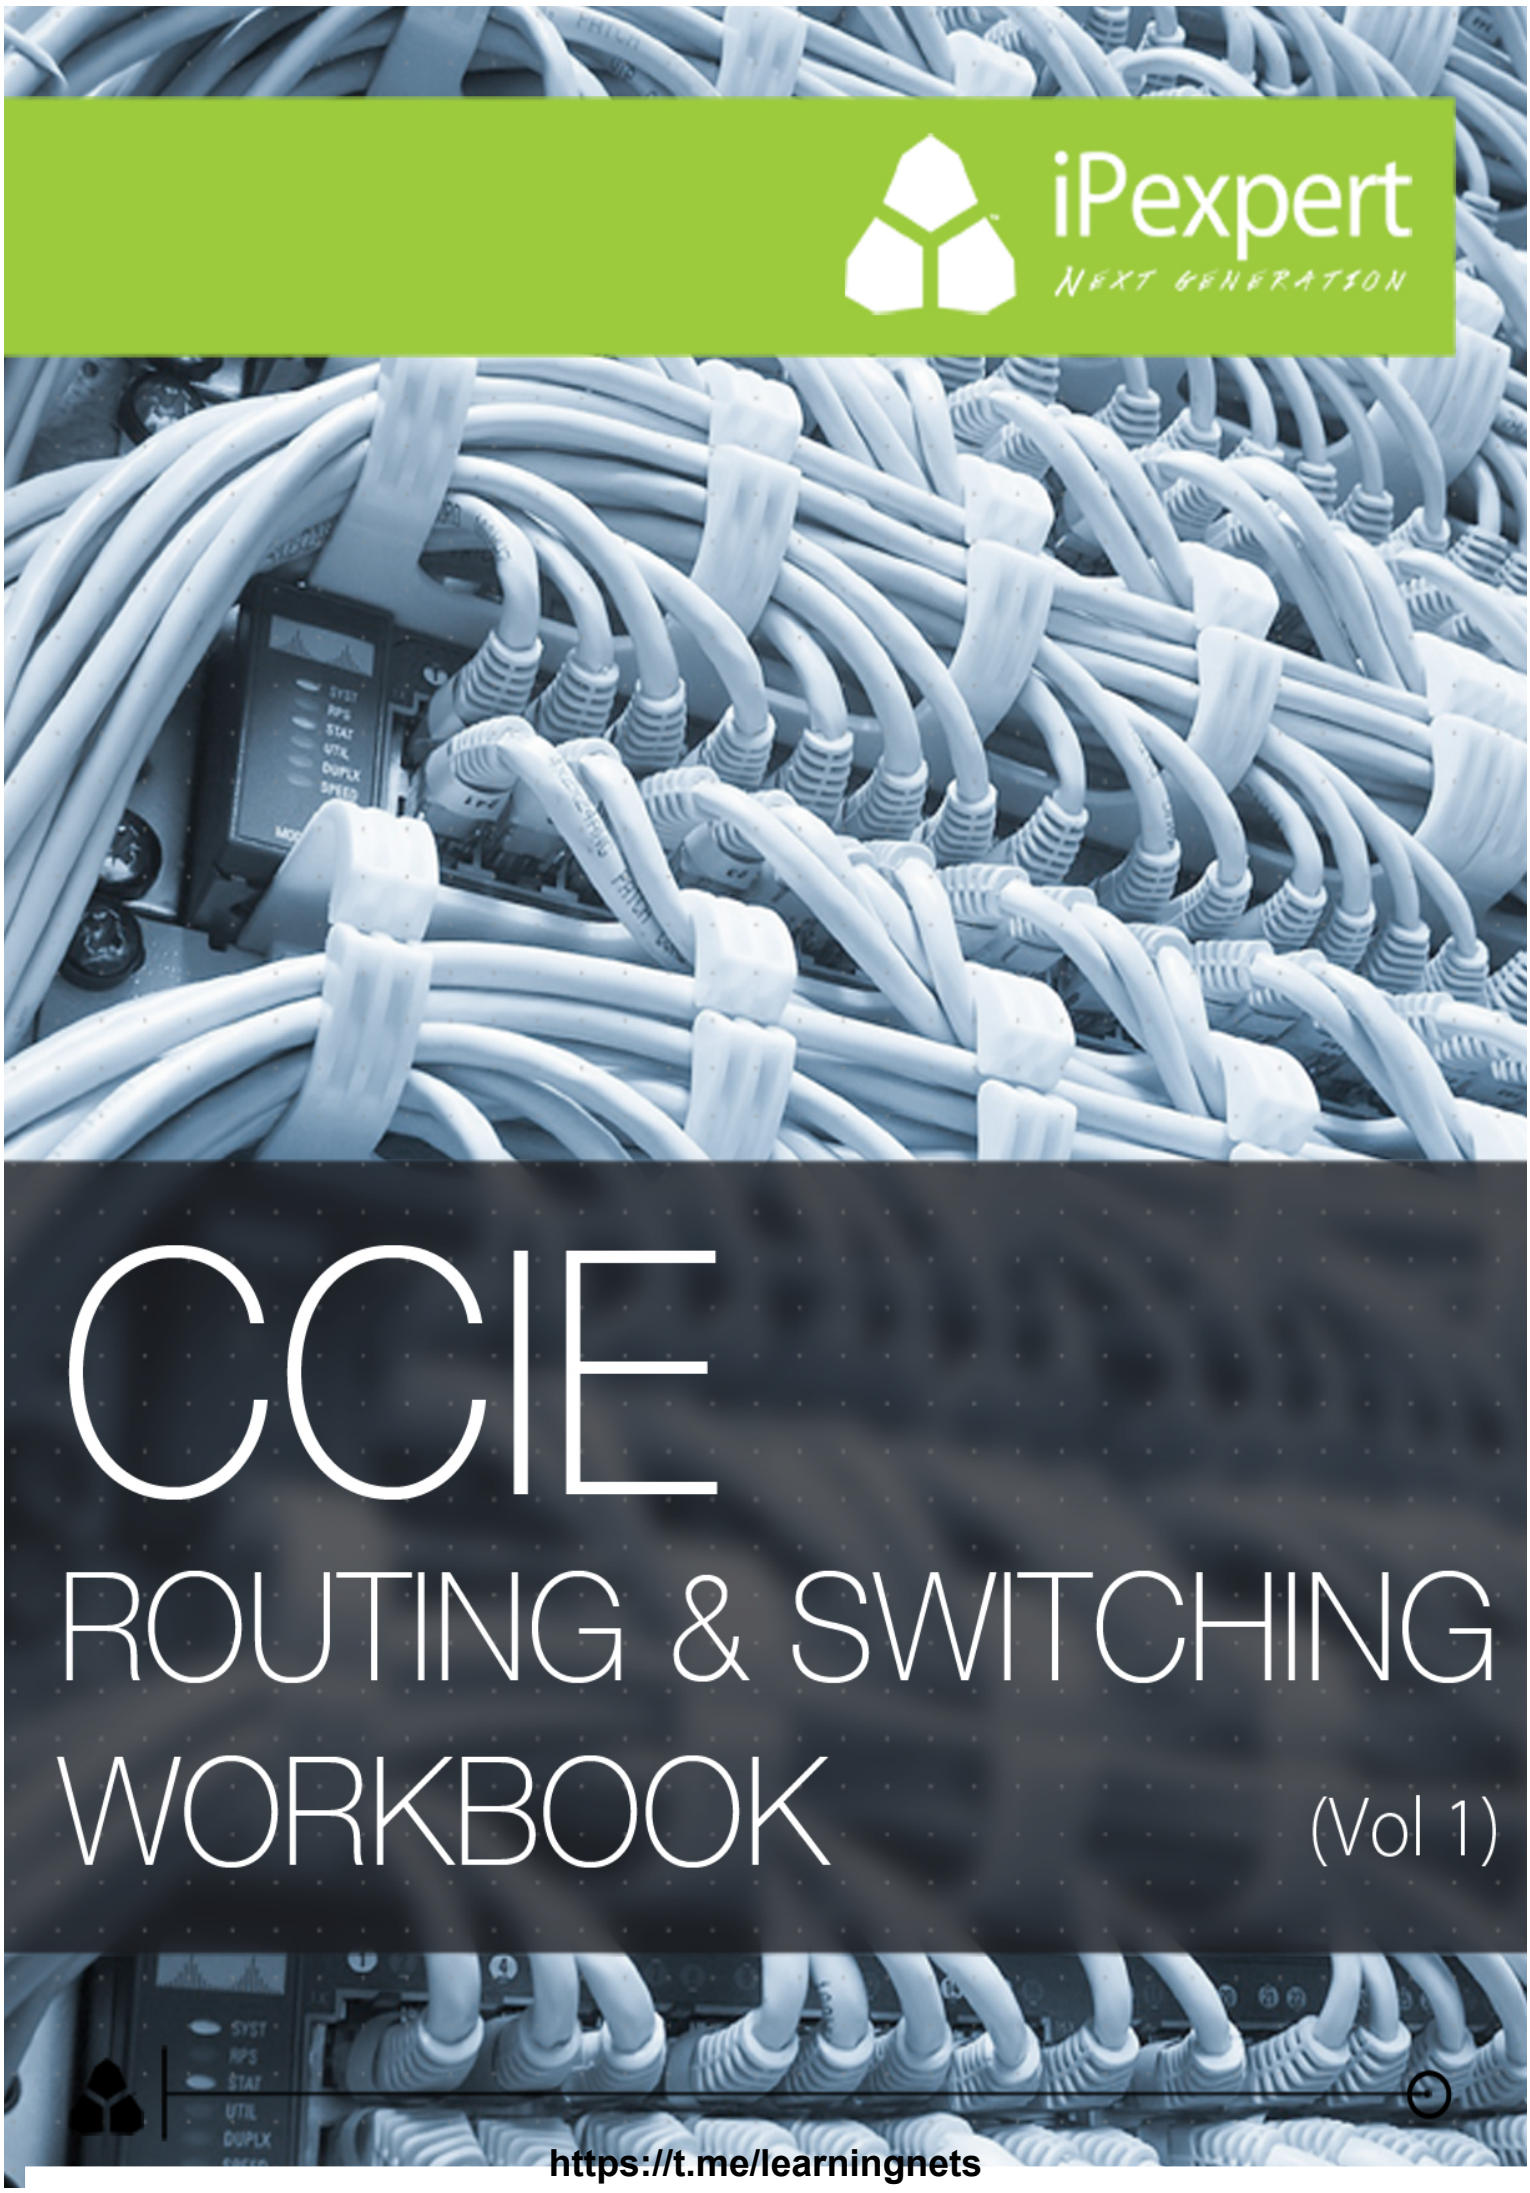

Table of Contents

Lab 1: Configure and troubleshoot switch port modes

Technologies covered

- CDP
- Access ports
- VLAN database
- VLAN
- Trunking
- dot1Q
- Native VLAN
- Manual pruning
- Layer 3 native interfaces
- SVIs
- Router-on-a-stick

Overview

You have been tasked to configure the layer 2 part of the network and to enable the routing between 2 VLANs in a router-on-a-stick topology.

The topology used in the lab will be the following:

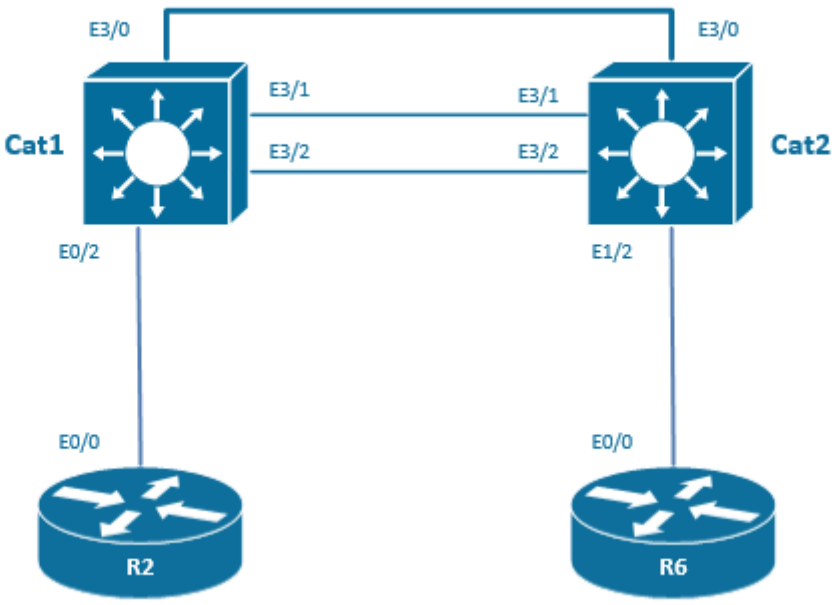

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 1.1: Disable CDP on R2.

Task 1.2: Disable CDP on the connection between R6 and Cat2.

Task 1.3: Between Cat1 and Cat2, CDP should only be running on the E3/1 and E3/2 interfaces. The updates should be sent every 20 seconds, and the neighbor should be declared lost after 5 missing updates.

Task 1.4: Between Cat1 and Cat2, the broadcast CDP packets should not report mismatched native VLAN IDs.

Task 1.5: Configure VLAN 101, 102, and 103 in the VLAN local database of Cat1 and Cat2 with the respective name of VLAN101, VLAN102, and VLAN103. The configuration of the VLANs should appear in the running-configuration and no VLAN distribution protocol should be running.

Task 1.6: Configure interface E3/0 in access mode VLAN 101 on Cat1 and Cat2.

Task 1.7: Configure the following IP addresses under the following interfaces:

Cat1 E0/2	10.1.0.1/24
R2 E0/0	10.1.0.2/24

Make sure that the ping is working.

Task 1.8: Configure an ISL trunk allowing VLAN 102 on E3/1. Leave it to DTP to negotiate, or not, a trunk.

Task 1.9: Configure a dot1q trunk allowing VLAN 103 on E3/2. Disable DTP on this connection. VLAN 103 should be sent untagged.

Task 1.10: Configure only the following SVIs:

Cat1 Vlan 103	10.103.0.1/24
Cat2 Vlan 101	10.101.0.2/24

Task 1.11: Configure the following sub-interfaces on E0/0 of R6:

E0/0.101	10.101.0.6/24
E0/0.103	10.103.0.6/24

Task 1.12: Ensure that you can ping from interface VLAN 103 on Cat1 to the interface VLAN 101 on Cat2 by using R6 as the inter-VLAN routing point. Do not use the "ip route" command.

You have completed Lab 1

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.

Lab 2: Configure and troubleshoot VTP

Technologies covered

- VTPv1
- VTPv2
- VTPv3
- VTP pruning

Overview

You have been tasked to automatically distribute the VLANs in the network using VTP. You have to propagate normal VLANs as well as extended VLANs. Your VTP set-up should be secured and high available.

The topology used in the lab will be the following:

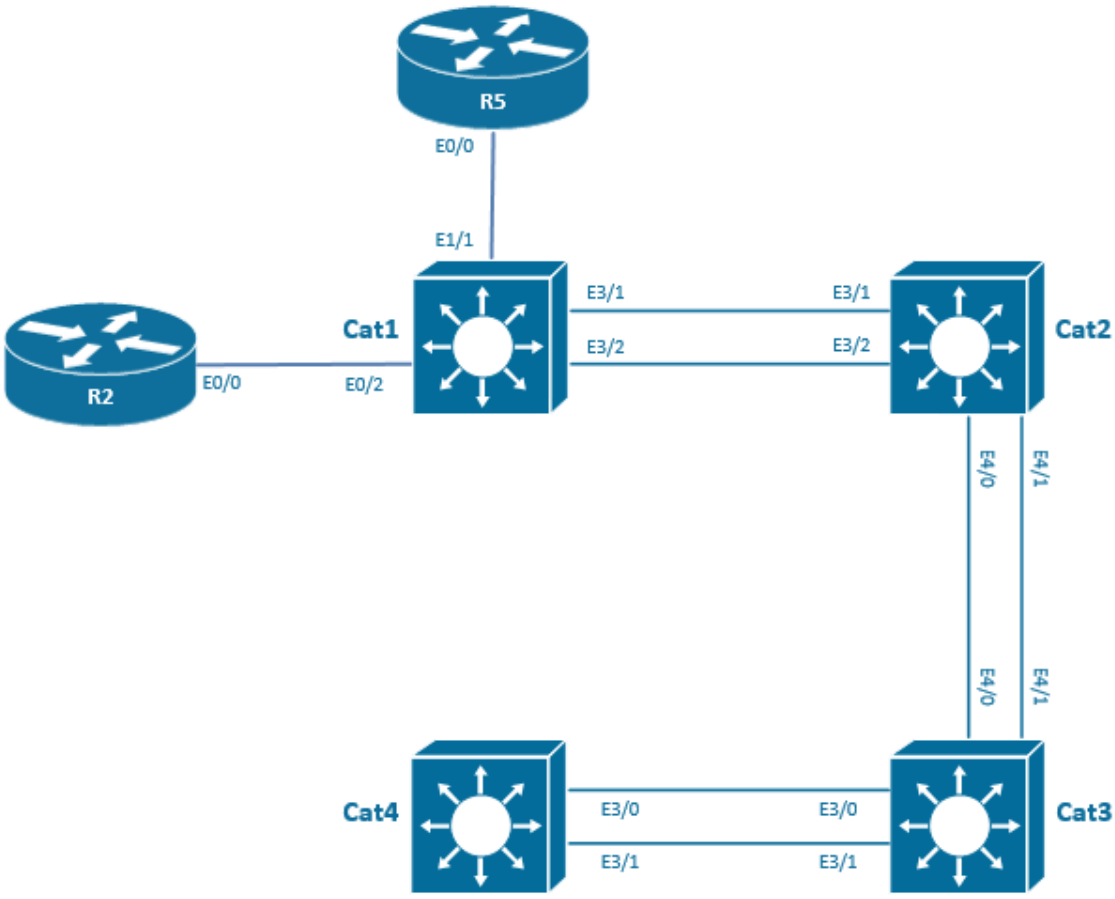

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 2.1: Configure a dot1q trunk allowing all VLANS on all the connections between Cat1 and Cat2, between Cat2 and Cat3, and between Cat3 and Cat4.

Task 2.2: Configure Cat4 as the server of the VTP domain iPexpert.

Task 2.3: Configure Cat3 not to update its VLAN database. VTP packets should be silently forwarded by Cat3.

Task 2.4: Configure Cat1 and Cat2 as client of Cat4.

Task 2.5: Add VLAN 150 and 151 on Cat4, and check that those VLANs are now present on Cat1 and Cat2, but not on Cat3.

Task 2.6: Add VLAN 1500 on Cat4, and make sure that it is propagated to Cat1 and Cat2, but not to Cat3.

Task 2.7: Configure the VTP domain with a password of "090909". This password should be stored in the NVRAM database.

Task 2.8: Ensure that the next VLAN created will not be propagated to switches where this VLAN is not allowed on any trunks.

Task 2.9: Ensure that Cat2 will take over the server role in the case of a failure of Cat4.

Task 2.10: Configure R2 in VLAN 150 and R5 in VLAN 1500 as client ports. As Cat1 is not having any client's port in VLAN 151, make sure that broadcast packets in VLAN 151 will never be transmitted to Cat1.

You have completed Lab 2

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.

Lab 3: Configure and troubleshoot Portchannels

Technologies covered

- LACP etherchannel
- PagP etherchannel
- Manual etherchannel
- L2 etherchannel
- L3 etherchannel
- Load-balancing
- Etherchannel misconfiguration guard

Overview

You have been tasked to configure seamless redundancy in the network by bundling several physical connections into a logical connection called port-channel. In addition, you should traffic-engineer the way that traffic is distributed on the different members of those port-channels.

The topology used in the lab will be the following:

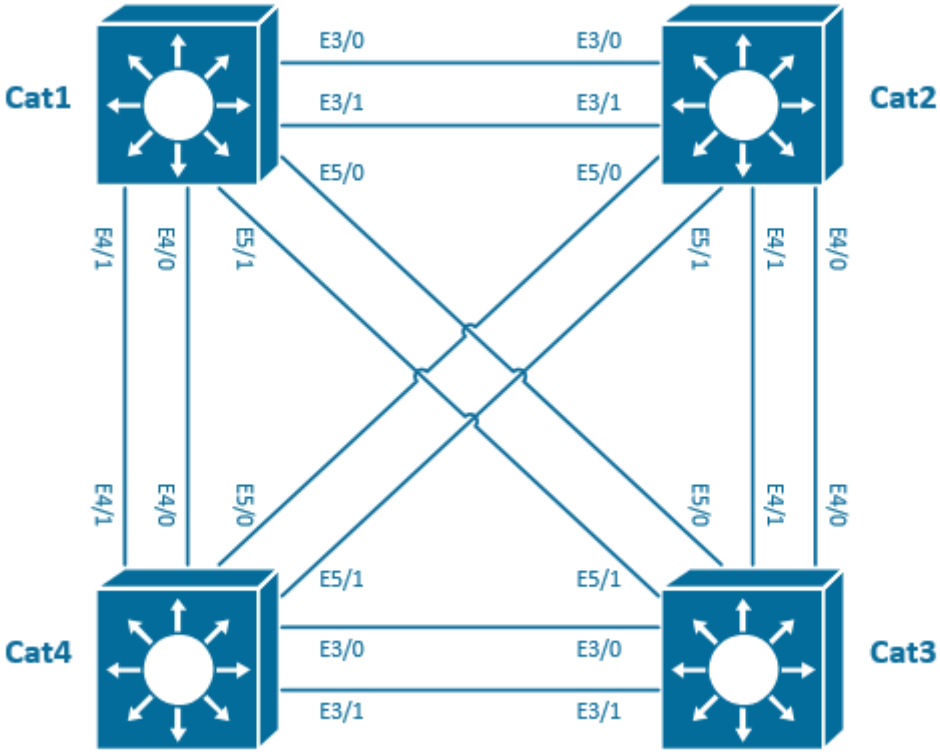

Estimated time to complete: 2-3 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the topology drawing. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 3.1: Between Cat2 and Cat3, configure a static port-channel Po23 trunking in a dot1q encapsulation the VLAN 101.

Task 3.2: Between Cat3 and Cat4, configure a PagP port-channel Po34 trunking in an ISL packet the VLAN 101. The Cat3 should not start the negotiation. Configure PagP in a way that the port-channel is protected against unidirectional failure.

Task 3.3: Between Cat2 and Cat4, configure a LACP port-channel Po24 trunking in the port in VLAN 102. Cat2 should never start the negotiation.

Task 3.4: Ensure that Cat4 is leading the LACP negotiation.

Task 3.5: Ensure that E5/0 will be used as LACP failover if 9 members are present in the Port-channel.

Task 3.6: Between Cat1 and Cat2, configure a static port-channel Po12 with the following IP address:

Cat1 Po12	10.12.0.1/24
Cat2 Po12	10.12.0.2/24

Task 3.7: Between Cat1 and Cat3, configure a PagP port-channel Po13 with the following IP address:

Cat1 Po13	10.13.0.1/24
Cat3 Po13	10.13.0.3/24

Task 3.8: Between Cat1 and Cat4, configure a LACP port-channel Po14 with the following IP address:

Cat1 Po14	10.14.0.1/24
Cat4 Po14	10.14.0.4/24

Task 3.9: On the Port-channel between the Cat1 and the Cat2, all the TCP flows from a source MAC address to the same destination MAC address should be using the same member in all the port-channels just configured.

Task 3.10: On the Port-channel between the Cat3 and the Cat4, make sure that all the flows coming from a MAC address are using the same PagP member when the packet returns to this MAC address.

Task 3.11: Configure the four switches with a mechanism to disable the port-channel in the case of a mis-configuration that is leading to the port-channel receiving Spanning-Tree BPDUs on two different members.

You have completed Lab 3

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.

Lab 4: Configure and troubleshoot Spanning-tree Protocol

Technologies covered

- PVST+
- Switch priority
- Port priority
- Path cost
- STP timers
- Port fast
- BPDUguard, BPDUfilter
- Loopguard
- Rootguard
- Backbonefast
- Loopfast
- UDLD

Overview

You have been tasked to guarantee in a redundant L2 network a loop-free topology by configuring the Spanning Tree protocol. Traffic engineering and optimization is also required. The 2 routers R6 and R9 will be considered as hosts that should not make part of the spanning-tree topology.

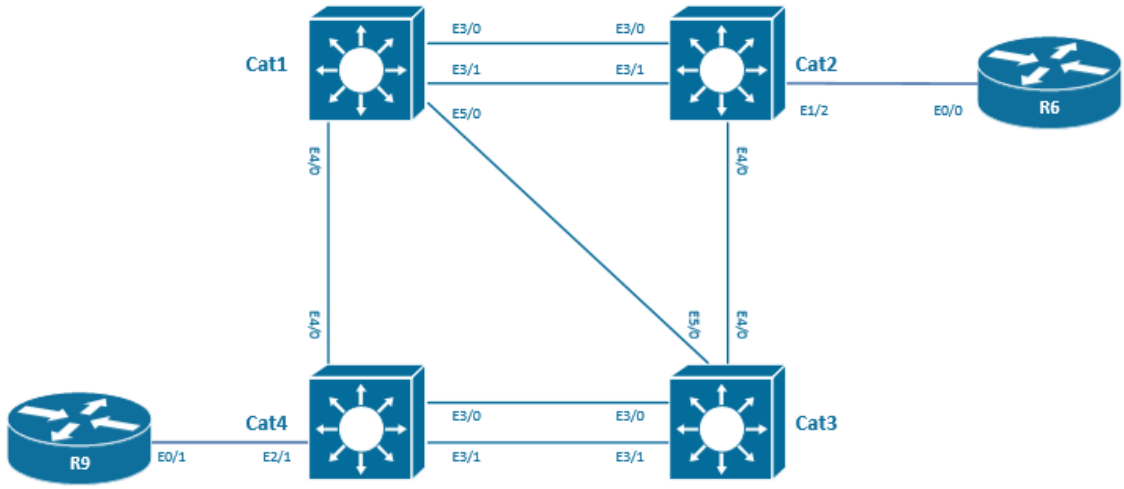

Estimated time to complete: 3-4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 4.1: Configure the 4 Catalysts to run PVST+ (and not rapid PVST+).

