and P routers as specified in the BGP diagram.
- Global Providers 1, 2, and 3 have agreed to pass VRF VPN information. Treat Global Service Provider 1 as a PE Router coming from R3.
- Global Service Provider 2 is a PE for VRF:DR.
- Global Service Provider 3 is a PE for VRF:RMT and VRF:DR.
- BGP California HQ (9 points)
 - Create all BGP ipv4 and vpnv4 peerings in California HQ.
 - Use the IP address of Loopback 0 on each device to form the adjacencies.
 - Unicast routing should be disabled under the main BGP process.
 - Create the EBGP peering to AS111 in VRF RMT. ISP1, ISP2, and ISP3 running MPLS and contain the same VRF information through the provider cloud. The ISP peering between Global Service Provider 1 and R3 should be made with the directly connected serial link, not the loopback 0 interface.
 - Create the EBGP peering to AS1010 in VRF RMT.
 - Create the VRFs listed on the BGP diagram on R1, R3, and R5 using the RDs and RTs listed.
 - The 172.16.22.0/24 subnet at the Disaster Recovery site requires connectivity to the Remote Offices in Texas ONLY. It should not be able to reach the Remote offices or Germany.
 - The Remote Offices, Germany, and Texas should have full reachability to all devices in the RMT VRF.
- DMVPN (7 points)
 - R5 is a DMVPN hub for connectivity to the Remote Offices in California. Use interface Tunnel 555 on R5, R21, R22, and R23.
 - Place these interfaces in VRF RMT on R5.
 - Assign the addressing specified in the IPv4 diagram for the tunnel.
 - After initial connectivity is established, Peer R21, R22, and R23 with R5 in EIGRP 555 via their tunnel interfaces and advertise the Loopback 2X2X interfaces into EIGRP.
 - Use EIGRP wide metrics.
 - Mutually redistribute routes between EIGRP and BGP on R5 in VRF RMT.
 - MTU on all tunnel interfaces should be set to allow for the VPN overhead.
 - Verify that the all IP addresses that are directly connected to R21, R22, and R23 are reachable from Texas and Germany. The serial interfaces do not need reachability.

IP Security

(5 Points)

- Access Control (2 points)
 - R16 should act as a HTTP server.
 - HTTP should use port 8000 for access.
 - Only allow this HTTP traffic through interface E0/1. It should be blocked on all other interfaces.
 - Only the e0/1 interfaces of R20 should be allowed to access R16 via HTTP port 8000.
 - Telnet from the e0/1 interface of R20 on port 8000 on R16 to verify connectivity.
- Unicast Source Verification (3 points)
 - R12 should allow packets arriving from R13 only if the packets are sourced from the networks that are listed in the R12 routing table. Apply the solution on the R12 interface that is connected to R13. Do not use any prefix-based filtering techniques.

IP Services

(5 Points)

- Quality of Service (3 points)
 - All traffic sourced from any 172.16.18.x network and leaving the Autonomous System should be marked with DSCP 2 and given 30% of the bandwidth in a priority queue.
 - This policy should be applied outbound on R18.
- Address Administration (2 points)
 - Configure SW1 as a DHCP server.
 - Assign the subnet of 2.2.2.0/24.
 - .100 should be the only IP that is generated. All other IP's should be reserved.
 - .2 should be the gateway.
 - 8.8.8.8 is the DNS server.
 - Interface e0/0 on R2 is pre-configured to use this IP pool. From R2, verify that e0/0 pulled ip address 2.2.2.100 and that it can ping 2.2.2.2. This does not need to be reachable from the rest of the network.

You have completed this Lab.

For Verification of your work, please refer to this Workbook's accompanying Detailed Solutions Guide. If you need assistance with any of this book's content, please visit our Member Community at <http://community.ipexpert.com> .