Task 4.2: Configure all the inter-switches connection as trunk dot1q trunking all the VLANs.

Task 4.3: Configure Cat1 as VTP server for the domain iPexpert and configure VLAN 21 and 22.

Task 4.4: Configure the primary root bridge on Cat2 for VLAN 21. Configure the secondary root bridge on Cat4 for VLAN 21. Do not use a command containing "priority" in order to achieve this. Optimize the timers to the number of switches.

Task 4.5: Configure the primary root bridge on Cat3 for VLAN 22. Configure the secondary root bridge on Cat1 for VLAN 22. Do not use a command containing "root" in order to achieve this.

Task 4.6: In VLAN 22, make sure that Cat2 and Cat4 will never become root of the network.

Task 4.7: On VLAN 22, change the hello timer to 5s, the max aging time to 20s and the forward delay to 15s.

Task 4.8: All the connections being up and running, on VLAN 21, the traffic from R6 to R9 should be forwarded using the following path: Cat2-Cat1-Cat3-Cat4.

Task 4.9: All the connections being up and running, on VLAN 22, the traffic from R6 to R9 should be forwarded using the following path: Cat2-Cat3-Cat4.

Task 4.10: All connections being up and running, on VLAN 21, the traffic from Cat1 to Cat2, and from Cat3, and Cat4 should flow over the E3/0 connections.

Task 4.11: All the connections being up and running, on VLAN 22, the traffic from Cat3 and Cat4 should flow over the E3/0 connection.

~~**Task 4.12:** With all connections up on VLAN 22, the traffic from Cat3 and Cat4 should flow over the E3/0 connection.~~

Task 4.12: Reduce the convergence time associated with indirect failures in the network.

Task 4.13: Enable the Uplinkfast feature on the switches where it cannot create loops. When a failure occurs on a switch with Uplinkfast feature on, a maximum of 100 dummy multicast packets have to generated every second in order to update the rest of the network bridging tables.

Task 4.14: Configure R6 as a client in VLAN 21 in access mode.

Task 4.15: Configure R9 as a client with a trunk connection allowing VLAN 22. VLAN 22 should be native of the dot1q trunk.

Task 4.16: Allow the port connected to the routers to transition immediately from blocked to forwarding.

Task 4.17: R6 could be sending BPDUs and we would like the port to be put in error-disabled in the case that it happens. Configure the port to re-enable itself automatically after 1 minute.

Task 4.18: VLAN R9 is sending BPDUs, but we would like to ignore them and to silently drop them.

Task 4.19: The link between Cat1 and Cat3 should be protected from a loop caused by a unidirectional link. Do not use UDLD.

Task 4.20: The link between Cat1 and Cat4 should be removed from the network topology if an unidirectional link is detected. The port on Cat1 should be put in err-disabled when an unidirectional event happens but not the port on Cat4. Configure the port to re-enable itself automatically after 5 minutes.

You have completed Lab 4

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.

Lab 5: Configure and troubleshoot Multi-Instance Spanning-tree Protocol (MST)

Technologies covered

- MST
- MST region
- CST
- RPVST+

Overview

The switches will run very CPU intensive processes. You have been tasked to optimize the spanning-tree protocol in order to create fewer burdens on the CPU of the switches. Running one SPT process for a group of VLANs is made possible with MST.

The topology used in the lab will be the following:

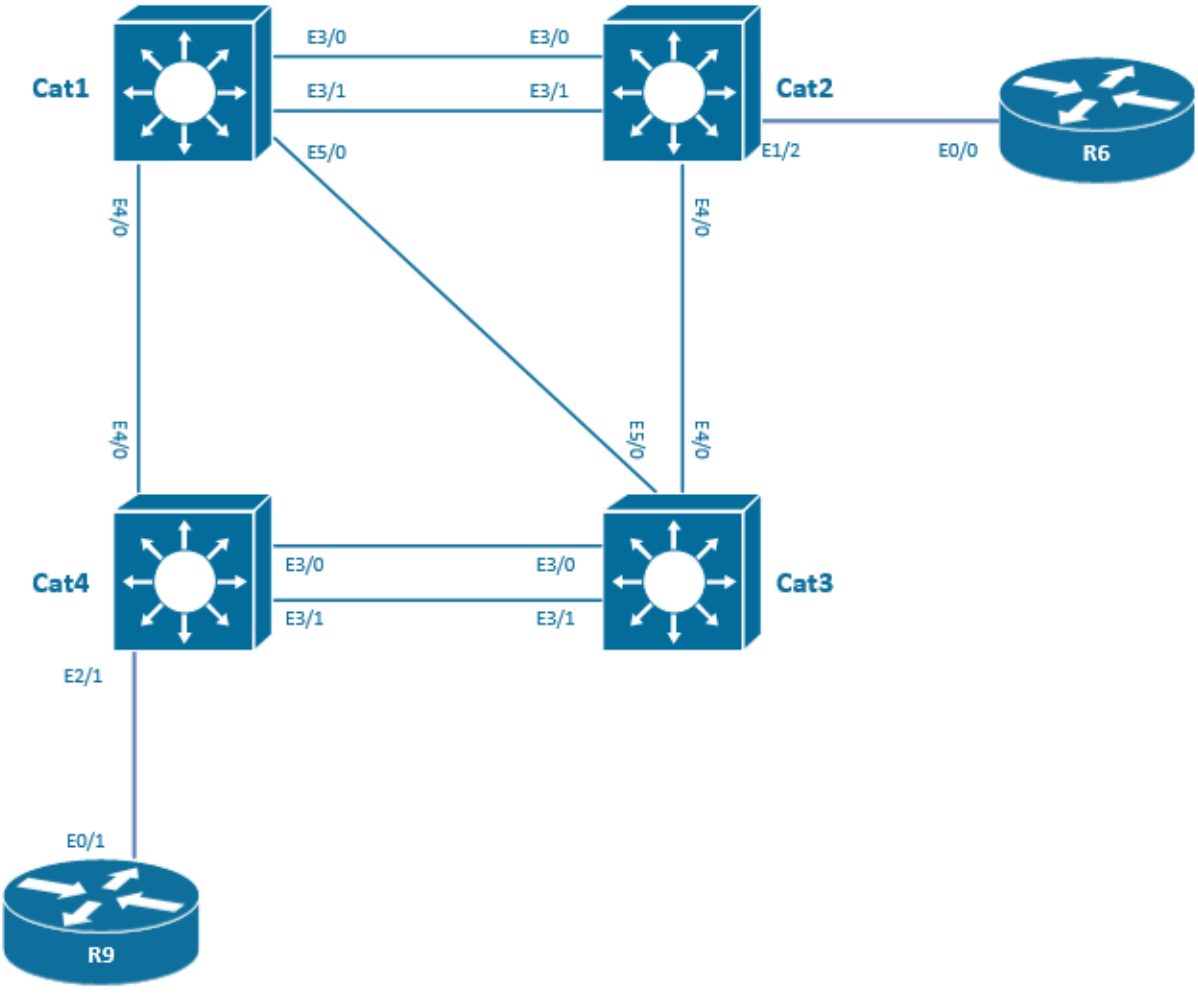

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 5.1	Configure the Cat1, Cat2, and Cat3 Switches to run the MST protocol with the name iPexpertRegion. Configure VLAN 100, 110, and VLAN 200, 210 on the Cat1, Cat2, and Cat3 Switches.
Task 5.2	Instance 10 with the name iPexpert10 will encompass the VLAN range 100-150.
Task 5.3	Instance 20 with the name iPexpert20 will encompass the VLANs 200, 210, 220, 230, 240, and 250.
Task 5.4	Configure all the inter-switches connection as trunk dot1q trunking all the VLANs.
Task 5.5	For instance 10, configure Cat2 to be the root primary and Cat3 to be the root secondary. Do not use the priority command.
Task 5.6	For instance 20, configure Cat3 to always be the root primary and Cat2 to be the root secondary.
Task 5.7	Between Cat1 and Cat2, make sure that the blocked path is on the E3/0 for instance 10.
Task 5.8	Configure VLAN 100, 110, 200, and 210 on Cat4.
Task 5.9	Configure the MST region iPexpertRegion to always be the root of the CST.
Task 5.10	Ensure that the port E4/0 on Cat4 is in BLK state.
Task 5.11	Ensure that the port E3/0 on the Cat4 is in BLK state.
Task 5.12	Make sure that the spanning-tree reconfiguration on Cat4 occurs in less than one second with 802.1w.

You have completed Lab 5

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.

Lab 6: Miscellaneous Layer 2 Topics

Technologies covered

- Managing MAC address table
- Protected ports
- Stormcontrol
- SPAN
- RSPAN
- ERSPAN
- Voice VLANs
- Smartports Macros
- Private VLAN

Overview

There are some application problems in the network. You have been tasked to troubleshoot and understand the performance issues by sniffing the problematic traffic and setting up a SPAN and RSPAN session. As Cisco IP phones will be hooked up to the network, you will be asking to configure those ports and guarantee voice quality.

The topology used in the lab will be the following:

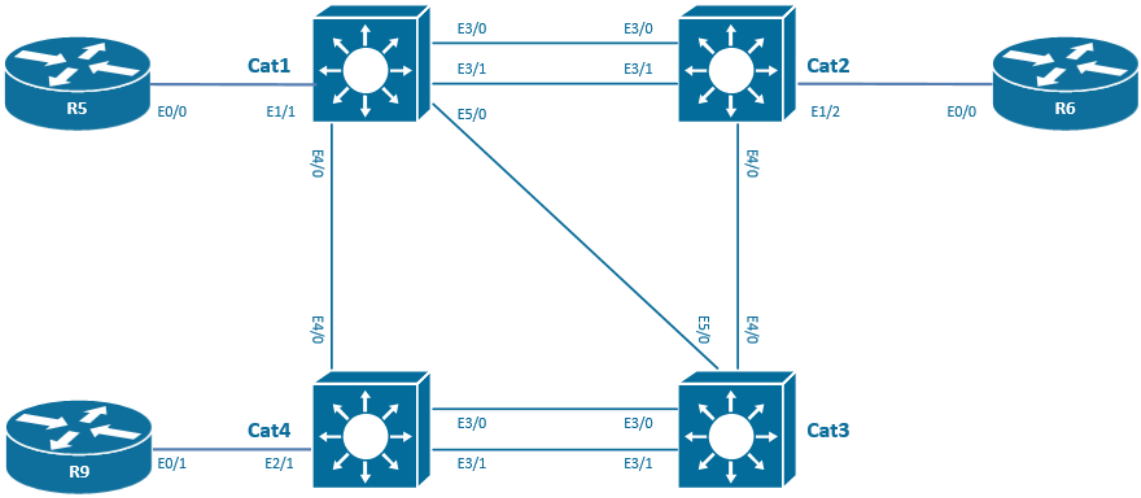

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 6.1: On Cat1, the dynamic MAC-address table entries should be removed from the table when they are not re-learned after 10 seconds.

Task 6.2: On Cat1, for troubleshooting reasons, enable the MAC address change notification feature. Configure the switch to send SNMP traps to server

10.1.99.99 with the community iPexpert1 as soon as a MAC address is removed or added on interface E1/1. Keep up to 500 entries in the MAC notification table.

Task 6.3: On Cat1 configure interface E1/1 as an access port in VLAN 120.

Task 6.4: On Cat1, disable MAC address learning in VLAN 120 and add a static entry that indicates the MAC address of the interface E0/0 of R5 is located in VLAN 120 behind interface E1/1.

Task 6.5: On Cat1, enable unicast MAC address filtering in VLAN 120 and configure the switch to drop packets that have a source or destination address of cafe.2222.2222.

Task 6.6: On Cat1, ensure that a server connected to the interface E5/1 and a server connected to the interface E1/2 cannot send traffic to each other at layer 2. Do not use port-security.

Task 6.7: On Cat4, prevent traffic on the LAN from being disrupted from a broadcast and unicast storm on the interface E2/1. A storm is considered a storm when more than 50% of the bandwidth is used by broadcast packets and when more than 80% of the bandwidth is used by unicast packets.

Task 6.8: Multicast packets should always be dropped on the interface E2/1. Use storm-control.

Task 6.9: Configure a dot1q trunk between Cat2 and R6. This trunk should be allowed on VLAN 121 and VLAN 122.

Task 6.10: A laptop called Laptop1 with a Wireshark sniffer is connected on Cat2 on the port E1/3. Configure this port with dot1q trunk encapsulation allowing all the VLANs.

Task 6.11: Configure the Cat2 switch to mirror all the traffic transiting in VLAN 121 on Cat2 on E1/2 to the port where the sniffer Laptop1 is connected. Use session number 60.

Task 6.12: The port where the Sniffer is connected should accept incoming traffic with a dot1q encapsulation. Default ingress VLAN is VLAN 121.

Task 6.13: Configure a LACP port-channel between Cat1 and Cat2. Bundle interface E3/0 with E3/1 on both sides. This port-channel is a dot1q trunk allowing VLAN 121, VLAN 122, and VLAN 500.

Task 6.14: A laptop called Laptop2 with a Wireshark sniffer is connected on Cat1 on the port E0/3. Configure this port with an access port in VLAN 1.

Task 6.15: Configure the mirroring of the sent traffic transiting in VLAN 122 on Cat2 on E1/2 to the port where the sniffer Laptop2 is connected. Use session number 61 and VLAN 500 as RSPAN VLAN.

Task 6.16: On Cat3, there will be a Cisco IP phone connected to the port E1/0. Enable QOS on the Cat3 and configure the port E1/0 to trust COS.

Task 6.17: Configure a VLAN of 33 reserved for voice traffic on Cat3. The voice traffic on E1/0 should use this voice VLAN.

Task 6.18: The incoming Data frames coming from a computer connected on the Cisco IP phone should be tagged by the switch with a COS of 2.

Task 6.19: On Cat3, configure a macro called “Bounce-int” to bounce (shut followed by a no shut) an interface. Use a variable called $int. Test and run the macro for E1/0.

Task 6.20: On Cat3, there will be an additional Cisco IP phone connected to the E1/1. Use the preconfigured macro called “cisco-phone” to configure the port. Voice VLAN has to be VLAN 2 and Data VLAN has to be VLAN 1.

Task 6.21: On Cat1 and on Cat4, configure VLAN 120 as the primary VLAN, VLAN 130 as the isolated VLAN, and VLAN 140 as the community VLAN. Configure E4/1 Cat4 as the PVLAN promiscuous port. Configure interface E4/0 and int E5/0 Cat1 as the PVLAN host port for VLAN 130, interface E5/1, and interface E3/0 Cat1 as the PVLAN host port for VLAN 140. The connection between Cat1 and Cat4 has to be configured as a trunk port that will support the setup.

You have completed Lab 6

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.

Lab 7: HDLC and PPP/PPPoE

Technologies covered

- HDLC
- PPP PAP, CHAP
- PPPoE
- MLPPP
- PPP inter-leaving
- RTP reserve
- Virtual-assembly

Overview

You have been tasked to configure the serial connections of your network with HDLC and PPP encapsulation. PPP connection may have to be authenticated or aggregated in a bundle.

The topology used in the lab will be the following:

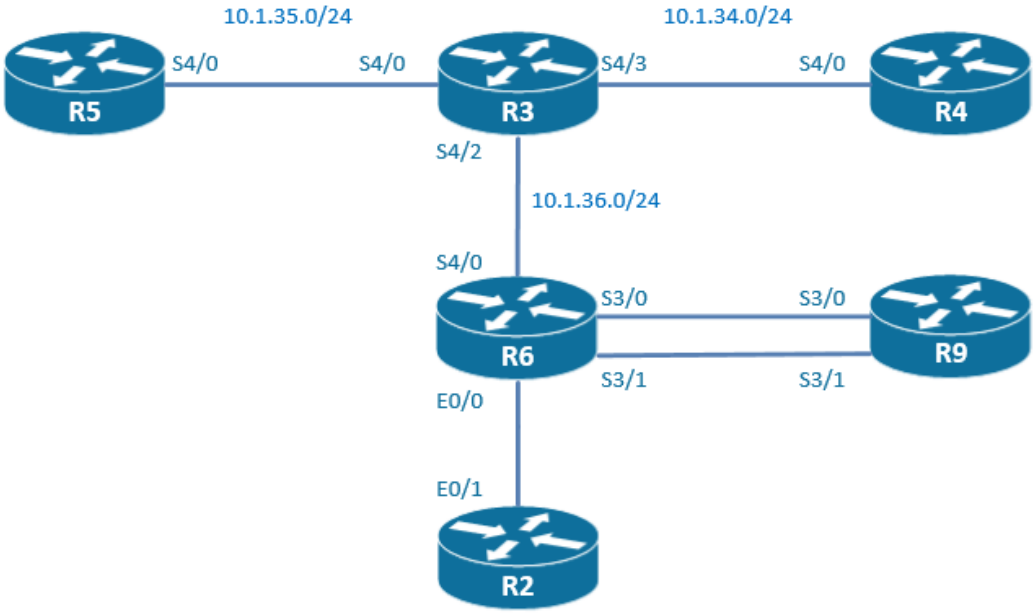

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 7.1: The link between R3 and R4 should be using the HDLC encapsulation. Check that you can ping from R3 to R4.

Task 7.2: The link between R3 and R5 should be using the PPP encapsulation. Turn on the CHAP authentication with the password of "Password35". Check that you can ping from R3 to R5.

Task 7.3: The link between R3 and R6 should be using the PPP encapsulation. Turn on the PAP authentication with the password of "Password361". If the PAP authentication is unsuccessful, CHAP authentication has to kick in with a password of "Password362". Check that you can ping from R3 to R6.

Task 7.4: Configure PPPoE between the R6 and the R2 routers. R6 is the server side and R2 is the client side. On the server side, a BBA is called "iPexpertgroup". The IP pool is called "iPexpertpool" and the range is from 10.1.26.10 to 10.1.26.20. The virtual-template number should use id 23 and the IP address configured on the virtual template is 10.1.26.6 255.255.255.0.

Task 7.5: Limit the number of sessions established (per client MAC address) to 3.

Task 7.6: On the client side, use the id 26 for both the dialer interface and the dialer-pool-number interface. Check that you can ping from R6 to R2.

Task 7.7: Make sure that unnecessary fragmentation is avoided.

Task 7.8: The client R2 should authenticate when connecting on the server. Create a local account username called R2 with the password "Password26".

Task 7.9: Bundle with PPP multilink the two serial connections between R6 and R9. Use a group ID of 69.

Task 7.10: Configure the IP address 10.1.69.6/24 on the R6 PPP multilink69. Configure the IP address of 10.1.69.9/24 on the R9 PPP multilink69. Check that you can ping from R6 to R9.

Task 7.11: Ensure that it is checked on the PPP multilink interfaces that all the fragments of an IP datagram are received on the virtual interfaces before forwarding them.

Task 7.12: There will be voice traffic running over the multilink PPP connection. Ensure that a small voice packet is delayed a maximum of 20 ms because of the transmission of a big data packet.

Task 7.13: Reserve 1 Mbps in a special queue for real-time packet flows designated to the UDP port starting 32768 and ending 32867.

You have completed Lab 7

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.

Lab 8: Configure and troubleshoot Basic IP routing

Technologies covered

- Static route
- Traffic engineering
- Floating static route
- Object tracking
- PBR
- GRE

Overview

You have been tasked to configure the routing in your network.

The topology used in the lab will be the following:

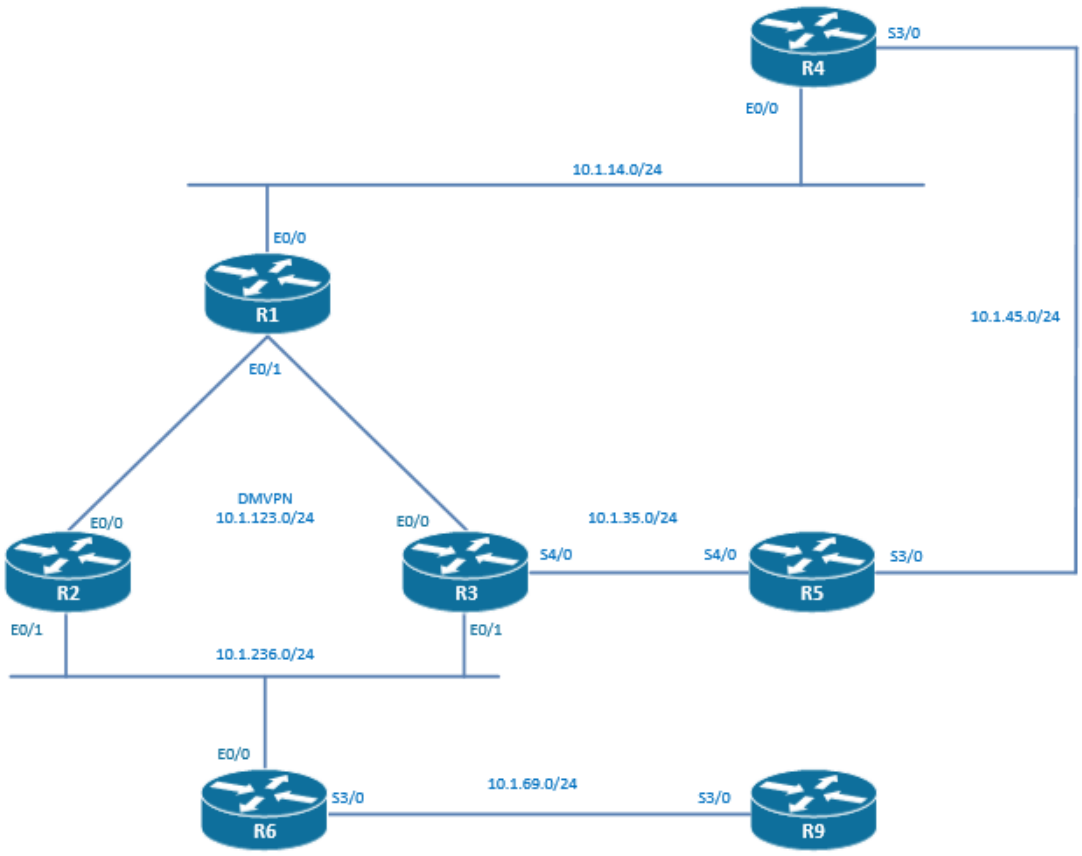

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 8.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. Configure DMVPN phase 2 as the underlying technology. Multicast support has to be configured.

Task 8.2: On R1, configure a static route to the loopback0 of R3 using the tunnel interface on R3 as the next-hop. Check that you can ping the loopback0 of R3 with a ping sourcing on the tunnel interface of R1.

Task 8.3: On R2 tunnel interface, disable proxy-arp.

Task 8.4: On R1, configure a static route to the loopback0 of R2 using the tunnel interface on R2 as the egress interface.

Task 8.5: On R1, ensure that you can ping the loopback0 of R2 with a ping sourcing on the tunnel interface of R1. Create a static arp entry to achieve this task.

Task 8.6: On R6, configure a static route to network 10.1.0.0/16 pointing to E0/0. Check that you can ping the loopback0 of R2 and R3.

Task 8.7: Disable proxy-arp on E0/1 of R2 and R3. Ensure that you can ping the loopback0 of R2 and R3 with a ping sourcing from the E0/0 ip address of R6.

Task 8.8: Configure a GRE tunnel interface Tunnel0 between the loopback0 of R6 and the loopback0 of R3. Use ip address 36.0.0.3/24 on R3 and 36.0.0.6/24 on R6. Configure default routes on R6 and R3 with an AD of 250.

Task 8.9: On R6, configure a static route to the loopback network of the router R3 using the Tu0 as egress with an AD of 5. The tunnel0 interface should go down because of a recursion issue. Leave this tunnel0 down as it is.

Task 8.10: Configure static routing so that you can ping the loopback0 of R1 with a ping sourcing from the loopback0 ip address of R6. The ping should follow the R6-R3-R1 route and use the DMVPN tunnel.

Task 8.11: Configure a GRE tunnel interface Tunnel16 between the loopback0 of R6 and the loopback0 of R1. Use ip address 16.0.0.1/24 on R1 and 16.0.0.6/24 on R6.

Task 8.12: On R3, configure a floating static route that will be used in the case that the tunnel interface to R1 goes down. This floating route should not point to R1, but to R5 as a next-hop. At this point, you are not asked to configure all the static routing that will make the backup path operational.

Task 8.13: On R4, configure a default-route using the next-hop of R1. On R1, configure a static route to the network 10.1.4.0/24 pointing to the next-hop on R4.

Task 8.14: On R4, configure a default-route using the next-hop of R5 with an AD of 5.

Task 8.15: The default-route using the next-hop of R5 should be used when the loopback0 of R1 has become unreachable. Use object tracking and IP SLA.

Task 8.16: On R5, configure default routing using policy-based routing. This default routing should be pointing to a next-hop of R3 IP address using PBR. When CDP detects that R5 to R3 connectivity is down, the traffic should be routed over R4. Do not use local policy-base routing.

Task 8.17: On R9, use local-policy based routing to route to the loopback interface of R6.

Task 8.18: On R6, use local-policy based routing to route to the loopback interface of R9. You should be able to ping the loopback0 of R6 with a ping sourcing from the loopback0 of R9.

You have completed Lab 8

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.

Lab 9: Configure and troubleshoot Routing Information Protocol (Part 1)

Technologies covered

- RIP version 2
- Split-horizon
- Auto-summarization
- Send and receive version
- Manual summarization
- Convergence timers
- Offset-list
- Distribute-list
- Per neighbor AD filtering

Overview

You have been tasked to configure routing in your network using the RIP version 2 protocol.

The topology used in the lab will be the following:

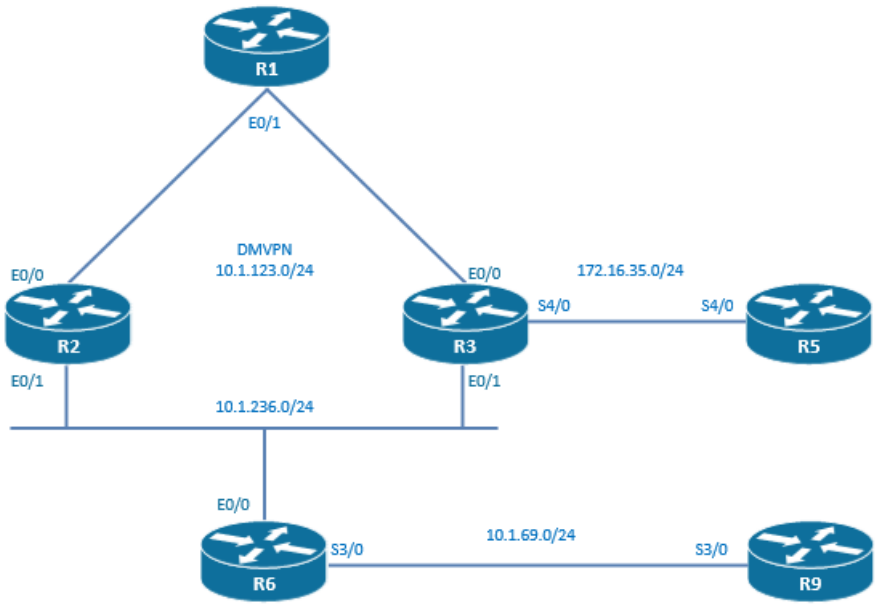

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 9.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Configure RIP version 2 in this DMVPN network.

Task 9.2: Advertise the loopbacks 10 of R1, R2, and R3 in the RIP process.

Task 9.3: Ensure full reachability in this hub and spoke technology. On R2, check that you can ping the loopback10 of R3 sourcing from the loopback10 of R2.

Task 9.4: Configure RIP version 2 between R5 and R3. Advertise the loopbacks of R5 in the RIP process.

Task 9.5: Ensure that there is a single 10.0.0.0/8 entry in the routing table of R5. Use manual summarization.

Task 9.6: Ensure that the network 200.0.0.0/24 is advertised to the router R3. Do not use manual summarization.

Task 9.7: Enable RIP on the 172.16.236.0/24 network.

Task 9.8: Advertise the loopbacks 0 and 1 of R6 in the RIP process. R6 is running version 1.

Task 9.9: Make sure that the interfaces part of network 172.16.236.0/24 can send and receive either version 1 or version 2 packets.

Task 9.10: Configure RIP MD5 authentication on the 11.1.1.0/24 network. Use a key chain of “iPexpertchain”, a key number 1, and a key-string of “iPpassword”.

Task 9.11: On R2, the network 200.0.0.0/8 received on Ethernet0/0 should be rejected, and the network 201.0.0.0/8 received on Ethernet0/1 should be rejected. Do not use distribute-list or administrative distance poisoning.

Task 9.12: On R1, all the traffic should be sent to R2 and R3 should never be considered as a next hop. Do not use offset-list or administrative distance poisoning. Configure 2 Prefix-lists.

Task 9.13: On R1, the network 23.0.0.0/8 should be routed via the tu23 and the network 24.0.0.0/8 should be routed via the E0/1. Use administrative distance poisoning.

Task 9.14: Configure RIP filtering so that R3 does not learn 5.0.0.0/24. Do not use any access-list, distribute-list, and do not change AD values. R5 should learn all RIP subnets.

Task 9.15: Configure the RIP timers on R1, R2, and R3 to 20 second updates, 40 second invalid, 10 second hold, and 80 second flush.

Task 9.16: On R3, configure Serial4/0 to send updates every 6 seconds towards R5.

You have completed Lab 9

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.

Lab 10: Configure and troubleshoot Routing Information Protocol (Part 2)

Technologies covered

- RIP default route
- RIP update
- Unicast update
- Broadcast update
- Triggered update
- Source validation

Overview

You have been tasked to configure routing in your network using the RIP version 2 protocol.

The topology used in the lab will be the following:

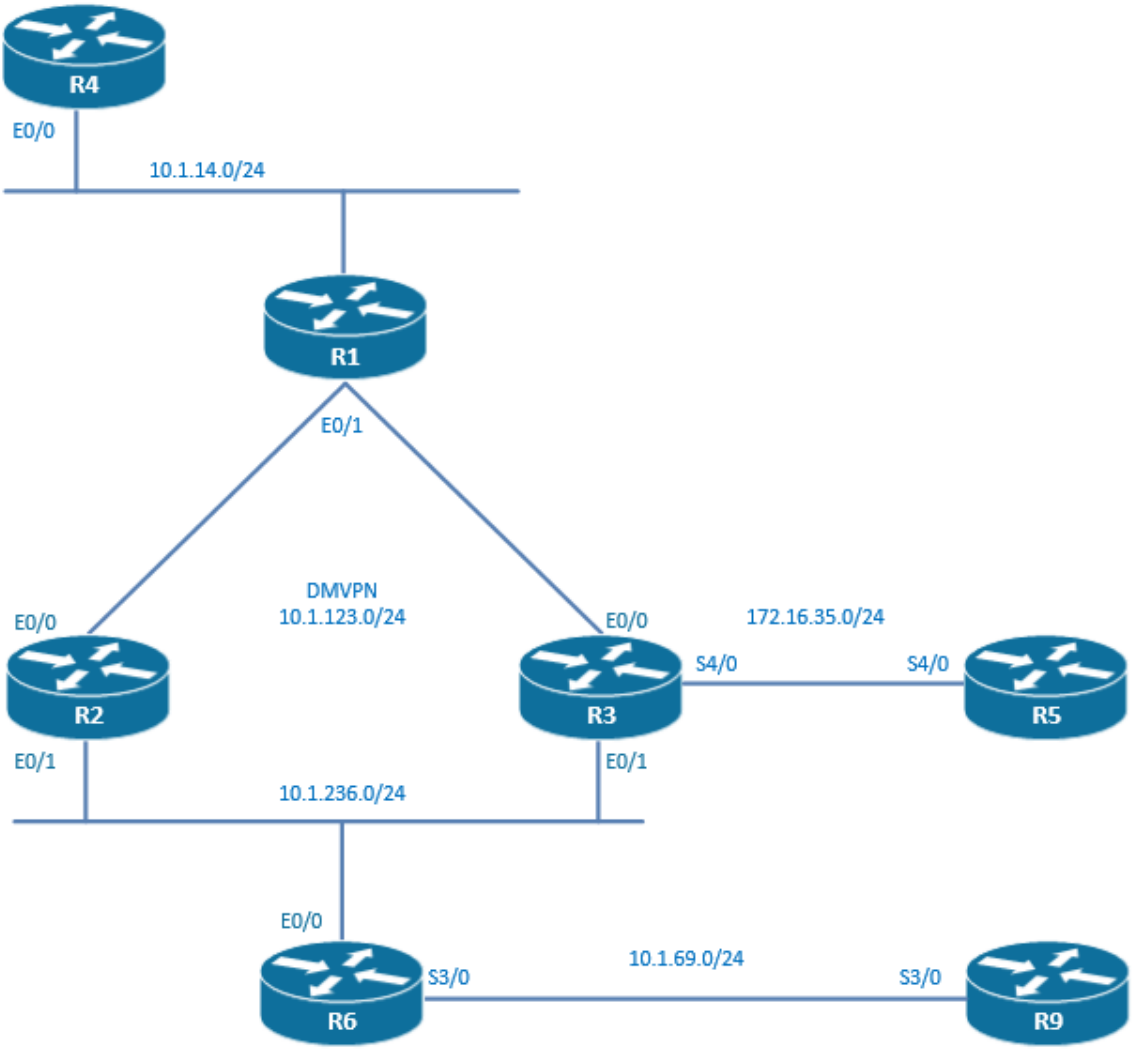

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 10.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Configure RIP version 2 in this DMVPN network.

Task 10.2: The RIP updates have to be sent as unicast packets on the DMVPN tunnels.

Task 10.3: Advertise the loopbacks 0 of R1, R2, and R3 in the RIP process.

Task 10.4: Ensure full reachability in this hub and spoke technology. On R2, check that you can ping the loopback of R3 sourcing from the loopback of R2.

Task 10.5: Configure RIP version 2 between R1 and R4. Advertise the loopback of R4 into the RIP process.

Task 10.6: R1 should advertise a default route to all its RIP neighbors with the exception of R4.

Task 10.7: If the E0/0 interface is going down, R1 will stop advertising this default route.

Task 10.8: Configure RIP version 2 on the LAN connecting R2, R3, and R6. Advertise the loopback of R6 into the RIP process.

Task 10.9: The RIP updates should be broadcasted on the LAN 10.1.236.0/24.

Task 10.10: Configure RIP version 2 on the serial connection between R3 and R5. Advertise the loopback 0 of R5 into the RIP process.

Task 10.11: The RIP updates between R3 and R5 should stay silent. Updates should be sent only when there is a change in the topology.

Task 10.12: Configure RIP version 2 on the serial connection between R6 and R9. Advertise the loopback of R9 into the RIP process.

Task 10.13: Configure PPP encapsulation on the serial connection between R6 and R9. Use IPCP for address allocation with PPP. R6 is the server side (IP address 10.1.69.6/24) and R9 is client side (IP address 10.1.69.9/32 assigned by server). Ensure that R6 is getting the RIP updates from R9 and that you can ping the loopback of R9 sourcing from the loopback of R6.

Task 10.14: R5 should advertise a default-route to R3. This default-route should only be advertised if the network 10.1.2.2/32 is present in the routing table.

You have completed Lab 10

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.

Lab 11: Configure and troubleshoot EIGRP (Part 1)

Technologies covered

- EIGRP AS mode
- EIGRP named mode
- Stub
- Summarization
- Authentication
- Key chain rotation
- Prefix number limiting

Overview

You have been tasked to configure the routing reachability in your network using the EIGRP protocol.

The topology used in the lab will be the following:

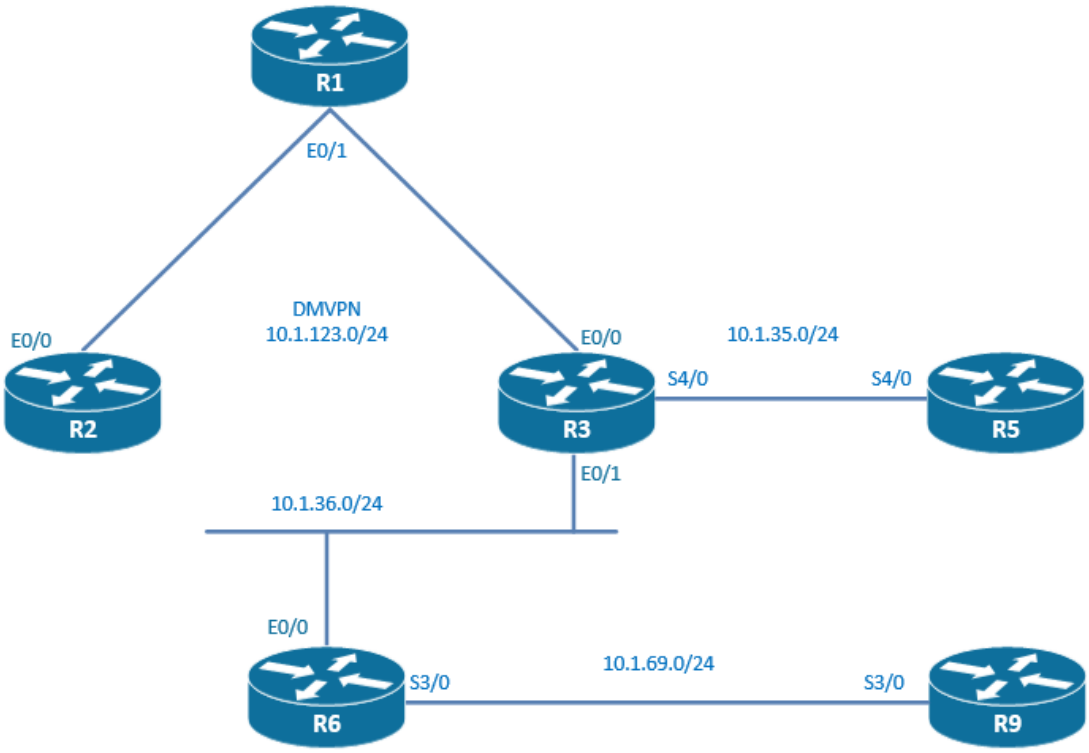

Estimated time to complete: 2-3 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 11.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Setup EIGRP routing in autonomous configuration mode with AS11 in this DMVPN network.

Task 11.2: Advertise the loopbacks of R2 and R3 in the EIGRP process. Only the 12.1.x.x/24 networks should be redistributed from connected into the routing protocol.

Task 11.3: Redistribute only the loopback0 on R1 in the EIGRP process.

Task 11.4: Make sure that there is full connectivity between loopbacks with the DMVPN network.

Task 11.5: Make sure that the traffic from the spoke to spoke is not transiting by the hub.

Task 11.6: R2 should advertise the 12.1.0.0/16 network out to R1 with a metric using the following parameters:

bandwidth	100 000 kilobits per s
delay	5 tens of microsecond
reliability	255
load	20
mtu	1500 bytes

Task 11.7: R2 is not transiting any traffic, so R2 should not receive EIGRP query packets anymore. Configuration for this task should be performed on R2, and loopbacks of R2 should stay reachable.

Task 11.8: On R6 and R9, setup EIGRP routing in named configuration mode using AS11 and the name of “iPexpert”. Advertise the loopbacks of R6 and R9 in the EIGRP process. On R9, ensure that you can ping the loopback1 of R2 from the loopback0 of R9.

Task 11.9: Configure the only possible EIGRP authentication mode between R6 and R3. Use a key chain called “keyiPexpert1” with 2 keys. Key 1 with a key-string of “Password1” is used since 03:00:00 Jan 1 2014 until 03:00:00 Jan 1 2015, but can already be used one month before and is still valid one month after. Key 2 with a key-string of “Password2” will be used from 03:00:00 Jan 1 2015 onwards, but can be used since 03:00:00 Dec 15 2014.

Task 11.10: Configure EIGRP HMAC-SHA-256 authentication between R6 and R9. Use a key-string of “Password3”.

Task 11.11: On R6, generate a syslog message when the maximum prefix limit of 10 has been accepted from the neighbor R9. Do not take any other action when this max limit of 10 is exceeded.

Task 11.12: On R6, tear down the EIGRP neighborship relations when more than 20 prefixes are received by the EIGRP process, and generate a syslog message when more than 10 prefixes have been accepted.

You have completed Lab 11

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.

Lab 12: Configure and troubleshoot EIGRP (Part 2)

Technologies covered

- Summarization with default routing
- Summarization with leak-map
- Summarization with floating default routing
- EIGRP metric weights
- TE
- Unequal cost load balancing
- EIGRP timers

Overview

You have been tasked to configure the routing reachability in your network using the EIGRP protocol.

The topology used in the lab will be the following:

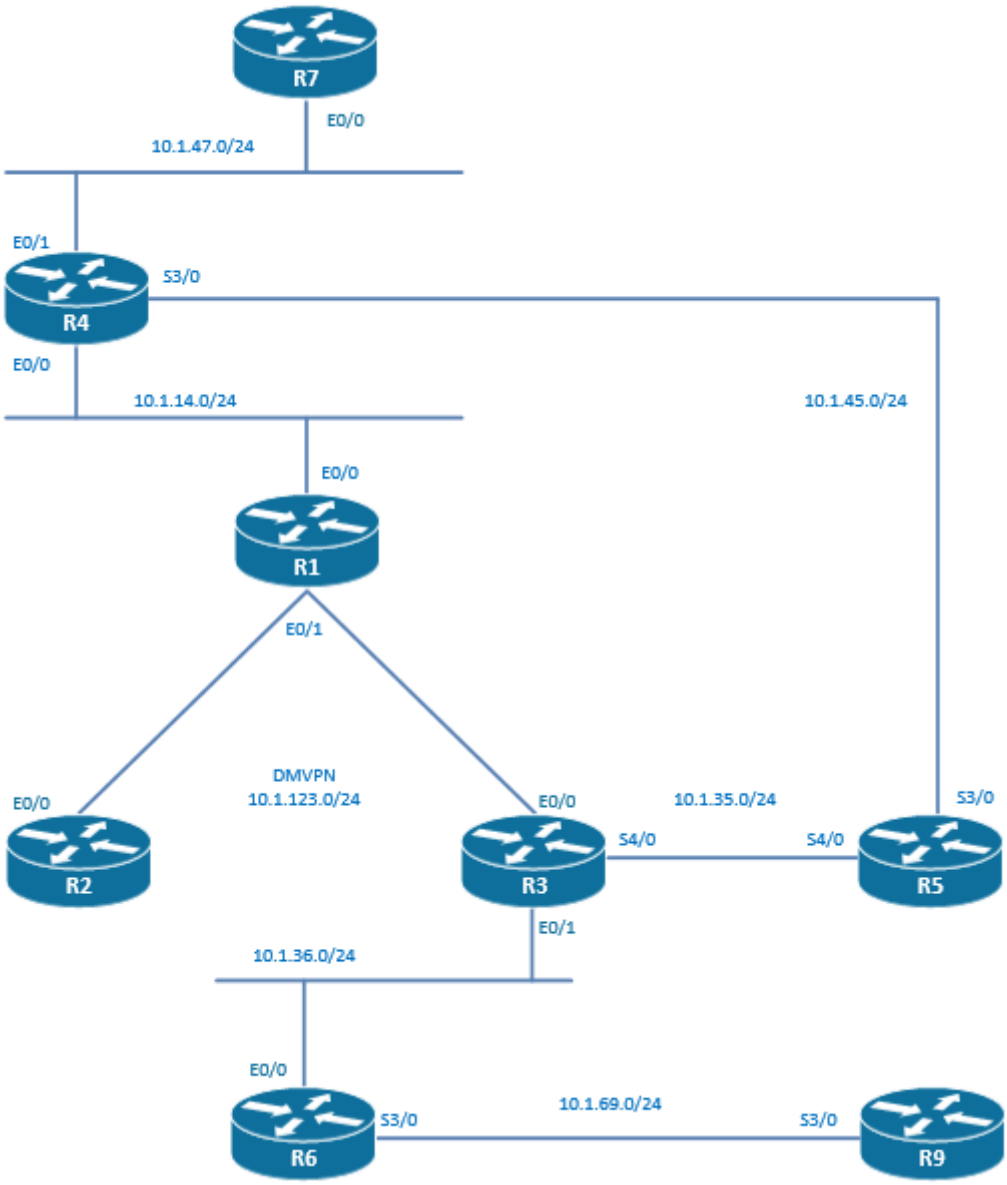

Estimated time to complete: 2-3 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 12.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Setup EIGRP routing in autonomous configuration mode with AS4 in this DMVPN network.

Task 12.2: Advertise loopback0 on R1, R2, and R3 in the EIGRP process using network statements.

Task 12.3: Setup EIGRP routing between R3 and R5. Advertise the loopback0 into the EIGRP process.

Task 12.4: On R3, configure summarization in a way that R5 only receives a default-route from R3. Leak also the loopback 10.1.4.4.

Task 12.5: Setup EIGRP routing between R3 and R6, and between the R6 and R9. Advertise the loopback0 of R6 and R9 into the EIGRP process. On R3, check that you can ping the loopback of R9 using the loopback of R3 as a source.

Task 12.6: On R3, configure summarization in a way that R6 only receives a default-route from R3.

Task 12.7: On R6, configure summarization in a way that R9 only receives a default-route from R6.

Task 12.8: On R3, check that you can ping the loopback of R9 using the loopback of R3 as a source. Use a floating route summarization.

Task 12.9: Setup EIGRP routing between R1 and R4, and between the R4 and R7. Advertise the loopback0 of R4 and R7 into the EIGRP process. On R1, check that you can ping the loopback of R7 using the loopback of R1 as a source.

Task 12.10: On R4, configure summarization in a way that R7 receives from R4 a default-route and the loopback0 networks of R1, R2, and R3.

Task 12.11: Setup EIGRP routing between R4 and R5.

Task 12.12: In the whole EIGRP domain, configure the metric calculation to use K1=0, K2=0, K3=1, K4=0, and K5=0.

Task 12.13: Configure a delay of 512 on the link between R4 and R5, a delay of 256 on the link between R4 and R1, a delay of 256 on the link between R1 and R3, and a delay of 128 on the link between R3 and R5.

Task 12.14: Configure bidirectional un-equal cost load-balancing between R4 and R5. Use off-set list when it is necessary.

Task 12.15: Configure R6 to send EIGRP hello packets every 1 s to R9.

Task 12.16: In the EIGRP domain, ensure that a router that has not replied to an EIGRP Query packets for 2 minutes is declared Stuck in Active.

Task 12.17: On R9, configure a NSF during 5 minutes when the R6 NSF-capable router is undertaking a switchover.

You have completed Lab 12

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.

Lab 13: Configure and troubleshoot EIGRP (Part 3)

Technologies covered

- Stub routing with leak-map
- Filtering with passive interfaces
- Filtering with distribute-list
- Filtering with offset-list
- Filtering with AD
- Filtering with route-maps
- Bandwidth pacing
- Neighbor logging
- Router-id
- Maximum hops

Overview

You have been tasked to configure the routing reachability in your network using the EIGRP protocol.

The topology used in the lab will be the following:

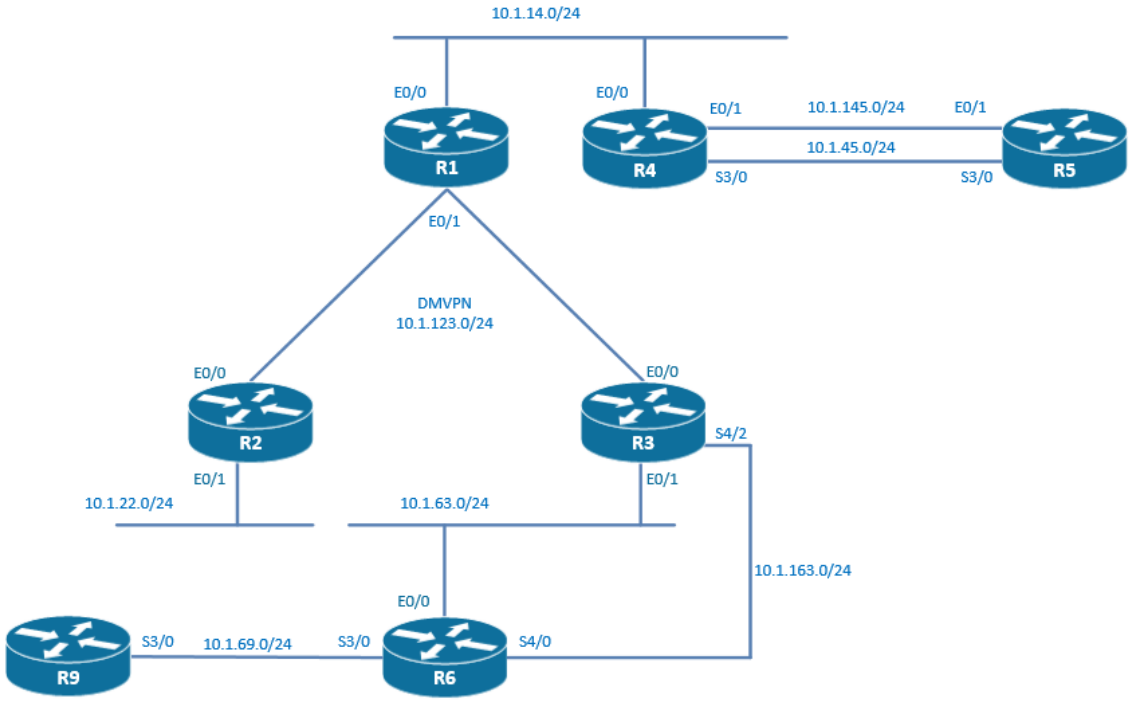

Estimated time to complete: 2-3 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 13.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Setup EIGRP routing in autonomous configuration mode with AS33 in this DMVPN network.

Task 13.2: Advertise the loopbacks of R1, R2, and R3 in the EIGRP process. Use network statements.

Task 13.3: Configure EIGRP on the LAN between R3 and R6.

Task 13.4: On R6, redistribute all the preconfigured loopbacks in the EIGRP process. Use network statements.

Task 13.5: Configure R2 and R3 as stub routers that advertise connected and summary routes.

Task 13.6: R3 should still advertise towards R1 the network 10.11.6.0/24, 10.22.6.0/24, and 10.33.6.0/24.

Task 13.7: Configure EIGRP on the serial connection between R3 and R6.

Task 13.8: Configure EIGRP on the LAN between R1 and R4. Advertise the loopbacks of R4 in the EIGRP process. Use a network statement.

Task 13.9: Configure a distribute-list with prefix-list to prevent R1 from advertising the network 10.1.4.4/32.

Task 13.10: Configure a distribute-list with prefix-list to prevent R1 from learning the network 10.33.6.0/24.

Task 13.11: Configure EIGRP on the connection between R4 and R5. Advertise the loopbacks of R5 in the EIGRP process. Use network statements. Make sure that the traffic is load-balanced on the 2 connections.

Task 13.12: On R4, create a filter based on ACL. R4 should use the Ethernet connection to reach 10.1.5.5/32. Use a standard access-list to achieve this.

Task 13.13: On R4, create filters based on ACL. R4 should use the serial connection to reach 10.11.5.0/24. Use an extended access-list to achieve this.

Task 13.14: Configure EIGRP on the serial connection between R6 and R9. Advertise the loopbacks of R9 in the EIGRP process except loopback 3. Use network statements. Between R3 and R6, make sure that the traffic is load-balanced between the serial interface and the ethernet interface.

Task 13.15: On R6, create a filter based on offset-list. R3 should use the serial 4/2 connection to reach 10.1.9.9/32.

Task 13.16: On R6, create a filter based on offset-list. R3 should use the E0/1 connection to reach 10.11.9.0/24.

Task 13.17: Configure R1 not to install the route 10.11.6.0/24 when received from R3. Manipulate AD.

Task 13.18: Configure R1 not to install the route 10.22.6.0/24 when received from R3. Manipulate AD.

Task 13.19: On R9, there is a preconfigured static route to 172.16.1.0/24. Redistribute this static route into EIGRP and tag this route with a tag of 666.

Task 13.20: Filter on R6 this route out based on the tag 666.

Task 13.21: On the serial connection between R4 and R5, make sure that EIGRP control traffic cannot exceed 25% of the bandwidth.

Task 13.22: The R4 and R5 routers should log EIGRP neighbor relationship changes.

Task 13.23: On R9, configure an EIGRP router-id as 9.9.9.9 and redistribute the loopback3 into EIGRP.

Task 13.24: On R6, configure the EIGRP process to reject the 10.22.9.0/24 network. You are only allowed to change the EIGRP router-id.

Task 13.25: On R6 and R9, configure the EIGRP process to reject EIGRP packets that have transited over more than 10 hops.

You have completed Lab 13

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.

Lab 14: Configure and troubleshoot OSPF (Part 1)

Technologies covered

- DR/BDR
- OSPF network types
- OSPF path selection
- OSPF per neighbor cost
- OSPF auto-cost reference bandwidth
- OSPF version 3 address-family support

Overview

You have been tasked to configure the routing in a network using OSPF.

The topology used in the lab will be the following:

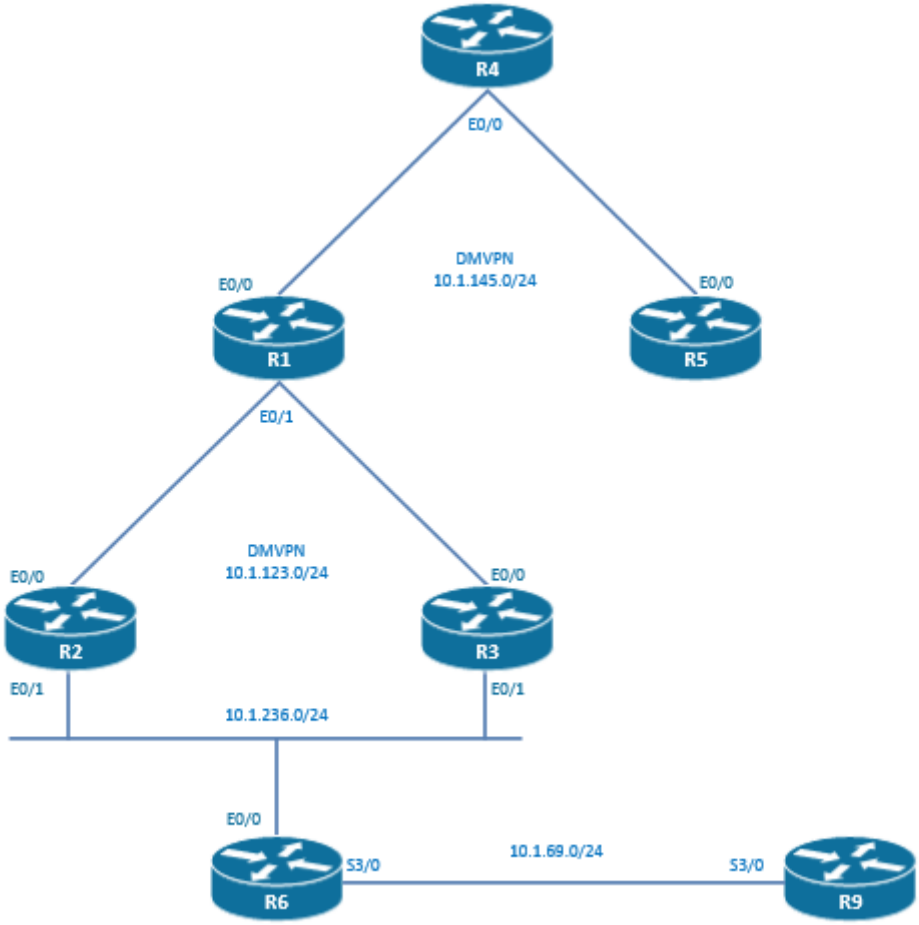

Estimated time to complete: 3-4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 14.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Configure OSPF process 1 area 0 in this network. The election of a DR should not take place. On routers R2 and R3, you are not allowed to change the default network type and not allowed to modify the timers.

Task 14.2: R4, R5, and R1 are also in a hub and spoke topology where R4 is the hub and R1 and R5 are the spokes. DMVPN is the underlying used technology. Configure OSPF process 1 area 0 in this network. The election of a DR should take place in this network. The DR should always be on the hub router. Multicast is not enabled on the DMVPN tunnels.

Task 14.3: On R1, R2, R3, R4, and R5, configure loopbacks 0 as the OSPF router-ids and advertise loopback0 of the routers into OSPF in the following areas:

R1	Area 1
R2	Area 2
R3	Area 3
R4	Area 4
R5	Area 5

Check that you have full reachability between the loopbacks, especially on R2, check that you can ping the loopback of R5 sourcing from the loopback of R2.

Task 14.4: Configure the network 10.1.236.0/24 into area 236 on R2, R3, and R6.

Task 14.5: The R2 should always be elected as the DR, and R3 should always be elected as the BDR.

Task 14.6: Advertise only the loopback 0 of R6 into OSPF area 236. Do not use a network statement. On R5, check that you can ping the loopback of R6 sourcing from the loopback of R5.

Task 14.7: We are going to have links faster than 100M in the network. In the whole OSPF network, a gigaethernet link should have a cost of 1 and a fast ethernet link should have a 10.

Task 14.8: Manipulate the OSPF cost so that R1 prefers R2 over R3 to reach the loopback of R6. Do not configure anything under the interfaces.

Task 14.9: Configure OSPF version 3 area 0 for IPv4 between R6 and R9.

Use the following global unicast addresses:

R6 s3/0	2001: :6/64
R9 s3/0	2001: :9/64

Task 14.10: Create the following IPv4 address loopback1:

R6	20.1.6.6/32
R9	20.1.9.9/32

Task 14.11: Advertise the IPv4 address loopback1 of R6 and R9 into area 0 of the OSPF version 3 processes.

If necessary, use the IPv6 following address for loopback0:

R6	2001:bd8: :6/64
R9	2001:bd8: :9/64

Task 14.12: On R6, make sure that you can ping the loopback of R9 sourcing from the loopback of R6.

You have completed Lab 14

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.

Lab 15: Configure and troubleshoot OSPF (Part 2)

Technologies covered

- Discontiguous area
- Virtual-links
- GRE tunnels
- Non-backbone transit area
- OSPF authentication
- Flood reduction
- Demand circuit
- Summarization
- Discard-route
- Flood reduction

Overview

You have been tasked to configure OSPF as the routing protocol of your network.

The topology used in the lab will be the following:

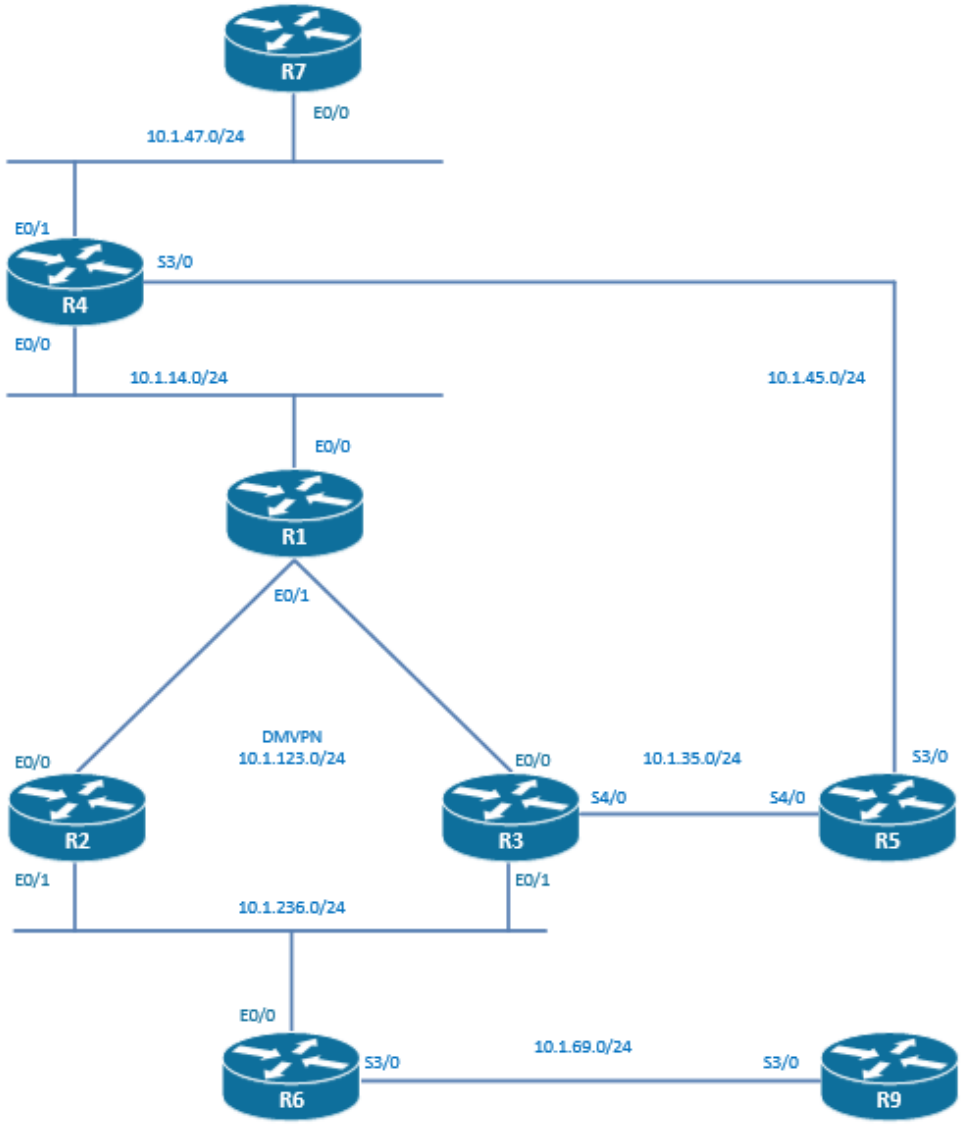

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 15.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Configure OSPF process 1 area 0 in this network. The election of a DR should take place in this network. The DR should always be on the hub router.

Task 15.2: The loopback0 networks of R1, R2, and R3 should present in the OSPF database of R1 as LSAs type 1.

Task 15.3: Configure the network 10.1.236.0/24 into area 236 on R2, R3, and R6. Redistribute only the loopback0 of R6 into the area 236.

Task 15.4: Configure the network 10.1.69.0/24 into area 69 on R6 and R9. Add the loopback0 of R9 into the area 69 process as a network statement.

Task 15.5: Configure area 236 as a stub area.

Task 15.6: Ensure that there is IP connectivity between loopback0 of R9 and the loopback0 of R1. Do not use a virtual-link, as the transit area is a stub area. The path through R3 should be used. Use an IP address of 36.0.0.3/24 and 36.0.0.6/24 when necessary.

Task 15.7: Configure the network 10.1.14.0/24 into area 14 on R1 and R4. Add the loopback0 of R4 into the area 14 process as a network statement.

Task 15.8: Configure the network 10.1.47.0/24 into area 47 on R4 and R7. Add the loopback0 of R7 into the area 47 process as a network statement.

Task 15.9: Ensure that there is IP connectivity between loopback0 of R7 and the loopback0 of R2.

Task 15.10: Configure the network 10.1.35.0/24 to be part of area 0.

Task 15.11: Configure the network 10.1.45.0/24 and the network 10.1.5.5/32 to be part of area 45.

Task 15.12: Configure an OSPF cost of 60000 on the interfaces belonging to the network 10.1.14.0/24.

Task 15.13: On R7, when performing a trace route from the loopback of R7 to the loopback of R3, we can observe that the trace route is following the path R7, R4, R5, and R3. The routing is using a non-backbone area, that is to say area 45, as a transit. Without modifying any OSPF costs, ensure that the trace route is using the R7, R4, R1, and R3 path.

Task 15.14: OSPF should not exchange periodic hellos and periodic refreshes of LSAs over the point-to-point connection between R6 and R9. Configuration can only be applied on R9.

Task 15.15: Configure plain-text authentication on the connection between R6 and R9. The key value should be set to “iPexpert”. Make sure that this authentication is enforced even if this is an on-demand circuit.

Task 15.16: Configure MD5 authentication on the connection between R5 and R3. The key value should be set to 2 and the password to “iPexpert2015”. On R5, configure authentication under the routing process.

Task 15.17: Protect the connection between R5 and R4 with the Null authentication.

Task 15.18: OSPF process is reflooding by default every LSAs every 30 minutes. This should not be necessary for LSAs sent out of the two serial interfaces on R5.

Task 15.19: Configure the following loopbacks on R9:

Loopback8	10.8.9.9/16
Loopback9	10.9.9.9/16
Loopback10	10.10.9.9/16

Task 15.20: Those 3 loopbacks should be seen in the area 0 routing table as a single summary network. Use internal summary.

Task 15.21: On R6, ensure that the summary route created in **Task 15.20** is not present in the routing table pointing to Null0.

Task 15.22: On R9, redistribute the pre-configured routes into OSPF and make sure that they appear as one routing entry in the routing table in all other OSPF routers.

Task 15.23: Configure area 45 in a way that LSAs never age out in this area.

You have completed Lab 15

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.

Lab 16: Configure and troubleshoot OSPF (Part 3)

Technologies covered

- Stub area
- Totally not so stubby area
- NSSA
- NSSA type 5 to type 7 translation
- LSA filtering
- FA Suppression
- Reliable conditional default routing

Overview

You have been tasked to configure OSPF as the routing protocol of your network.

The topology used in the lab will be the following:

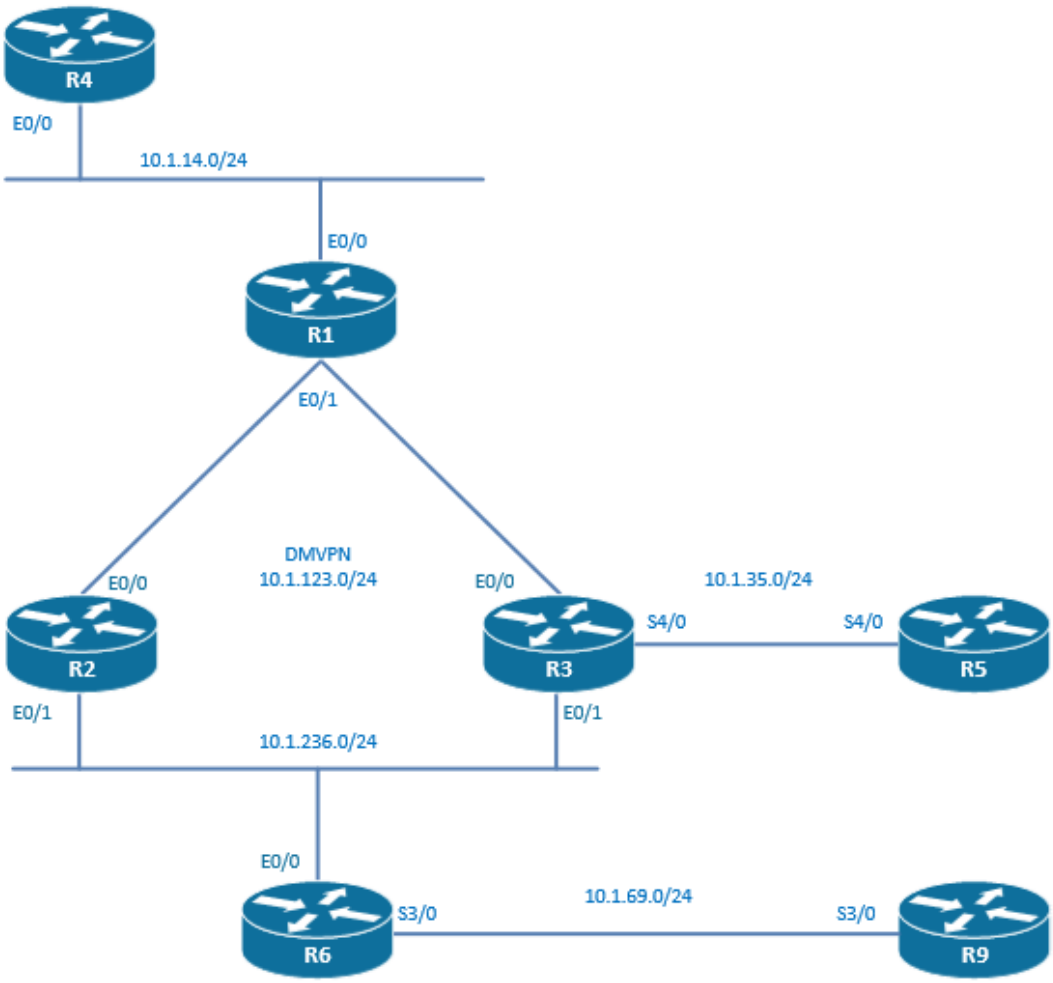

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.
This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 16.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Get OSPF routing up and routing with process 1 area 0 in this DMVPN network. The election of a DR should not take place in this network. Do not modify any OSPF timers.

Task 16.2: Add the loopback0 of R1, R2, and R3 into the area 0 process as network statements.

Task 16.3: On R2, R3, and R6, configure the network 10.1.236.0/24 as part of OSPF area 236. Add the loopback0 of R6 into the area 236 process as a network statement.

Task 16.4: In the R6 routing-table, the only IA OSPF-learned route should be a default route with the ABRs as the next-hop.

Task 16.5: On R6 and R9, configure static routing to ensure the reachability of the loopback0, loopback1, and loopback2 network of R9.

Task 16.6: On R6, redistribute the static routes configured in Task 16.5 (except loopback2) into OSPF. In the routing-table of R1, 10.1.9.0/24 should show as E1 and 10.11.9.0/24 should show as E2. On R1, ensure that you can ping the loopback0 and loopback1 of R9 from the loopback0 of R1 as a source.

Task 16.7: Area 236 is a totally Not-so-stub area having two ABRs to area 0. By manipulating OSPF cost, ensure that the default route in the R6 routing table is using R3 as a next hop. The cost of the default route to R2 should be modified and this cost should be the default cost +1.

Task 16.8: On R1 and R4, configure the network 10.1.14.0/24 as part of OSPF area 14. Add loopback0 of R4 into the area 14 process as a network statement.

Task 16.9: Configure Area 14 in a way that it does not receive any LSA 5 updates. Ensure full reachability and test that you can ping from R4 the loopback 0 of R9 from the loopback0 of R4 as a source.

Task 16.10: Area 35 is a totally NSSA area. On R3 and R5, configure the network 10.1.35.0/24 as part of OSPF area 35. Inject the loopback0 of R5 into the area 35 process as a network statement.

Task 16.11: Redistribute loopback1, loopback2, loopback3, and loopback4 of R5 into the area 35 each as a N2 route and each with a metric of 55. Make sure that on R5, you can ping to the loopback0 of R9 with the ping sourcing from loopback 4 of R5.

Task 16.12: Block the LSA 7 to LSA 5 translation for the network 10.11.5.0/24 using a summary-address command.

Task 16.13: Filter the forwarding address for the type-5 LSAs originated at R5 using the area 35 range no-advertise in command on the ABR.

Task 16.14: Instruct R3 to become the forwarding address itself and check that the IP address reachability is restored, that is to say check that you can ping to the loopback0 of R9 with the ping using as a source the loopback4 of R5.

Task 16.15: On R1, there is a default route pre-configured. This default route should be redistributed into OSPF only if the network 10.21.5.0/24 is present in the routing table of R1. Use IP SLA to track, in a reliable way, this network.

You have completed Lab 16

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.

Lab 17: Configure and troubleshoot OSPF (Part 4)

Technologies covered

- Filtering with distribute-lists
- Filtering with discard-route
- Filtering with administrative distance
- Filtering with route-maps
- NSSA ABR external prefix filtering
- Database filtering
- Stub router advertisement
- OSPF timers optimization
- Resource limiting

Overview

You have been tasked to configure OSPF as the routing protocol of your network.

The topology used in the lab will be the following:

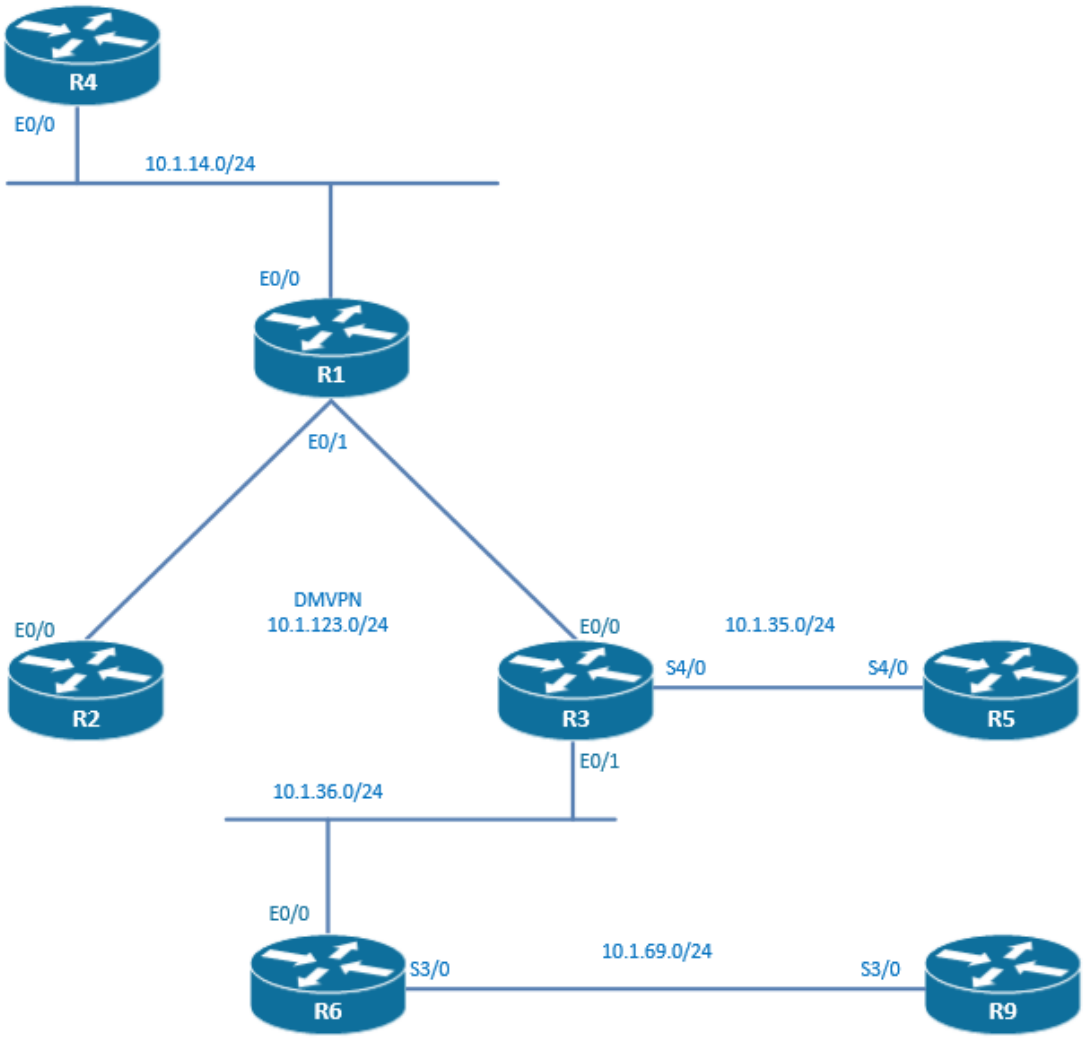

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 17.1: R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN is the underlying used technology. Configure OSPF process 1 area 0 in this network. Use point-to-multipoint network type on the hub and the 2 spokes.

Task 17.2: Configure the network 10.1.36.0/24 into area 0 on R3 and R6.

Task 17.3: The loopback 0 networks of R1, R2, R3, and R6 should present in the OSPF database of R1 as LSAs type 1.

Task 17.4: On R1, prevent the flooding of link-state advertisements to R2 by using the "database-filter all out" command applied to a neighbor. Make sure that R2 is still having full reachability.

Task 17.5: Configure the network 10.1.69.0/24 into area 69 on R6 and R9. Use network statement to advertise loopback0. Distribute loopback1, loopback2, loopback3, and loopback4 of R9 into the area 69 process as E2 type.

Task 17.6: Configure the following router-ids and make sure that they are in use by the process.

R1	1.1.1.1
R2	2.2.2.2
R3	3.3.3.3
R6	6.6.6.6
R9	9.9.9.9

Task 17.7: Ensure that the loopback0 network of R1 is not included by the OSPF process in the routing table of R9. Use prefix-list and distribute-list.

Task 17.8: On R9, the network 10.21.9.9/32 should be filtered out and not be propagated. Use distribute-list and access-list.

Task 17.9: On R3, configure a default route pointing to R5. On R5, configure a default route pointing to R3. Confirm that you can ping from R3 the loopback0 of R5 10.1.5.5 from the loopback 0 of R3.

Task 17.10: Redistribute this default route into OSPF area 0.

Task 17.11: On the ABR R6, configure the area 0 to advertise a summary network of 10.1.0.0/16 within the area 69.

Task 17.12: Try to ping loopback0 of R5 from loopback0 of R9. Because of the presence of a 10.1.0.0/16 route on the ABR, the default route is not being used and the ping is failing. Ensure that this 10.1.0.0/16 is suppressed.

Task 17.13: On R1, the network 10.41.9.9/32 should be present in the OSPF database but not in the routing table. Manipulate the administrative distance to achieve this.

Task 17.14: Configure R6 so that R1 doesn't receive the 10.1.9.9/32 prefix. Use prefix-list and area filter-list.

Task 17.15: Configure a NSSA area 14 between R1 and R4. On R4, redistribute all connected interfaces into OSPF.

Task 17.16: On R1, filter the 10.11.4.4/32 and 10.22.4.4/32 out and let the other networks coming from area 14 advertise to the area 0. Use summary-address command.

Task 17.17: Configure on all the routers the feature that will remove the transit networks from the OSPF database. Check that IP reachability is still working between the OSPF advertised prefixes once this feature is enabled.

Task 17.18: On R9, configure the minimum interval for accepting the same LSA to 80 ms.

Task 17.19: On R9, set the following rate-limit values for LSA advertisement:

Start-interval	10 ms
Hold-interval	100 ms
Max-interval	5000 ms

Task 17.20: On R9, configure OSPF throttling timers:

Spf-start	10 ms
Spf-hold	4800 ms
Spf-max-wait	90000 ms

Task 17.21: On R9, configure OSPF Update flood packet-pacing to 5 ms.

Task 17.22: On R9, in order to improve convergence, enable incremental SPF.

Task 17.23: R9 should fire up a syslog message when more than 3 prefixes are redistributed. First warning should be sent when 80% of the threshold is reached.

Task 17.24: On R9, limit to 1000 the number of nonself-generated LSAs the OSPF routing process can keep in the OSPF database.

You have completed Lab 17

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.

Lab 18: Configure and troubleshoot BGP (Part 1)

Technologies covered

- EBGP peering
- EBGP multihop
- EBGP Disable-connected-check
- Update source
- iBGP peering
- Route Reflector

Overview

You have been tasked to configure the routing in your network using OSPF, EIGRP, and RIP, static route, iBGP and eBGP.

The topology used in the lab will be the following:

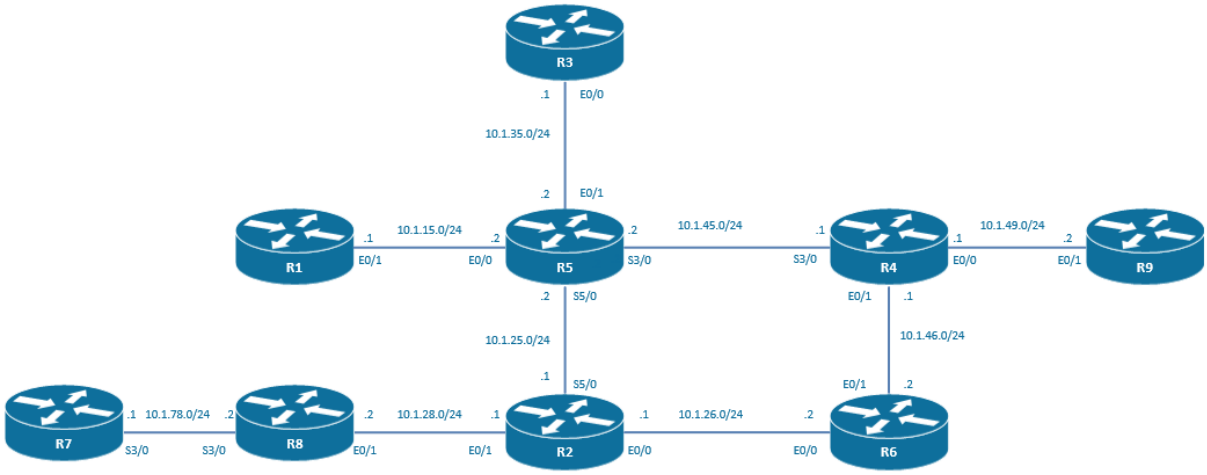

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 18.1 Routing between R1 and R5 should be configured with RIP version 2. Loopback0 reachability has to be achieved thanks to this protocol.

Task 18.2 Configure an eBGP peering between R1 in AS 1 and R5 in AS 65001. This peering should be established between the loopback0 of each router.

Task 18.3 On the peering between R1 and R5, do not use the ebgp multihop command.

Task 18.4 Advertise the loopback0 of R1 in BGP using a network statement.

Task 18.5 Routing between R3 and R5 should be configured with static routes. Loopback0 reachability has to be achieved thanks to this protocol.

Task 18.6 Configure an eBGP peering between R3 in AS 3 and R5 in AS 65001. This peering should be established between the loopback0 of each router. On the peering between R3 and R5, use the ebgp multihop command.

Task 18.7 Advertise the loopback0 of R3 in BGP using a network statement.

Task 18.8 Routing between R2 and R7 should be configured with EIGRP. Loopback0 reachability has to be achieved thanks to this protocol.

Task 18.9 Configure an eBGP peering between R2 in AS 65001 and R7 in AS 7. This peering should be established between the loopback0 of each router. Use the minimum number of hops necessary in the ebg-multihop command.

Task 18.10 Advertise the loopback1 of R7 in BGP using a network statement. Check that you can ping from R1 to the loopback1 of R7. Use of static routes on R8 is required.

Task 18.11 Configure OSPF area 0 on the R2 to R5 connection. Advertise the loopback0 of R2, R5, and into OSPF.

Task 18.12 Configure iBGP peering between R2 and R5. This peering should be established between the loopback0 of each router. Make sure that the ping from R3 to R7 is up and running. Do not use the redistribute command into BGP at this point of the lab.

Task 18.13 Enable synchronization on R5. Using a route-map and a prefix-list, redistribute BGP into OSPF on R2. The full IP reachability should be established between the loopback0 of R1, R3, and R7.

Task 18.14 Configure OSPF area 0 on the R5 to R4 connection. Advertise the loopback0 of R4 into OSPF.

Task 18.15 Routing between R4 and R9 should be configured with static routes. Loopback0 reachability has to be achieved thanks to this protocol.

Task 18.16 Configure an eBGP peering between R4 in AS 65001 and R9 in AS 9. This peering should be established between the loopback0 of each router.

Task 18.17 Advertise the loopback0 of R9 in BGP using a network statement.

Task 18.18 Configure iBGP peering between R4 and R2. This peering should be established between the loopback0 of each router. Configure R2 as a route-reflector for R4 and R5.

Task 18.19 On R7, make sure that you can ping from the loopback1 of R7 to the loopback0 of R1, R3, and R9.

Task 18.20 Configure OSPF area 0 on the R2 to R6 connection and on the R4 to R6 connection. Advertise the loopback0 of R6 into OSPF.

Task 18.21 For redundancy, configure R2 and R6 as part of a RR cluster with cluster-id 1.

You have completed Lab 18

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.

Lab 19: Configure and troubleshoot BGP (part 2)

Technologies covered

- Next-hop-self
- BGP next-hop with route-map
- BGP Confederation
- GRE tunnels

Overview

You have been tasked to configure the routing in your network using OSPF, EIGRP, RIP, static route, iBGP, and eBGP.

The topology used in the lab will be the following:

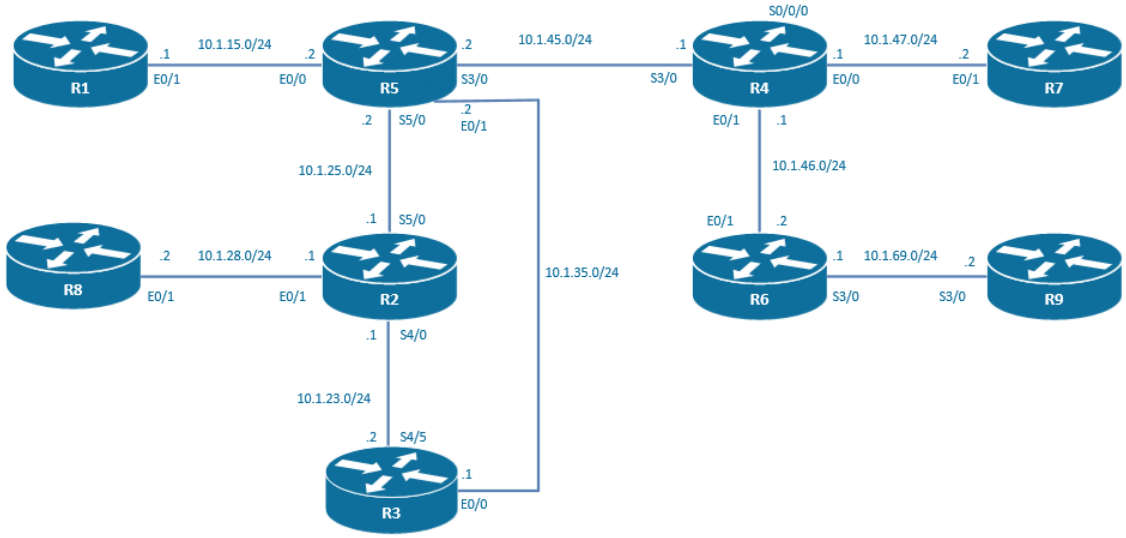

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 19.1	Routing between R4 and R7 should be configured with EIGRP. Loopback0 reachability has to be achieved thanks to this protocol.
Task 19.2	Configure an eBGP peering between R4 in AS 65019 and R7 in AS 7. This peering should be established between the loopback0 of each router.
Task 19.3	Advertise loopback0 of R7 in BGP using a network statement.
Task 19.4	Routing between R6 and R9 should be configured with static routes. Loopback0 reachability has to be achieved thanks to this protocol.
Task 19.5	Configure an eBGP peering between R6 in AS 65019 and R9 in AS 9. This peering should be established between loopback0 of each router.
Task 19.6	Advertise loopback0 of R9 in BGP using a network statement.

Task 19.7 Configure OSPF area 0 only between R4 and R6. Advertise the loopback0 of R4 and R6 into OSPF using a network statement. Do not advertise anything else into OSPF.

Task 19.8 Configure iBGP between R4 and R6.

Task 19.9 Use next-hop-self to enable IP connectivity between loopback0 of R7 and the loopback0 of R9.

Task 19.10 Routing between R5 and R1 should be configured with RIP. Loopback0 reachability has to be achieved thanks to this protocol.

Task 19.11 Configure eBGP peering between R1 in AS 1 and R5 in AS 65019. This peering should be established between the loopback0 of each router.

Task 19.12 Advertise loopback0 of R1 in BGP using a network statement.

Task 19.13 Routing between R8 and R2 should be configured with EIGRP. Loopback0 reachability has to be achieved thanks to this protocol.

Task 19.14 Configure eBGP peering between R2 in AS 65019 and R8 in AS 8. This peering should be established between the loopback0 of each router.

Task 19.15 Advertise loopback0 of R8 in BGP using a network statement.

Task 19.16 Configure OSPF area 0 only between R5 and R2. Advertise the loopback0 of R5 and R2 into OSPF using a network statement. Do not advertise anything else into OSPF.

Task 19.17 Configure an OSPF cost of 10 on this link.

Task 19.18 Configure iBGP between R5 and R2.

Task 19.19 Use a route-map to enable the IP connectivity between loopback0 of R1 and the loopback0 of R8.

Task 19.20 Configure OSPF area 0 between R5 and R4.

Task 19.21 R2 and R5 are part of confederation with ID 25, R6, and R4 are part of confederation with ID 46.

Task 19.22 Configure the confederation ID 25 and 46 to be part of AS 65019. Ensure full reachability between R1, R7, R8, and R9. As an example, you should be able to ping from R8 to loopback0 of R7with the ping sourced from the loopback0 of R8. Use of 2 static routes is allowed.

Task 19.23 Configure OSPF area 0 on the connection between R5 and R3 with an OSPF cost of 1.

Task 19.24 Configure OSPF area 0 on the connection between R2 and R3 with an OSPF cost of 1.

Task 19.25 Restore the IP connectivity between R8 and R1, R8 and R7, R8 and R7, and R8 and R9. You are not allowed to redistribute BGP routes into OSPF. Use the network 10.1.145.0/24 for the tunnel interfaces. Check that you are again able to ping from R8 to loopback0 of R1 with the ping sourced from the loopback0 of R8.

You have completed Lab 19

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.

Lab 20: Configure and troubleshoot BGP (part 3)

Technologies covered

- Weight
- Local Preference
- As-path prepending
- Origin
- MED
- Always compare MED
- AS-path ignore
- Maximum AS Limit

Overview

You have been tasked to configure the routing in your network using OSPF, EIGRP, RIP, static route, iBGP, and eBGP.

The topology used in the lab will be the following:

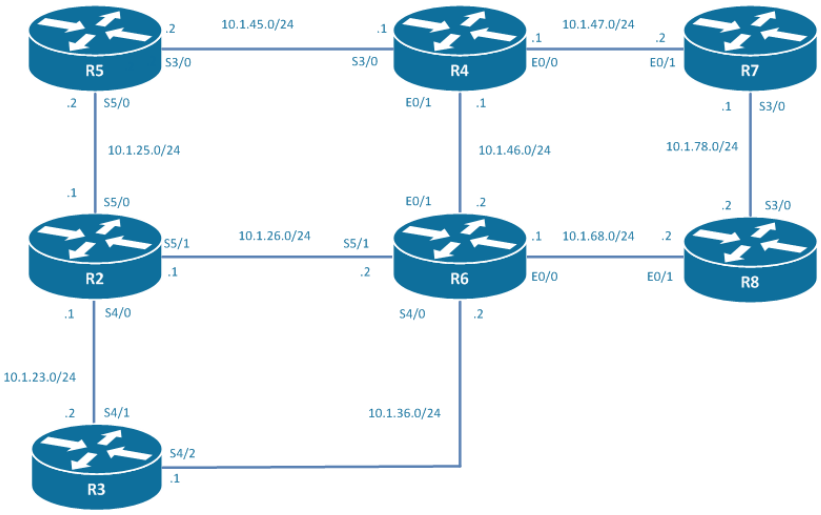

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 20.1 Configure an iBGP peering between R4 and R7 in AS 65001. Make sure that the 10.1.46.0/24 network and that the network 10.1.78.0/24 is carried in the BGP updates with an origin of i.

Task 20.2 Configure an eBGP peering between R4 in AS 65001 and R6 in AS 65002.

Task 20.3 Configure an eBGP peering between R6 in AS 65002 and R8 in AS 8.

Task 20.4 Configure an eBGP peering between R8 in AS 8 and R7 in AS65001.

Task 20.5 The loopback0 of R4 should be present in the BGP database with an origin attribute of incomplete. The loopback0 of R7 should be present in the BGP database with an origin attribute of internal.

Task 20.6 On R8, manipulate the weight attribute so that the route to 10.1.4.4/32 is pointing towards R6. Use a prefix-list called WEIGHT_PL and a route-map called WEIGHT_RM.

Task 20.7 The loopback0 of R6 should be present in the BGP database with an origin attribute of incomplete.

Task 20.8 In order to reach the 10.1.6.6/32 loopback, the routers in AS 65001 should route the traffic over R8 through AS 8. Change the configuration in R7 and use a route-map called LOCALPRF_RM.

Task 20.9 The loopback0 of R8 should be present in the BGP database with an origin attribute of IGP.

Task 20.10 Configure R8 so that the traffic originated on R6 is going through AS 65001 to reach the network 10.1.8.8/32. On R6, 10.1.8.8 should have the following AS-pat attribute 8 8 8 8 i. Use a prefix-list called PREPEND_PL and a route-map called PREPEND_RM.

Task 20.11 Configure OSPF area 0 between R6 and R2.

Task 20.12 Configure iBGP connection between R6 and R2.

Task 20.13 The loopback0 of R2 should be present in the BGP database with an origin attribute of incomplete.

Task 20.14 Configure an eBGP connection between R6 and R3 in AS 3. Redistribute the EBGP next-hop in OSPF area 0.

Task 20.15 Configure an eBGP connection between R2 and R3 in AS 3. Redistribute the EBGP next-hop in OSPF area 0.

Task 20.16 Advertise loopback0 and loopback1 of R3 using network statements.

Task 20.17 Ensure that the traffic is routed via R2 to reach network 10.1.3.3. Configure R3 and use the prefix-list called MED_PL 2 and a route-map called MED_RM2. Use a MED value of 200.

Task 20.18 Ensure that the traffic is routed via R6 to reach network 10.11.3.3. Configure R3 and use the prefix-list called MED_PL 6 and a route-map called MED_RM6. Use a MED value of 300.

Task 20.19 In R2 and R6, advertise the network 10.1.26.0/24 with a network statement.

Task 20.20 On R3, modify the origin of route 10.1.26.0/24 and ensure that this route is reached primarily through R6. Use a prefix-list called ORIGIN_PL and a route-map called ORIGIN_RM.

Task 20.21 On R6, advertise the network 10.22.6.0/24 using a network statement.

Task 20.22 This network should be advertised to the router R4 using the MED 500 and prepending one more AS in the AS-path. Use a prefix-list called ALWAYSCOMPMED_PL and a route-map called ALWAYSCOMPMED_RM.

Task 20.23 Configure R4 and ensure that R4 always prefers the route with the lowest MED, that is to say the route to R6 is pointing to R7 on R4.

Task 20.24 Configure an eBGP connection between R2 and R5 in AS 5 and between R4 and R5 in AS 5. Advertise the loopback of R5 into BGP with an origin of "?".

Task 20.25 On R4, prepend the AS 65001 4 times when advertising the network 10.1.7.7/32 to R5. The route from R5 to the loopback0 should now be transiting through AS 65002.

Task 20.26 On R5, the AS-path attribute should be ignored and the route to the 10.1.7.7/32 network has to point towards R4 and not transit through AS 65002 anymore. Use MED to achieve this.

Task 20.27 On the peering between R2 and R5, shut down the peering if more than 50 BGP updates are advertised from R5 to R2. A syslog message should be sent when more than 40 BGP updates are advertised from R5 to R2.

You have completed Lab 20

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.

Lab 21: Configure and troubleshoot BGP (part 4)

Technologies covered

- Aggregation
- Summary-only
- Suppress-map
- Unsuppress-map
- AS-set
- Attribute-map
- Advertise-map
- Community no-export
- Community local-AS
- Community no-advertise

Overview

You have been tasked to configure the routing in your network using OSPF, EIGRP, and RIP, static route, iBGP and eBGP.

The topology used in the lab will be the following:

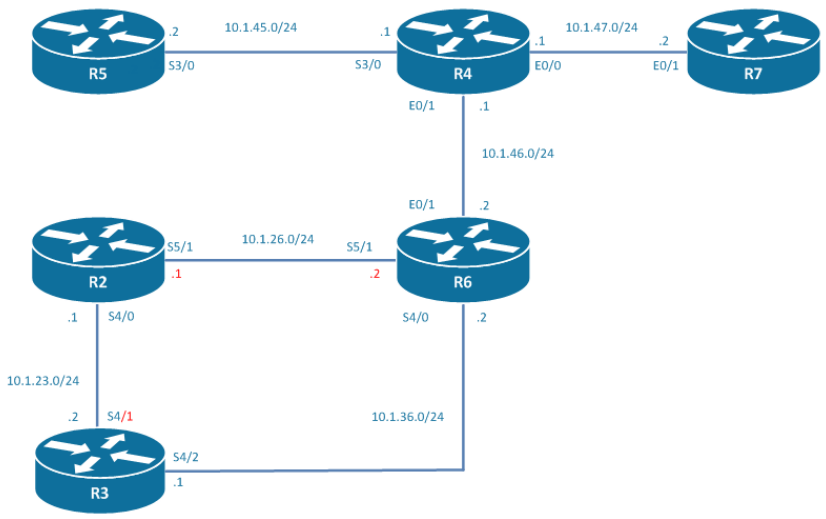

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 21.1 Configure an iBGP peering between R2 and R6 in AS 65001.

Task 21.2 Configure an eBGP peering between R3 in AS 3 and R6.

Task 21.3 Configure an eBGP peering between R3 in AS 3 and R2.

Task 21.4 R3 has to advertise a summary route representing the loopback1, loopback2, loopback3 and the loopback4 addresses of R3. The aggregate address command cannot be used.

Task 21.5 R3 has to advertise a summary route representing the loopback11, loopback12, loopback13, and the loopback14 addresses of R3. More specific networks should also be advertised. Use redistribution and a prefix-list with one single line.

Task 21.6 R3 has to advertise a summary route representing loopback21, loopback22, loopback23, and loopback24 addresses of R3. Specific subnets should not be advertised. Use network statements.

Task 21.7 In the addition to the summary route, loopback21 network should be the only specific network advertised towards R2. Use an unsuppress-map.

Task 21.8 In the addition to the summary route, loopback22 network should be the only specific network advertised towards R6. Use an unsuppress-map.

Task 21.9 In the addition to the summary route, loopback14 network should be the only specific network advertised towards R2. Use a suppress-map.

Task 21.10 Configure an eBGP peering between R4 in AS 4 and R6 in AS 65001.

Task 21.11 Configure an eBGP peering between R4 in AS 4 and R7 in AS 7. Advertise the network 200.1.1.0/24 into BGP using a network statement.

Task 21.12 Configure an eBGP peering between R4 in AS 4 and R5 in AS 5. Advertise the network 200.2.1.0/24 into BGP using a network statement.

Task 21.13 On R4, configure the aggregate 200.0.0.0/14. The more specific networks should not be advertised to R6. This aggregate should have in its AS-path attribute all the ASs that were contained in the AS-path attribute of the more specific networks.

Task 21.14 On R3, advertise the networks 153.153.153.0/24 and 153.153.154.0/24 into BGP using network statements.

Task 21.15 On the peerings with R2 and R6, the network 153.153.153.0/24 has to be sent with the No-Export community. Use a route-map called NOEXPORT_RM.

Task 21.16 On R6, configure an aggregate for the network 153.153.152.0/22 with the summary-only and with the AS-SET option on.

Task 21.17 Ensure that this aggregate is advertised to R4. Use a route-map called ATTRIBUTEMAP_RM.

Task 21.18 On R5, advertise the networks 200.200.0.0/16 and 200.201.0.0/16 into BGP using network statements.

Task 21.19 When advertising out the network 200.200.0.0/16 to R4, configure the community of no-advertise.

Task 21.20 On R4, configure an aggregate for the network 200.0.0.0/8 with the summary-only and with the AS-SET option on.

Task 21.21 Ensure that the network 200.0.0.0/8 will be advertised to R7 and R6. You are not allowed to use an attribute-map to remove the community. Use a route-map called ADVERTISEMAP_RM.

Task 21.22 On R4, advertise the network 10.22.4.0/24 into BGP using a network statement.

Task 21.23 Ensure that the network 10.22.4.0 will be advertised to R6 with a community that will prevent it to be advertised to other eBGP peers.

You have completed Lab 21

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.

Lab 22: Configure and troubleshoot BGP (part 5)

Technologies covered

- Local AS
- Replace AS
- Dual AS
- Remove Private AS
- Dampening
- ORF
- BGP allowas-in

Overview

You have been tasked to configure the routing in your network using OSPF, EIGRP, RIP, static route, iBGP, and eBGP.

The topology used in the lab will be the following:

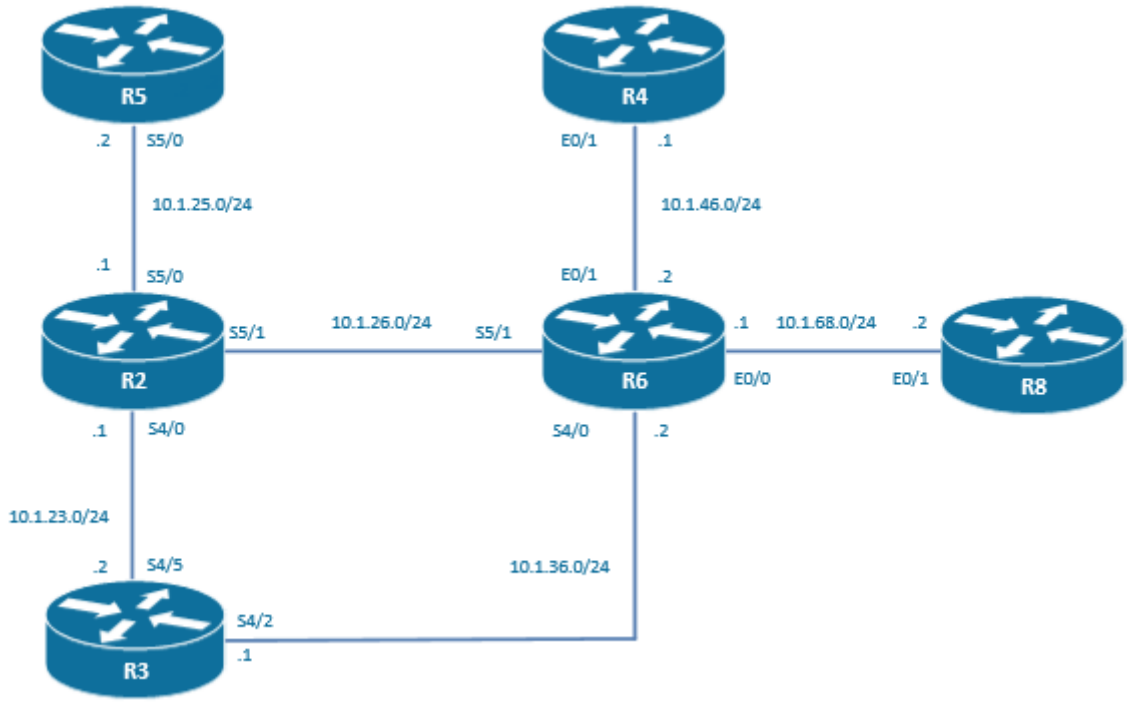

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 22.1 Configure an iBGP peering in AS 65001 between R2 and R6.

Task 22.2 Configure an eBGP peering between R6 and R3 in AS 3.

Task 22.3 Configure an eBGP peering between R2 and R3 in AS 3.

Task 22.4 On R3, redistribute the network 153.153.153.0/24 and the network 153.153.154.0/24 using network statements.

Task 22.5 On R3, on the peering between R3 and R2, filter out 153.153.153.0/24. Use prefix-list.

Task 22.6 On R3, on the peering between R3 and R6, filter out 153.153.154.0/24. Use access-list.

Task 22.7 R3 should appear to R2 and R6 as if it is using AS 65003 but R3 should still be in AS 3.

Task 22.8 Regarding the routes advertised from R3 to R2, the AS 65003 should not appear in the AS-path.

Task 22.9 Configure an eBGP peering between R2 and R5 in AS 5.

Task 22.10 R5 should appear to R2 and R6 as if it is using AS 65005 but R5 should still be in AS 5.

Task 22.11 Advertise the loopback0 of R5 into BGP. Regarding the routes advertised from R5 to R2, the AS 65005 should not appear in the AS-path.

Task 22.12 On R5, the route 153.153.154.0/24 should not contain the AS 3 as well as the AS 65003 in the AS-path.

Task 22.13 Configure an eBGP peering between R6 and R4 in AS 4.

Task 22.14 On R6, in all advertisements sent toward R4, the private AS numbers have to be stripped off from the AS-path before being sent.

Task 22.15 On R3, configure the 153.153.153.0/24 network to use the following dampening parameters:

- Max-Suppress=60 minutes
- Suppress=2000 points
- Reuse=800 points
- Half-Time=15 minutes

Task 22.16 On R3, configure the 153.153.154.0/24 network to use the following dampening parameters:

- Max-Suppress=50 minutes
- Suppress=2500 points
- Reuse=600 points
- Half-Time=10 minutes

Task 22.17 Between R6 and R4, configure the BGP peering to use fast session deactivation.

Task 22.18 On R4, advertise the loopbacks in BGP using network statements.

Task 22.19 On R6, filter in the network 10.11.4.0/24 on the peering towards R4.

Task 22.20 Make sure that the two routers exchange information via the ORF capability and that R4 will be filtering the network 10.11.4.0/24 and not sending updates for networks that are filtered when arriving on R6.

Task 22.21 Configure an eBGP peering between R6 and R8 in AS 4.

Task 22.22 On R8, advertise the loopback0 into BGP using a network statement.

Task 22.23 Make sure that you can ping from loopback0 of R8 which is originated in AS 4 to the loopback0 of R4 which is always originated in AS 4. Use the allowas-in command.

You have completed Lab 22

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.

Lab 23: Configure and troubleshoot Multiprotocol Label Switching (Part 1)

Technologies covered

- IPv4 VPN address-family
- LSP
- LDP
- L3VPN
- CE
- PE
- P
- Export map

Overview

You have been tasked to configure a MPLS L3 VPN service on an existing MPLS backbone. The CEs are managed by the Service Provider and the loopbacks of the CEs should be leaked from the VRF of the customer into the management VRF of the Service provider.

The topology used in the lab will be the following:

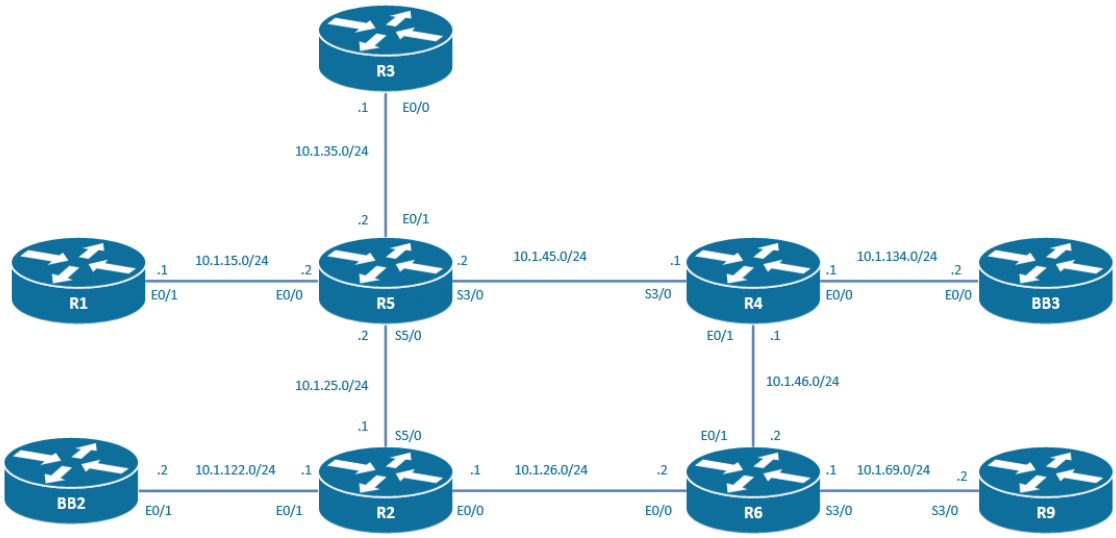

Estimated time to complete: 3-4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 23.1 The network is pre-configured with OSPF and LDP and the PEs are the R5, R4, R6, and R2 routers. In order to optimize the building of the MPLS forwarding-table, make sure that only LSPs for the loopback interfaces will be built.

Task 23.2 Configure the following L3 MPLS VPN routing tables on the R5 and on the R6:

AS	VPN name	rd	rt export	rt import
1	Customer_A	1	10	10
1	Customer_B	2	20	20

Task 23.3 Configure the following loopbacks for the VPN Customer_A and Customer_B.

R5	Loopback10	10.10.5.5/32	Customer_A
R5	Loopback20	10.20.5.5/32	Customer_B
R6	Loopback10	10.10.6.6/32	Customer_A
R6	Loopback20	10.20.6.6/32	Customer_B

Task 23.4 Configure the BGP routing sessions that will permit to exchange the VPNv4 information between the PEs. Use BGP AS 1.

Task 23.5 Redistribute the loopbacks created in the Task 23.3 in their respective VPNs and check that you can ping from loopback to loopback within the same VPN.

Task 23.6 Make sure that the loopbacks redistributed at PE router R5 has a known origin.

Task 23.7 Customer_ A and Customer _B companies are merging.

Task 23.8 The engineer was too quick and the merge between Customer_A and Customer_B is not going ahead.

Task 23.9 Configure R1 and R9 to be part of VRF Customer_A and R3 to be part of VRF Customer_B.

Configure the following loopbacks:

R1 loopback0	10.1.1.1/32
R9 loopback0	10.1.9.9/32
R3 loopback0	10.1.3.3/32

Task 23.10 Route the loopback0 interfaces of the CEs statically and make sure that those loopbacks are routed in their respective VRF. Verify that R1 loopback0 can ping R9 loopback0.

Task 23.11 The service provider is offering a service where the CEs are managed. Customer_A has chosen a managed service for its CEs. The management CE of the Service provider is the router called BB2. Create the management VRF on the router R2.

AS	VPN name	rd	rt export	rt import
1	SP_Management	100	1000	1000,1001

Task 23.12 The management network is using the network 192.168.1.128/25. Create on BB2 a loopback 100 with the following IP address: 192.168.1.129/25 and route it statically into the SP_Management VPN.

Task 23.13 Configure the multi-protocol BGP environment to enable the exchange of the RT information. As we are using iBGP, we create a full-mesh peering topology between R2, R5, and R6.

Task 23.14 The R1 CE and the R9 CE from Customer A has to be reachable from the service provider management network. Use an export map called CE_Loopback_Export on R5 and on R6, and make sure that the management network can only see the loopback of R1 and R9.

You have completed Lab 23

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.

Lab 24: Configure and troubleshoot Multiprotocol Label Switching (Part 2)

Technologies covered

- PE-CE static routing
- PE-CE RIP routing
- PE-CE OSPF routing
- OSPF Domain-ID
- OSPF sham-link
- PE-CE EIGRP routing
- EIGRP SoO

Overview

You have been tasked to configure a MPLS L3 VPN service on an existing MPLS backbone. You will have to configure the routing between the CEs and the PEs for two customer L3 VPNs.

The topology used in the lab will be the following:

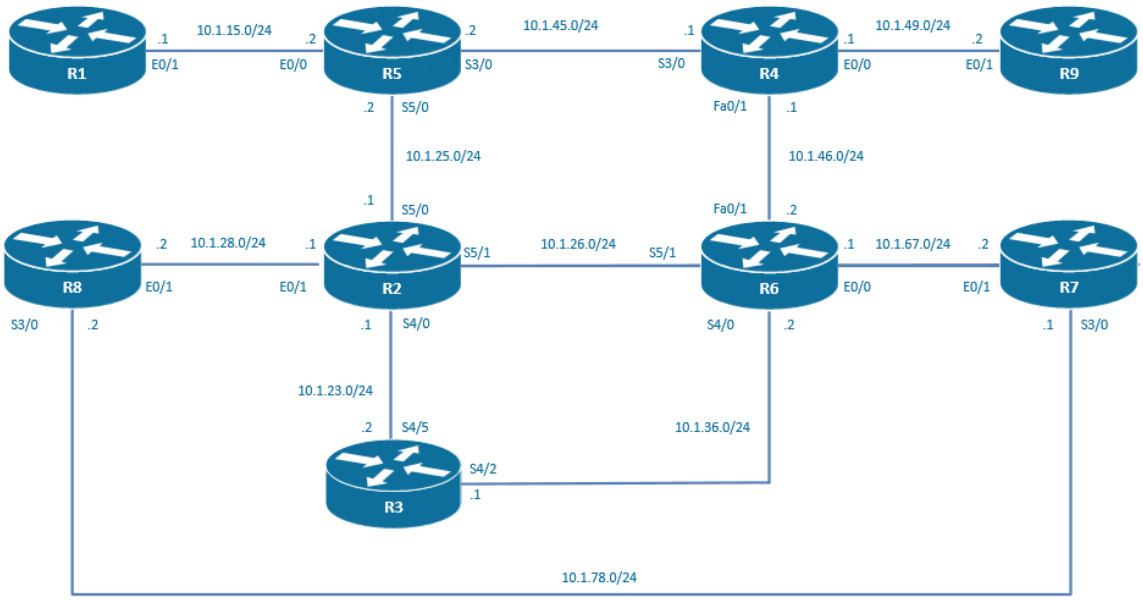

Estimated time to complete: 3-4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 24.1 Configure R5, R4, R6, and R2 as PE routers. The MPLS cloud is using BGP AS 1. Establish MP-BGP sessions between the PEs. Use the loopbacks 0 for the source of the peerings. Use R4 as a route-reflector for all the PEs.

Task 24.2 Create the following L3 VPNs on all PEs.

AS	VPN name	rd	rt export	rt import
1	Customer_A	10	10	10
1	Customer_B	20	20	20

Task 24.3 Configure the following loopbacks for the VPN Customer_A and Customer_B. Make sure that the loopbacks are routed in the VPN MPLS cloud using network statements.

R5	Loopback15	10.10.5.5/32	Customer_A
R5	Loopback25	10.20.5.5/32	Customer_B
R6	Loopback16	10.10.6.6/32	Customer_A
R6	Loopback26	10.20.6.6/32	Customer_B
R2	Loopback12	10.10.2.2/32	Customer_A
R2	Loopback12	10.20.2.2/32	Customer_B
R4	Loopback14	10.10.4.4/32	Customer_A
R4	Loopback14	10.20.4.4/32	Customer_B

Task 24.4 Make sure that you have full reachability between Lo15, Lo16, Lo12, and Lo14 in VPN Customer_A.

Task 24.5 Make sure that you have full reachability between Lo25, Lo26, Lo22, and Lo24 in VPN Customer_B.

Task 24.6 R1 is a CE in VRF Customer_A. The loopback of the router R1 should be routed statistically within the VPN Customer_A.

Task 24.7 R9 is a CE in VRF Customer_B. The loopback of the router R9 should be routed using RIP version 2 within the VPN Customer_B. Do not redistribute BGP into RIP.

Task 24.8 R7 is a CE connected to PE R6 in VRF Customer_A. The loopback of the router R7 should be routed using OSPF process ID 7 in area 0 within the VPN Customer_A. Ensure that you have IP reachability between lo0 of R1 and lo0 of R7.

Task 24.9 R8 is a CE connected to PE R2 in VRF Customer_A. The loopback of the router R8 should be routed using OSPF process ID 8 in area 0 within the VPN Customer_A. Ensure that you have IP reachability between lo0 of R7, R8, and R1.

Task 24.10 On R8, the network 10.1.7.0/24 should be present in the OSPF database as a LSA type 3. If necessary, use a domainID of 78.

Task 24.11 Configure the connection between R7 and R8 in OSPF area 0 with an IP ospf cost of 4000.

Task 24.12 Make sure that the path over the MPLS backbone is the preferred path for traffic going from R7 to R8. Use the loopback22 with IP address 2.2.2.2/32 on R2. Use the loopback66 with IP address 6.6.6.6/32 on R6.

Task 24.13 R3 is a CE connected to PE R2 in VRF Customer_B. The loopback of the router R3 should be routed using EIGRP ID 1 with AS 200 within the VPN Customer_B. Use metric 1 1 1 1 1 when redistributing BGP into EIGRP on the PE. Ensure that you have IP reachability between lo0 of R9 and lo0 of R3.

Task 24.14 R3 is a CE connected to PE R6 in VRF Customer_B. Routing between R3 and R6 is using EIGRP ID 1 with AS 200. Use metric 1 1 1 1 1 when redistributing BGP into EIGRP on the PE.

Task 24.15 By using the extended community 1:11 and 1:12, ensure that it is not allowed that an EIGRP route that has been distributed into BGP on R2 cannot be learnt via R6 when BGP is redistributed into EIGRP on R6, and vice-versa.

You have completed Lab 24

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.

Lab 25: Configure and troubleshoot Ipsec Virtual Private Networks

Technologies covered

- GRE tunnels
- IPsec tunnels
- GRE over IPsec
- IPsec VTIs

Overview

You have been tasked to configure an IPsec encryption on different connections of your network.

The topology used in the lab will be the following:

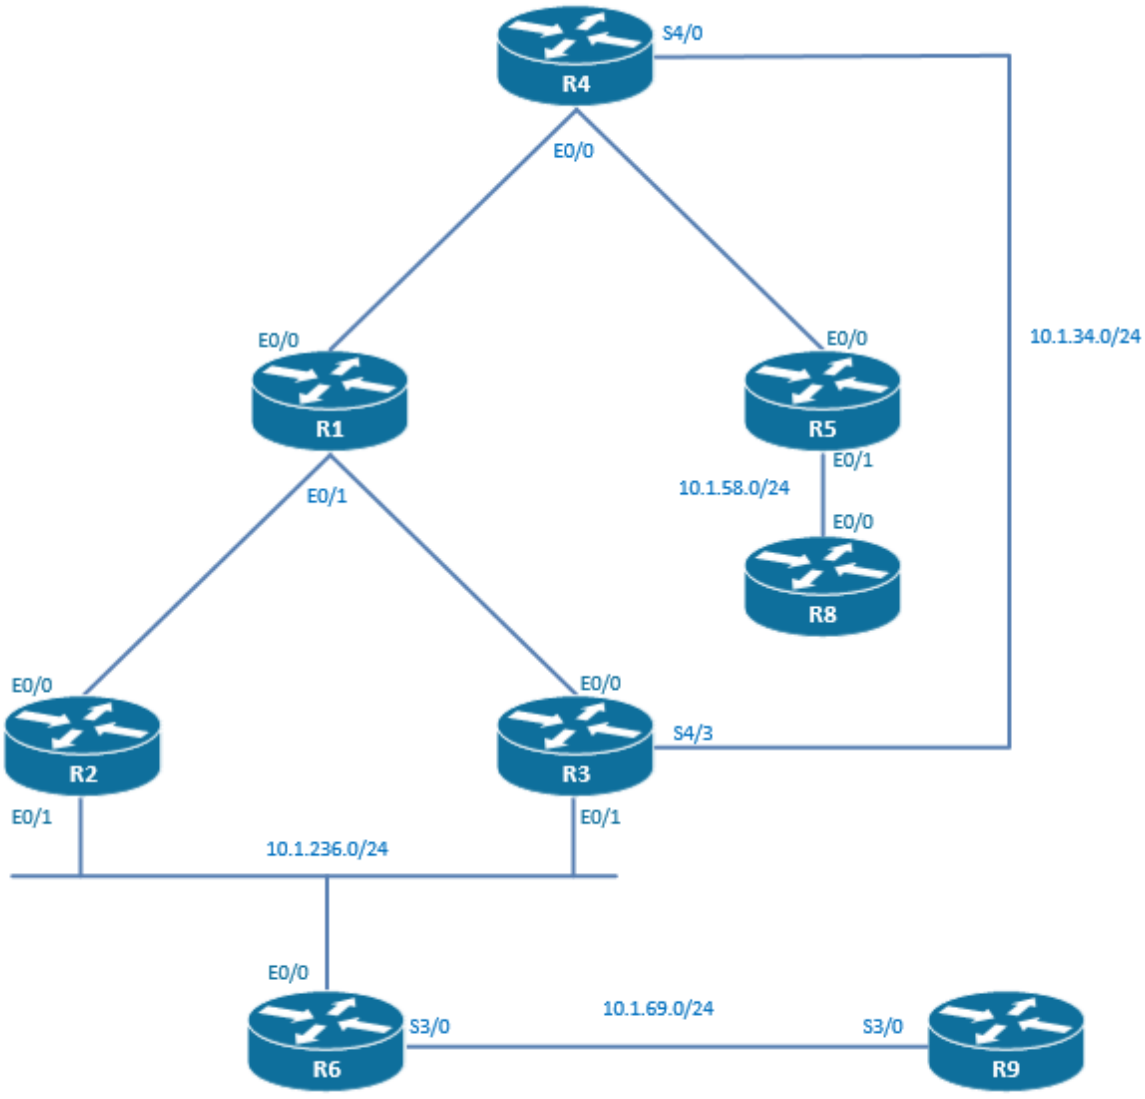

Estimated time to complete: 3-4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 25.1 Configure a LAN-to-LAN IPsec tunnel on the serial connection between R4 and R3. Use a hash of MD5 and pre-shared key of "iPexpert" during the phase 1 negotiation.

Task 25.2 Between R4 and R3, use esp-des encryption and an esp-md5-hmac authentication during the phase 2 negotiation.

Task 25.3 Traffic going from loopback0 of R4 to loopback0 of R5 should be encrypted in both directions. You are not allowed to use a dynamic routing protocol or a default route.

Task 25.4 Configure a GRE tunnel on the serial connection between R2 and R9. The tunnel1 interface has an IP address of 192.168.29.2/24 on R2 and an IP address of 192.168.29.9/24 on R9. Use the E0/1 of R2 and S3/0 of R9 as source/destination of the tunnel. You are not allowed to configure anything on the R6 router.

Task 25.5 You are not allowed to use a dynamic routing protocol or a default route. Traffic going from loopback0 of R2 to loopback0 of R9 should transit through this GRE tunnel.

Task 25.6 There is a Web server which is connected to a client and the traffic is running over Tunnel 1. The client cannot communicate with the server. The web server is sending IP packets with a size of 1500 bytes and the DF-bit set. Configure the tunnel to restore connectivity between the server and the client. You are not allowed to clear the DF-bit or to intervene in the TCP negotiation.

Task 25.7 Encrypt the GRE traffic tunnel between R2 and R9. Use a GRE over IPsec tunneling. Use a hash of MD5 and pre-shared key of "iPexpert" during the phase 1 negotiation.

Task 25.8 Between R2 and R9, use esp-3des encryption and an esp-md5-hmac authentication during the phase 2 negotiation. Make sure that the IP connectivity between the loopback0 of R2 and the loopback0 of R9 is still up and running.

Task 25.9 Configure IPsec encryption on the ethernet connection between R5 and R8. Use an encryption of AES, a DH group number 2 ans pre-shared key of " iPexpert" during the phase 1 negotiation.

Task 25.10 Between R5 and R8, use esp-3des encryption and an esp-sha-hmac authentication during the phase 2 negotiation.

Task 25.11 Create a VTI on both ends. IP address on R5 is 192.168.58.5/24 and IP address on R8 is 192.168.58.8/24.

Task 25.12 Traffic going from loopback0 of R5 to loopback0 from R8 should be encrypted in both directions. You are not allowed to use a dynamic routing protocol or a default route.

You have completed Lab 25

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.

Lab 26: Configure and troubleshoot IPsec Virtual Private Networks (Part 2)

Technologies covered

- DMVPN phase 1 EIGRP
- DMVPN phase 1 OSPF
- DMVPN phase 2 EIGRP
- DMVPN phase 2 OSPF
- DMVPN phase 1 with IPSec
- DMVPN phase 2 with IPSec

Overview

You have been tasked to configure an IPsec encryption on different connections of your network.

The topology used in the lab will be the following:

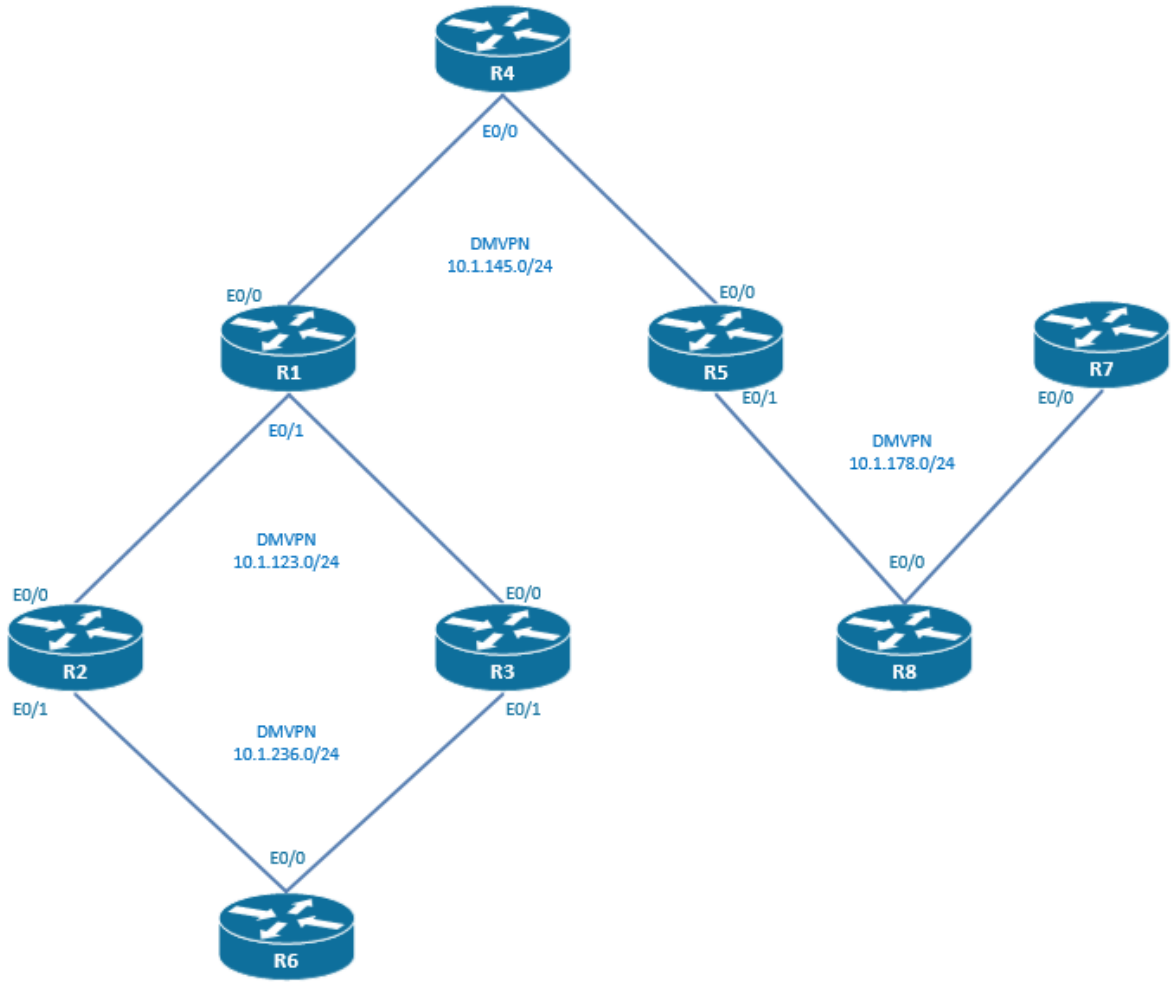

Estimated time to complete: 3-4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 26.1 Configure EIGRP AS 1 on the network between R2, R3, and R6. EIGRP should enable the IP connectivity between the loopback0 of R2, R3, and R6.

Task 26.2 Configure DMVPN phase 1 between R2, R3, and R6. The tunnels number 11 is sourced from the loopback0. The Hub has to act as a NHS. The network-ID of the NHRP network is 11. Use a tunnel key of 11. Use the following IP addresses:

R2	11.0.0.2/24	Spoke
R3	11.0.0.3/24	Spoke
R6	11.0.0.6/24	Hub

Task 26.3 A new registration request should be sent every 10 seconds. A registration request sent by the spokes to the NHS should be kept for 60 seconds if no new update for this entry is received.

Task 26.4 Configure the following loopbacks:

R2	Loopback11	10.11.2.2/32
R3	Loopback11	10.11.3.3/32
R6	Loopback11	10.11.6.6/32

Task 26.5 Configure EIGRP AS 11 on the DMVPN tunnels, configure the spokes as EIGRP stub and advertise the loopback 11 of each router with a network statement. Make sure that there is IP reachability between the loopback11 of R2, R3, and R6.

Task 26.6 Secure the traffic with IPSec on the DMVPN tunnels. Use a hash of MD5, a DH group number 2 and a wild-card pre-shared key of “iPexpert” during the phase 1 negotiation. Use esp-des encryption and an esp-md5-hmac authentication during the phase 2 negotiation.

Task 26.7 Configure OSPF process 2 area 0 on the network between R1, R2, and R3. OSPF should enable the IP connectivity between the loopback0 of R1, R2, and R3.

Task 26.8 Configure DMVPN phase 1 between R1, R2, and R3. The tunnels number 22 are sourced from the loopback0. The network-ID of the NHRP network is 22. Use dynamic mapping. Use a tunnel key of 22. Use the following IP addresses:

R1	22.0.0.1/24	Hub
R2	22.0.0.2/24	Spoke
R3	22.0.0.3/24	Spoke

Task 26.9 Authenticate the NHRP network with an ID of 22 with the key “iPexpert”.

Task 26.10 Configure the following loopbacks:

R1	Loopback22	10.22.1.1/32
R2	Loopback22	10.22.2.2/32
R3	Loopback22	10.22.3.3/32

Task 26.11 Configure OSPF process 22 area 0 on the DMVPN tunnels and advertise the loopback 22 of each router with a network statement. There should not be any DR elected. Make sure that there is IP reachability between the loopback22 of R2, R3, and R6.

Task 26.12 Secure the traffic with IPSec on the DMVPN tunnels. Use an encryption of AES and a wild-card pre-shared key of "iPexpert" during the phase 1 negotiation. Use esp-aes encryption and an esp-sha-hmac authentication during the phase 2 negotiation.

Task 26.13 On the LAN between R1, R4, and R5, setup EIGRP routing in named configuration mode using AS3 and the name of iPexpert. EIGRP should enable the IP connectivity between the loopback0 of R1, R4, and R5.

Task 26.14 Configure DMVPN phase 2 between R1, R4, and R5. The tunnels numbers 33 are sourced from the loopback0. The network-ID of the NHRP network is 33. Do not use dynamic mapping. Use a tunnel key of 33. Use the following IP addresses:

R1	33.0.0.1/24	Spoke
R4	33.0.0.4/24	Hub
R5	33.0.0.5/24	Spoke

Task 26.15 Configure the following loopbacks:

R1	Loopback33	10.33.1.1/32
R4	Loopback33	10.33.4.4/32
R5	Loopback33	10.33.5.5/32

Task 26.16 Configure EIGRP process 33 on the DMVPN tunnels and advertise the loopback 33 of each router with a network statement. Make sure that a ping from the loopback 33 of R1 to the loopback 33 of R5 is always going through the hub.

Task 26.17 Secure the traffic with IPSec on the DMVPN tunnels. Use an encryption of 3-DES and a wild-card pre-shared key of "iPexpert" during the phase 1 negotiation. Use esp-des encryption and an esp-md5-hmac authentication during the phase 2 negotiation.

Task 26.18 On the LAN between R5, R7, and R8, setup OSPF process 4 area 0. OSPF should enable the IP connectivity between the loopback0 of R5, R7, and R8.

Task 26.19 Configure DMVPN phase 2 between R5, R7, and R8. The tunnels numbers 44 are sourced from the loopback0. The network-ID of the NHRP network is 44. No NHRP configuration should be done on the hub. Use a tunnel key of 44. Use the following IP addresses:

R5	44.0.0.5/24	Spoke
R7	44.0.0.7/24	Spoke
R8	44.0.0.8/24	Hub

Task 26.20 Configure the following loopbacks:

R5	Loopback44	10.44.5.5/32
R7	Loopback44	10.44.7.7/32
R8	Loopback44	10.44.8.8/32

Task 26.21 Configure OSPF process 44 area 0 on the DMVPN tunnels and advertise the loopback 44 of each router with a network statement. The election of a DR should take place in this network. The DR should always be on the hub router. Multicast should be enabled on the DMVPN tunnels. Do not use OSPF type broadcast. Make sure that a ping from the loopback 44 of R7 to the loopback 44 of R5 is going directly from R7 to R5.

Task 26.22 Secure the traffic with IPSec on the DMVPN tunnels. Use an encryption of AES, a DH group number 1 and a wild-card pre-shared key of “iPexpert” during the phase 1 negotiation. Use esp-aes encryption and an esp-sha-hmac authentication during the phase 2 negotiation.

You have completed Lab 26

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.

Lab 27: Configure and troubleshoot Protocol Independent Multicast Operations (Part 1)

Technologies covered

- PIM dense mode
- PIM sparse-dense mode
- PIM sparse mode
- RPF failure
- Accept RP
- Accept Register
- DR election
- NMBA mode

Overview

You have been tasked to configure the multicast routing reachability in your network.

The topology used in the lab will be the following:

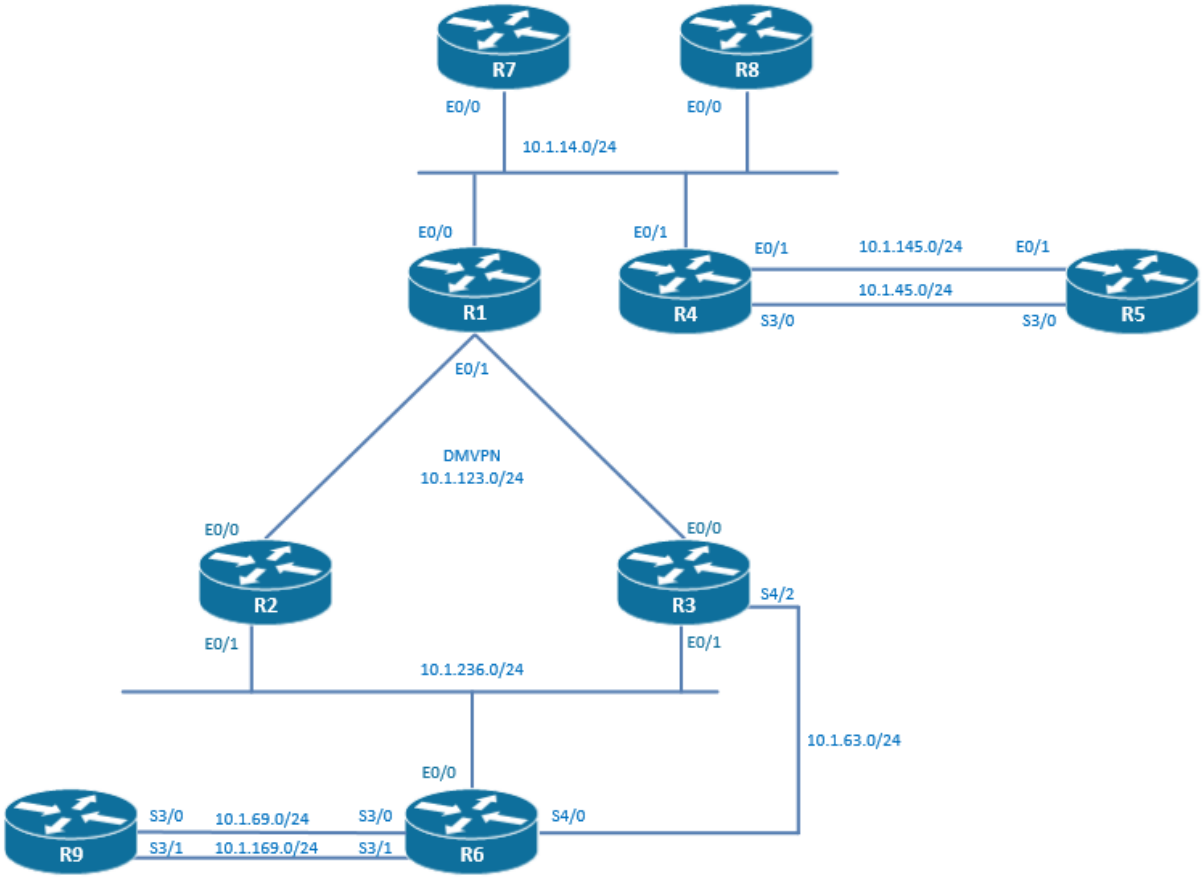

Estimated time to complete: 3 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 27.1 R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN phase 2 without IPSec is the underlying used technology. Setup OSPF in area 0 in this DMVPN network. Configure the OSPF network type as NBMA.

Task 27.2 Advertise the loopbacks of R1, R2, and R3 in the OSPF process. Use network statements. Make sure that you can ping from the loopback0 of R2 to the loopback0 of R3.

Task 27.3 Configure OSPF in area 55 on all the connections between R1, R4, and R5. R1 is the ABR. Cost out the network 10.1.45.0/24 with an OSPF cost of 2000.

Task 27.4 Advertise the loopbacks of R4 and R5 in the OSPF process. Use network statements.

Task 27.5 Configure OSPF in area 99 on all the connections between R2, R3, R6, and R9. R3 is the ABR. Cost out the network 10.1.236.0/24 with an OSPF cost of 2000.

Task 27.6 Advertise the loopbacks of R6 and R9 in the OSPF process. Use network statements.

Task 27.7 There is a multicast server connected on R5 that is sending a stream with the IP address 225.5.5.5. The listeners for this group are located on R1 and R4 only. Configure the network to route this multicast stream from the source to the listeners without the use of any RP. Do not enable multicast on the 10.1.145.0/24 network.

Task 27.8 Configure R1 E0/0 to join 225.5.5.5 and make sure that you can ping this multicast group from R5. If necessary, the use of mroute is allowed.

Task 27.9 There is a multicast server connected on R9 that is sending a stream with the IP address 229.9.9.9. The listeners for this group are located on R5 on network 10.1.45.0/24. Configure the network to route this multicast stream from the source to the listeners with the use of a static RP. Do not enable multicast on the 10.1.163.0/24 network.

Task 27.10 Make sure that R1 is the RP only for the group 229.9.9.9. Use the loopback0 interface for the RP IP address.

Task 27.11 Configure R3 to send the PIM join message to the RP on behalf of the 10.1.236.0/24 network.

Task 27.12 Configure R5 E0/1 to join 229.9.9.9 and make sure that you can ping this multicast group from R9. The use of mroute is allowed.

Task 27.13 There is a multicast server connected on R3 that is sending a stream with the IP address 233.3.3.3. The listeners for this group are located on R2. Shut down the interface e0/1 on R2. Configure the network to route this multicast stream from the source to the listeners with the use of a static RP.

Task 27.14 Make sure that R1 is allowed to be the RP for the group 233.3.3.3. Use the loopback0 interface for the RP IP address.

Task 27.15 Ensure that R2 and R3 send registers (*,G) entries for the group 233.3.3.3 only to the router R1.

Task 27.16 Make sure that you can ping multicast group 233.3.3.3 from R3.

Task 27.17 There is a plan to add a new multicast datastream. The multicast group will be 227.7.7.7 and the source is going to be the server 10.1.63.200. Configure the router R3 so that when he becomes the RP for this multicast group, the only

allowed source is the IP address 10.1.63.200. All other servers trying to register this group should be denied

You have completed Lab 27

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.

Lab 28: Configure and troubleshoot Protocol Independent Multicast Operations (Part 2)

Technologies covered

- Auto-RP
- Auto-RP filtering
- Auto-RP listener
- Multiple RP candidates
- Multicast boundary
- BSR
- BSR Propagation filtering

Overview

You have been tasked to configure the multicast routing reachability in your network.

The topology used in the lab will be the following:

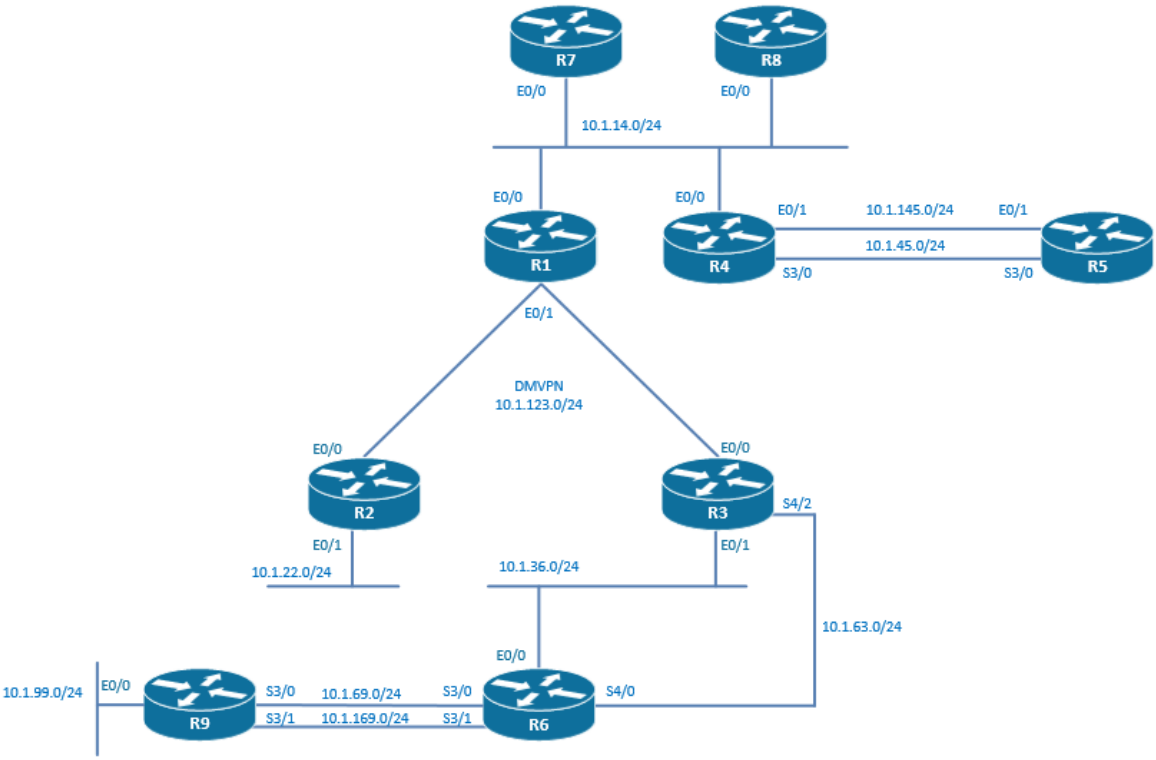

Estimated time to complete: 3 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 28.1 R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN phase 1 without IPsec is the underlying used technology. Setup EIGRP AS 10 in this DMVPN network.

Task 28.2 Advertise the loopbacks of R1, R2, and R3 in the EIGRP process. Use network statements. Make sure that you can ping from the loopback0 of R2 to loopback0 of R3.

Task 28.3 Extend the EIGRP routing domain to include the network 10.1.14.0/24, the network 10.1.145.0/24, and the network 10.1.45.0/24. Advertise the loopbacks of R7, R8, R4, and R5 in the EIGRP process using network statements.

Task 28.4 Extend the EIGRP routing domain to include the networks 10.1.36.0/24, the network 10.1.22.0/24, the network 10.1.63.0/24, the network 10.1.169.0/24, and the network 10.1.69.0/24. Advertise the loopbacks of R6 and R9 in the EIGRP process using network statements.

Task 28.5 Configure PIM on the networks 10.1.14.0/24, the network 10.1.145.0/24, and the network 10.1.45.0/24. Auto-RP will be used on these networks. You are not allowed to use "ip pim auto-rp listener" command.

Task 28.6 Enable R1, R7, and R8 as auto-RP candidates for the following multicast groups: 228.1.1.228, 228.2.2.228, and 228.3.3.228. Their loopback0 should be used in the advertisements.

Task 28.7 Auto-RP advertisements should be sent every 5 seconds to R1, R7, and R8.

Task 28.8 R4 should be configured as the mapping agent. The loopback0 has to be used in the advertisements.

Task 28.9 Configure E0/1 on R5 to join the group 228.1.1.228 and check that you can ping this multicast group from R7, and that R8 has been chosen to be the PIM DR.

Task 28.10 Create a "rp-announce-filter" that makes sure that R7 will never become a RP.

Task 28.11 Create 2 "rp-announce-filters" that make sure that R8 will only become the RP for multicast group 228.1.1.228, and that R1 will only become the RP for multicast groups 228.2.2.228 and 228.3.3.228.

Task 28.12 Configure R1 so that it never does send and receive on interface E0/1 multicast traffic from group 228.1.1.228, 228.2.2.228, and 228.3.3.228. Make sure that the auto-RP advertisements regarding those groups are also filtered.

Task 28.13 Configure E0/1 on R5 to join the group 228.2.2.228, and check that you can ping this multicast group from R7, and that R1 has been chosen to be the RP for 228.1.1.228.

Task 28.14 Ensure that R1, R4, R5, R7, and R8 don't fall back to PIM dense mode for unknown multicast addresses.

Task 28.15 The 2 connections between R9 and R6 have to be configured with PIM sparse-mode (no PIM sparse-dense mode). R9 has to be configured as an auto-RP candidate for all multicast groups, and R6 has to be configured as the mapping agent.

Task 28.16 R9 should not become the RP for routers that are more than 1 hop away.

Task 28.17 Configure the interface S3/0 on R9 to join the group 229.229.229.229, and check that you can ping this multicast group from R6, and that R9 has been chosen to be the RP.

Task 28.18 Enable PIM sparse mode on all interfaces on the network 11.1.1.0/24.

Task 28.19 Configure R2 as the BSR. Use the interface that is always up on a router.

Task 28.20 Configure R1 as the primary RP and configure R3 as a backup RP. One of the two should be configured with the default priority. Use the interfaces that are always up on a router.

Task 28.21 Enable PIM sparse mode on the network 10.1.36.0/24 and 10.1.63.0/24. Ensure that R6 doesn't receive information about RPs elected by PIM bootstrap router process.

Task 28.22 Ensure that R7, R8, and R4 don't receive information about RPs elected by the PIM bootstrap router process.

Task 28.23 Configure E0/1 on R2 to join the group 225.225.225.225 and check that you can ping this multicast group from R1.

You have completed Lab 28

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.

Lab 29: Configure and troubleshoot Protocol Independent Multicast Operations (Part 3)

Technologies covered

- Multicast stub routing
- IP IGMP helper-address
- SSM
- IGMP filtering
- IGMP timers
- Multicast helper map
- PIM bidirectional
- Multicast rate limiting

Overview

You have been tasked to configure the multicast routing reachability in your network.

The topology used in the lab will be the following:

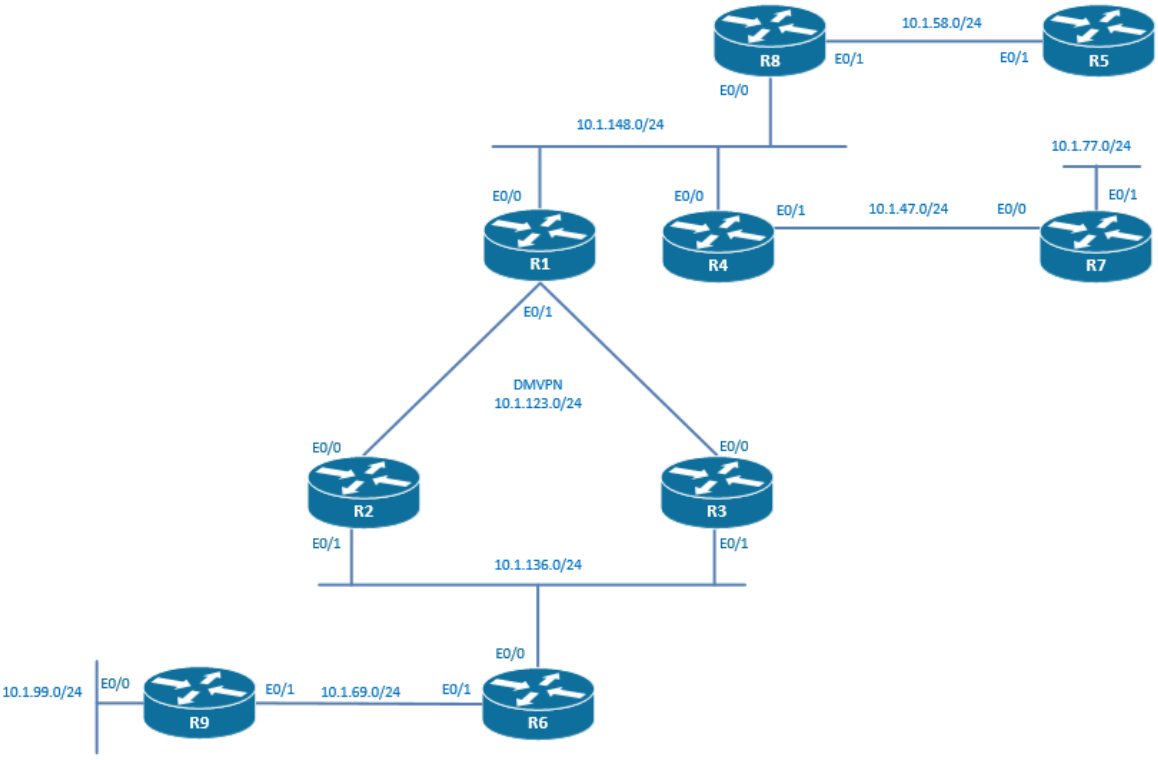

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 29.1 R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. DMVPN phase 1 without IPsec is the underlying used technology. Setup OSPF area 0 in this DMVPN network. Use the point-to-multipoint OSPF on the 2 two spokes.

Task 29.2 Advertise the loopbacks of R1, R2, and R3 in the OSPF process. Use network statements. Make sure that you can ping from the loopback0 of R2 to the loopback0 of R3.

Task 29.3 Introduce R4, R5, R7, R8, R6, and R9 into the OSPF area 0. Advertise the loopbacks of R4, R5, R7, R8, R6, and R9 in the OSPF process. Use network statements. Make sure that you can ping from the loopback0 of R7 to the loopback0 of R9.

Task 29.4 Advertise the networks 10.1.77.0/24 and 10.1.99.0/24 in the OSPF process. Use network statements. Make sure that no OSPF neighborships will never be formed on those networks.

Task 29.5 Configure PIM sparse mode on the networks 10.1.69.0/24, 10.1.236.0/24, 11.1.1.0/24, and 10.1.148.0/24.

Task 29.6 Configure IP PIM dense mode on 10.1.47.0/24. No PIM adjacency should be formed over this connection. Use the command “ip pim neighbor-filter” on R4.

Task 29.7 The source of the multicast stream 224.2.2.2 is located on VLAN 77. The receiver of this multicast stream is on VLAN 10.1.148.0/24. Enable multicast connectivity between this source and this receiver. You are not allowed to remove the filter configured in the previous question, and consequently not allowed to build a PIM adjacency over the connection between R4 and R7. On R1, R4, R7, and R8, configure statically the loopback0 of R1 as the RP for all multicast groups.

Task 29.8 Configure E0/1 on R7 to join the group 224.2.2.2 and check that you can ping this multicast group from R8.

Task 29.9 Make sure that E0/1 on R5 can receive traffic multicast for the group 224.3.3.3 only if it is sourced from loopback0 of R1. Do not enable PIM on this interface.

Task 29.10 Verify that you can ping this multicast group 224.3.3.3 from R1 only when the ping is sourced from the loopback0 of R1.

Task 29.11 R9 has to be protected from an IGMP DOS attack. On the interface E0/0 of R9, allow the maximum number of IGMP states to be 25.

Task 29.12 R6 should only accept on the interface E0/1 multicast clients that want to join a group in the range 225.0.0.0/8.

Task 29.13 Configure interface E0/1 of R9 to join multicast groups 225.2.2.2 and 226.2.2.2. Check on R6 that the filtering configured in the previous question is working.

Task 29.14 On the network 10.1.99.0, there is only one client receiving several multicast streams. As soon as this client is sending an IGMP leave group message, the router should immediately stop forwarding this multicast stream on the LAN and not try to send a group-specific query for this multicast group.

Task 29.15 On VLAN 136, configure IGMP to send membership queries every 30 seconds. The backup querier should become the querier for this LAN if it hasn’t seen a query packet within 1 minute.

Task 29.16 On R9, IGMP protocol should communicate to the multicast clients that they should report their group's membership in a maximum of 30 seconds after receiving a query.

Task 29.17 There is a server that is connected to the network 10.1.136.0/24. This server is sending broadcast UDP traffic to port 2500 to a client connected to the network 10.1.148.0/24. This broadcast traffic should be transported by the multicast group 227.7.7.7 when crossing the connection between R2 and R1, and the connection between R3 and R1.

Task 29.18 The multicast traffic should be converted back to a broadcast when reaching the network 10.1.148.0/24.

Task 29.19 Configure bidirectional PIM for a multicast stream of 224.22.22.22 on the network 11.1.1.0/24 and 10.1.148.0/24. The Loopback0 of the R1 has to be configured as the RP and the mapping agent in this PIM bidirectional setup.

Task 29.20 Configure R6 to limit to total bandwidth for multicast traffic to 20 M on all its interfaces in the egress direction.

Task 29.21 Configure R1 to limit to 5M the bandwidth that the multicast stream with a destination of 224.22.22.22 can use out of the tunnel interface.

You have completed Lab 29

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.

Lab 30: Configure and troubleshoot Protocol Independent Multicast Operations (Part 4)

Technologies covered

- RPF failure
- Multicast BGP extension
- BSR propagation filtering
- MSDP
- Catalyst IGMP snooping

Overview

You have been tasked to configure the multicast routing reachability in your network.

The topology used in the lab will be the following:

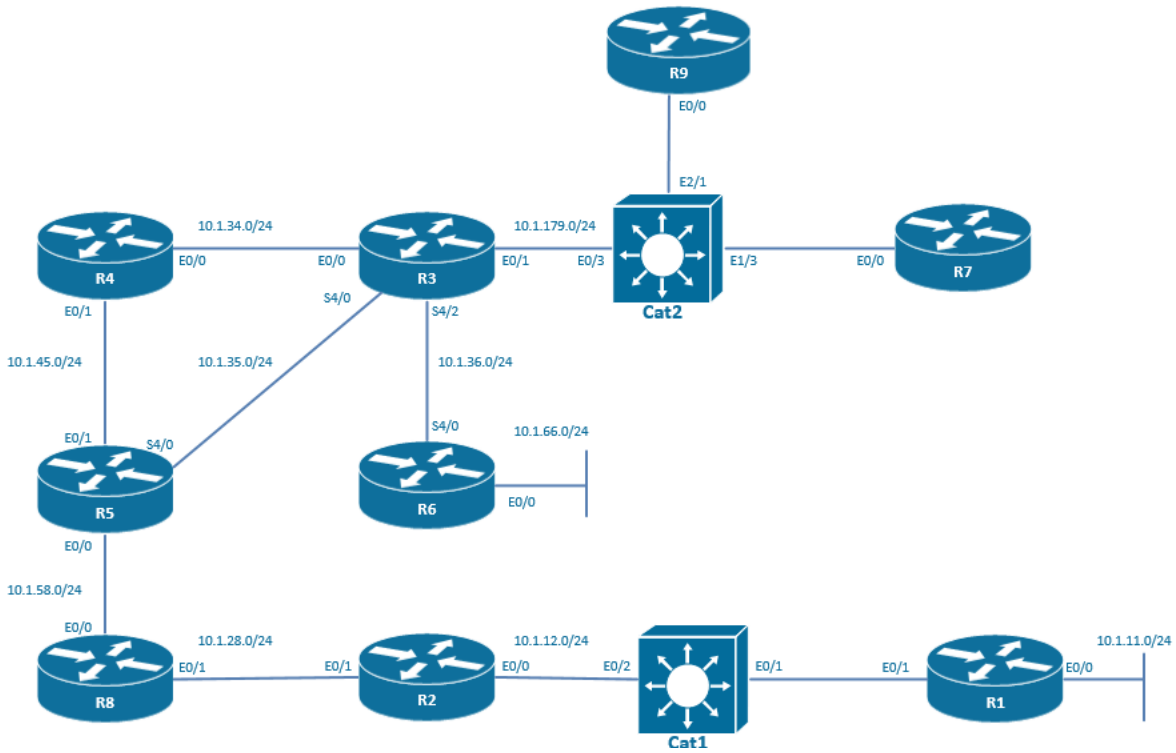

Estimated time to complete: 3 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 30.1 Configure OSPF area 0 routing on the ethernet connections between R5 and R4, R4 and R3, and on the serial connection between R3 and R6. Advertise the loopbacks of R5, R4, R3, and R6 in the OSPF process. Use network statements.

Task 30.2 Configure PIM sparse-mode on the ethernet connections between R5 and R4, R4 and R3, and R3 and R6.

Task 30.3 R3 should be configured as the BSR and the RP for the all multicast groups. Use the PIM bootstrap router solution to advertise the RP. Use the loopback 0 of R3 as the RP IP address.

Task 30.4 On R5, configure on the interface E0/0 an IGMP join for the group 225.7.7.7. Verify that you can ping from R6 to the multicast group 225.7.7.7.

Task 30.5 Configure OSPF area 0 routing on the serial connection between R5 and R3. Do not enable PIM on this link.

Task 30.6 Manipulate this OSPF cost to ensure that the direct link between R5 and R3 is the preferred path for OSPF.

Task 30.7 Verify that you cannot ping from R6 to the multicast group 225.7.7.7 because of a RPF failure. To solve the RPF failure, you are not allowed to configure ip mroutes.

Task 30.8 We are going to use multicast BGP. Remove OSPF from all the routers where it is running and shut down the direct connection between R5 and R3.

Task 30.9 Configure an iBGP peering between R5 and R4 in AS20. Use the Physical IP addresses for the peering's.

Task 30.10 Configure an iBGP peering between R3 and R6 in AS10. Use the Physical IP addresses for the peering's.

Task 30.11 Configure an eBGP peering between R4 and R3. Use the Physical IP addresses for the peering's.

Task 30.12 Configure on each BGP router an "address-family ipv4 multicast". Advertise all the circuits where there is a PIM neighborship into BGP with network statements.

Task 30.13 Advertise The RP IP address into the address-family used for multicast.

Task 30.14 Verify that the feed from R6 to the multicast group 225.7.7.7 is again reaching R5 after the migration from OSPF to BGP.

Task 30.15 On Cat1, configure the E0/1 and the E0/2 interfaces into VLAN 12.

Task 30.16 Configure OSPF area 0 routing on the connection between R5 and R8, on the connection between R8 and R2, and on the connection between R1 and R2.

Task 30.17 Configure PIM in sparse mode on the connection between R5 and R8, on the connection between R8 and R2, and on the connection between R1 and R2.

Task 30.18 R2 should be configured as the BSR and the RP for the all multicast groups. Use the PIM bootstrap router solution to advertise the RP. Use the loopback 0 of R2 as the RP IP address.

Task 30.19 Separate the two BSR domains and make sure that the propagation of the BSR packets is filtered on the connection between R5 and R8.

Task 30.20 On R6, configure on the interface E0/0 an IGMP join for the group 228.7.7.7. Make sure that when you ping from R4 to the group 228.7.7.7, the router R6 is replying.

Task 30.21 On R1, configure on the interface E0/0 an IGMP join for the group 228.7.7.7. Make sure that when you ping from R4 to the group 228.7.7.7, the router R6 and R1 are replying. Use MSDP. Enable OSPF process 2 on the R3, R4, and R5 path.

Task 30.22 As soon as there is one receiver for a multicast group on VLAN 12 connected to Cat1, this multicast group stream should be replicated on all the ports in VLAN 12 even if the servers connected to those ports are not multicast listeners.

Task 30.23 On Cat2, configure the E0/3, E1/3, and the E2/1 interfaces into VLAN 99.

Task 30.24 Configure R3 as the PIM DR for the network 10.1.179.0/24.

Task 30.25 Configure IGMP on Cat2 to prevent R7 to join group 229.7.7.7.

Task 30.26 On R7, configure on the interface E0/0 an IGMP join for the group 229.7.7.7. On R9, configure on the interface E0/0 an IGMP join for the group 229.7.7.7. On Cat2, verify that the IGMP filtering configured in the previous question is working.

You have completed Lab 30

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.

Lab 31: Configure and troubleshoot IP version 6 (Part 1)

Technologies covered

- IPv6 addressing
- DMVPN for IPv6
- RIPng
- RIPng prefix filtering
- RIPng summarization
- RIPng offset-list
- RIPng default route

Overview

You have been tasked to configure the IPv6 routing in your network.

The topology used in the lab will be the following:

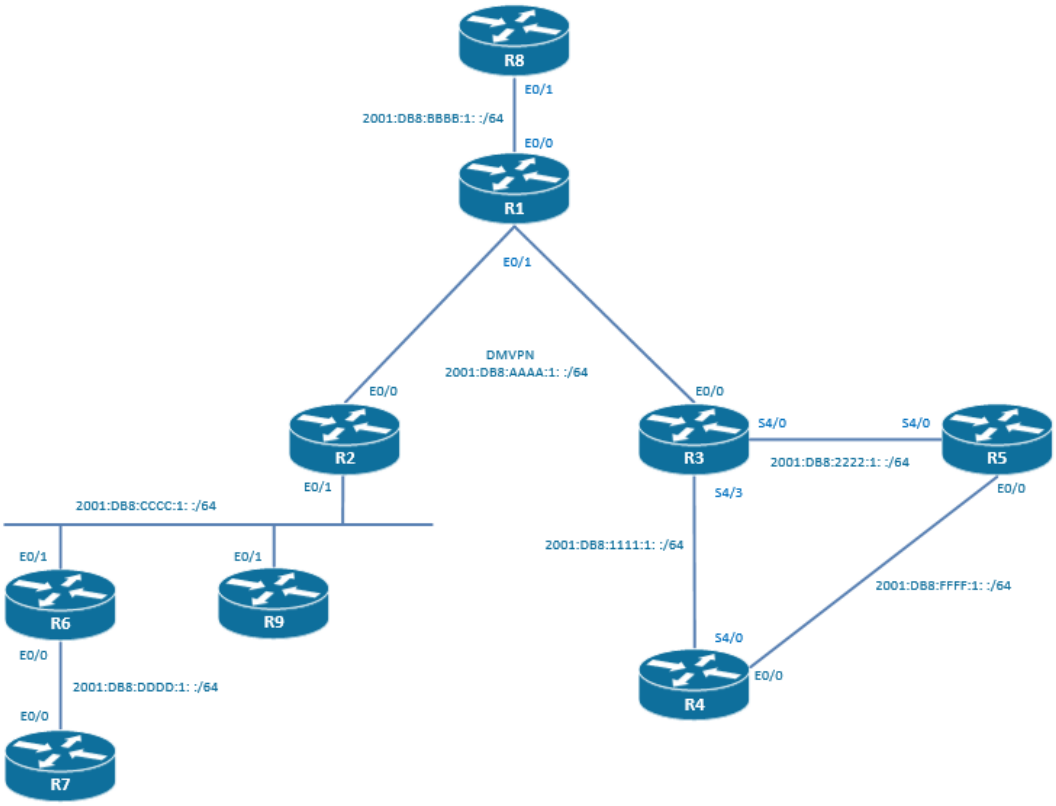

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 31.1 R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. Configure the DMVPN phase 3 tunnel infrastructure for IPv6. Do not implement encryption. Use the following addresses:

R1	E0/1	10.1.123.1/24
R2	E0/0	10.1.123.2/24
R3	E0/0	10.1.123.3/24

	Link Local Unicast	Global Unicast
R1 interface Tunnel23	FE80: :1	2001:DB8:AAAA:1: :1/64
R2 interface Tunnel23	FE80: :2	2001:DB8:AAAA:1: :2/64
R3 interface Tunnel23	FE80: :3	2001:DB8:AAAA:1: :3/64

Task 31.2 Configure the following IPv6 addresses:

	Link Local Unicast	Global Unicast
R1 interface E0/0	EUI-64 format	2001:DB8:BBBB:1: :/64 EUI-64 format
R2 interface E0/1	EUI-64 format	2001:DB8:CCCC:1: :/64 EUI-64 format

Task 31.3 Use RIPng with the identifier of "iPexpert" to enable IP routing between the interface E0/0 of R1 and the interface E0/1 of R2.

Task 31.4 On R1, create an IP host mapping called R2LAN for the IPv6 global address of the E0/1 of R2. Check that you can ping R2LAN from R1.

Task 31.5 On R2, create an IP host mapping called R1LAN for the IPv6 global address of the E0/0 of R1. Check that you can ping R1LAN from R2.

Task 31.6 Configure the following interfaces to automatically assigned IPv6 addresses to their interfaces:

R6	E0/1
R8	E0/1
R9	E0/1

Task 31.7 Enable RIPng with the identifier "iPexpert" on R6, R8, and R9. Check that R8 can reach the IPv6 global address that has been previously assigned to the E0/1 of R6 and to the E0/1 of R9.

Task 31.8 Configure the following IPv6 address on the connection between R6 and R7:

	Link Local Unicast	Global Unicast
R6 interface E0/0	FE80: :1	2001:DB8:DDDD:1: :6/64
R7 interface E0/0	FE80: :2	2001:DB8:DDDD:1: :7/64

Task 31.9 On R7, configure the following IPv6 loopback addresses:

	Global Unicast
R7 interface Loopback4	2001:DB8:EEEE:4: :7/64
R7 interface Loopback5	2001:DB8:EEEE:5: :7/64
R7 interface Loopback6	2001:DB8:EEEE:6: :7/64
R7 interface Loopback7	2001:DB8:EEEE:7: :7/64

Task 31.10 Enable RIPng with the identifier of "iPexpert" on the connection between R6 and R7, and on the 4 loopbacks on R7.

Task 31.11 Ensure that R6 receives from R7 a summary route encompassing all the loopbacks.

Task 31.12 Enable RIPng on the tunnel interface of the routerR3. Ensure that R3 is able to ping the IPv6 address of loopback4 of R7.

Task 31.13 Configure the following IPv6 address on the connection between R3 and R4:

	Link Local Unicast	Global Unicast
R3 interface S4/3	FE80: :1	2001:DB8:1111:1: :3/64
R4 interface S4/0	FE80: :2	2001:DB8:1111:1: :4/64

Task 31.14 Configure the following IPv6 address on the connection between R3 and R5:

	Link Local Unicast	Global Unicast
R3 interface S4/0	FE80: :1	2001:DB8:2222:1: :3/64
R5 interface S4/0	FE80: :2	2001:DB8:2222:1: :5/64

Task 31.15 Configure the following IPv6 address on the connection between R4 and R5:

	Link Local Unicast	Global Unicast
R4 interface E0/0	FE80: :1	2001:DB8:FFFF:1: :4/64
R5 interface E0/0	FE80: :2	2001:DB8:FFFF:1: :5/64

Task 31.16 Enable RIPng with the identifier of 345 on the connections between R3 and R4, on the connection between R3 and R5, and on the connection between R4 and R5.

Task 31.17 Enable full IPv6 connectivity between the 2 RIPng domains, iPexpert and 345.

Task 31.18 Ensure that R4 and R5 have a default route pointing towards R3. You have to configure R3 only to complete this task and you are not allowed to configure static routes.

Task 31.19 The default route and the summarized route for the loopbacks of R7 should be the 2 only RIP process iPexpert entries in the IPv6 routing table of R4 and R5. Configure R3 to achieve this task. Use an IPv6 prefix-list called "SUMMARYR7".

Task 31.20 The clients on the VLAN 2001:DB8:FFFF:1: :/64 should always be routed over the connection R5-R3. The connection R4-R3 should only be used in case the connection R5-R3 is going down.

You have completed Lab 31

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.

Lab 32: Configure and troubleshoot IP version 6 (Part 2)

Technologies covered

- EIGRPv6
- EIGRPv6 summarization
- EIGRPv6 default route
- EIGRPv6 authentication
- EIGRPv6 unequal load balancing

Overview

You have been tasked to configure the IPv6 routing in your network.

The topology used in the lab will be the following:

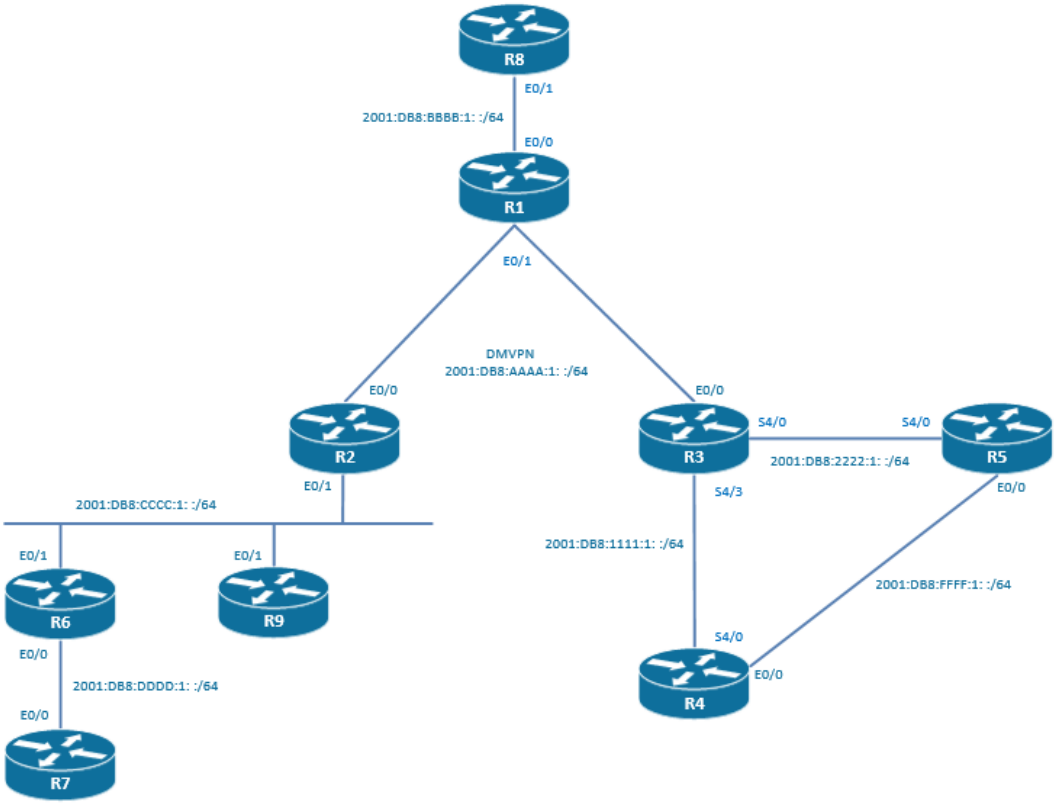

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 32.1 R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. Configure the DMVPN phase 3 tunnel infrastructure for IPv6. Do not implement encryption. Use the following addresses:

R1	E0/1	10.1.123.1/24
R2	E0/0	10.1.123.2/24
R3	E0/0	10.1.123.3/24

	Link Local Unicast	Global Unicast
R1 interface Tunnel23	FE80: :1	2001:DB8:AAAA:1: :1/64
R2 interface Tunnel23	FE80: :2	2001:DB8:AAAA:1: :2/64
R3 interface Tunnel23	FE80: :3	2001:DB8:AAAA:1: :3/64

Task 32.2 Configure an IPv6 NHRP authentication of iPexpert and a NHRP network-id of 123.

Task 32.3 Configure the following loopback IPv6 addresses:

	Global Unicast
R1 interface lo0	2001:DB8:A:A: :1/128
R2 interface lo0	2001:DB8:A:A: :2/128
R3 interface lo0	2001:DB8:A:A: :3/128

Task 32.4 Enable EIGRPv6 with an AS of 123 on the DMVPN network between R1, R2, and R3.

Task 32.5 Make sure that there is IPv6 connectivity between the loopbacks of R1, R2, and R3.

Task 32.6 Configure EIGRPv6 with an AS of 123 on the LAN 2001:DB8:CCCC:1: :/64. Check that you can ping the loopback0 of R3 from R6 and R9.

Task 32.7 In the routing table of R6 and R9, there should be no specific entries for the loopbacks of R1, R2, and R3. There should only be a routing entry to reach the summary route 2001:DB8:A:A: :/126. Check that you can ping the loopback0 of R3 from R6 and R9.

Task 32.8 On R2, create an IPv6 static default route pointing to Null0 and make sure that R2 will be the default router for all packets with an unknown IPv6 addresses in the EIGRP domain AS 123.

Task 32.9 Configure EIGRPv6 with an AS of 123 on the LAN 2001:DB8:BBBB:1: :/64.

Task 32.10 The router R1 should not advertise any specific networks to R8. Only a default route should be advertised. Use the “ipv6 summary-address eigrp” on R1 to resolve this task. Check that you can ping the loopback0 of R3 and the loopback0 of R2 from R8.

Task 32.11 Configure EIGRPv6 authentication between R1 and R8. Use a key chain called “iPexpertchain”, a key number of 2, and a key-string of “iPexpert”.

Task 32.12 Configure the following IPv6 address on the connection between R3 and R4:

	Link Local Unicast	Global Unicast
R3 interface S4/3	FE80: :1	2001:DB8:1111:1: :3/64
R4 interface S4/0	FE80: :2	2001:DB8:1111:1: :4/64

Task 32.13 Configure the following IPv6 address on the connection between R3 and R5:

	Link Local Unicast	Global Unicast
R3 interface S4/0	FE80: :1	2001:DB8:2222:1: :3/64
R5 interface S4/0	FE80: :2	2001:DB8:2222:1: :5/64

Task 32.14 Configure the following IPv6 address on the connection between R4 and R5:

	Link Local Unicast	Global Unicast
R4 interface E0/0	FE80: :1	2001:DB8:FFFF:1: :4/64
R5 interface E0/0	FE80: :2	2001:DB8:FFFF:1: :5/64

Task 32.15 Configure EIGRPv6 with an AS of 345 on the connections between R3 and R4, between R3 and R5, and between R4 and R5.

Task 32.16 Configure the following loopback IPv6 addresses:

	Global Unicast
R4 interface lo0	2001:DB8:A:A: :4/128
R5 interface lo0	2001:DB8:A:A: :5/128

Task 32.17 Make sure that there is IPv6 connectivity between the loopbacks of R2 and R4.

Task 32.18 In the routing table of R3, the routing entry towards the loopback of R5 should contain 2 next-hops, one next-hop being R4 and the other being R5 directly. The cost of the direct path should not be made equal to the cost of the indirect path (via R3). Use the variance command.

You have completed Lab 32

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.

Lab 33: Configure and troubleshoot IP version 6 (Part 3)

Technologies covered

- OSPPFv3
- OSPFv3 traffic engineering
- OSFPv3 virtual link
- OSPFv3 summarization
- IPv6 NAT-PT
- Protocol redistribution

Overview

You have been tasked to configure the IPv6 routing in your network.

The topology used in the lab will be the following:

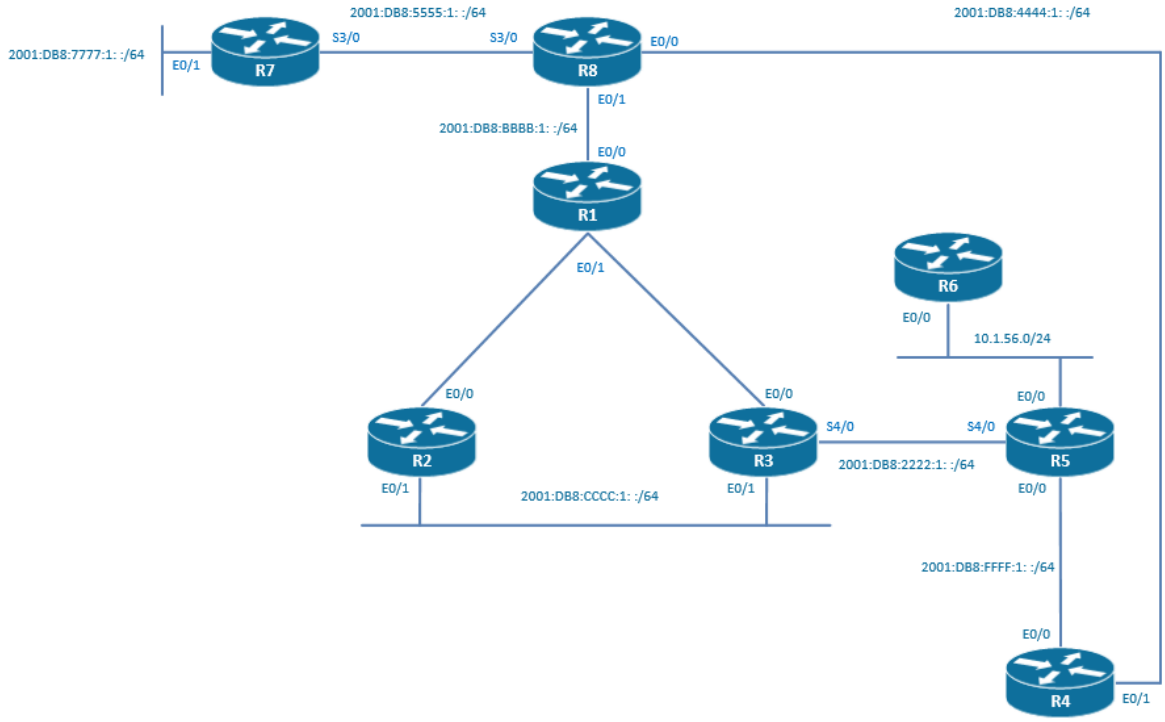

Estimated time to complete: 4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 33.1 R1, R2, and R3 are in a hub and spoke topology where R1 is the hub and R2 and R3 are the spokes. Configure the DMVPN phase 1 tunnel infrastructure for IPv6. Do not implement encryption. Use the following addresses:

R1	E0/1	10.1.123.1/24
R2	E0/0	10.1.123.2/24
R3	E0/0	10.1.123.3/24
R1	lo10	1.1.1.1/32
R2	lo10	2.2.2.2/32
R3	lo10	3.3.3.3/32
R4	lo10	4.4.4.4/32
R5	lo10	5.5.5.5/32

	Link Local Unicast	Global Unicast
R1 interface Tunnel23	FE80: :1	2001:DB8:AAAA:1: :1/64
R2 interface Tunnel23	FE80: :2	2001:DB8:AAAA:1: :2/64
R3 interface Tunnel23	FE80: :3	2001:DB8:AAAA:1: :3/64

Task 33.2 Configure an IPv6 NHRP authentication of "iPexpert" and a NHRP network-id of 123.

Task 33.3 Configure the following loopback IPv6 addresses:

	Global Unicast
R1 interface lo0	2001:DB8:A:A: :1/128
R2 interface lo0	2001:DB8:A:A: :2/128
R3 interface lo0	2001:DB8:A:A: :3/128

Task 33.4 Enable OSPFv3 process 99 in area 0 on the DMVPN network between R1, R2, and R3. DR election should not be taking place. On R1, R2, and R3 use the loopback10 IPv4 address as the OSPF router-ID.

Task 33.5 Make sure that there is IPv6 connectivity between the loopbacks of R1, R2, and R3.

Task 33.6 Configure the following IPv6 address on the connection between R3 and R2:

	Link Local Unicast	Global Unicast
R2 interface E0/1	FE80: :1	2001:DB8:CCCC:1: :2/64
R3 interface E0/1	FE80: :2	2001:DB8:CCCC:1: :3/64

Task 33.7 Enable OSPFv3 process 99 in area 0 on the network 2001:DB8:CCCC:1: :/64.

Task 33.8 R1 should always route via R2 to reach network 2001:DB8:CCCC:1: :/64. Only in case of a failure of the connectivity between R1 and R2, should the path via R3 be chosen. You have to configure R1 to achieve this task.

Task 33.9 Configure the following IPv6 address on the connection between R3 and R5:

	Link Local Unicast	Global Unicast
R3 interface S4/0	FE80: :1	2001:DB8:2222:1: :3/64
R5 interface S4/0	FE80: :2	2001:DB8:2222:1: :5/64

Task 33.10 Configure the following IPv6 address on the connection between R5 and R4:

	Link Local Unicast	Global Unicast
R5 interface E0/1	FE80: :1	2001:DB8:FFFF:1: :5/64
R4 interface E0/0	FE80: :2	2001:DB8:FFFF:1: :4/64

Task 33.11 Configure the following loopback IPv6 addresses:

	Global Unicast
R5 interface lo0	2001:DB8:A:A: :5/128
R4 interface lo0	2001:DB8:A:A: :4/128

Task 33.12 Enable OSPFv3 process 99 in area 55 on the network 2001:DB8:2222:1: :/64.

Task 33.13 Enable OSPFv3 process 99 in area 44 on the network 2001:DB8:FFFF:1: :/64.

Task 33.14 Make sure that there is IPv6 connectivity between the loopbacks of R1, R2, R3, R4, and R5.

Task 33.15 Configure the following IPv6 address on the connection between R1 and R8:

	Link Local Unicast	Global Unicast
R1 interface E0/0	FE80: :1	2001:DB8:BBBB:1: :1/64
R8 interface E0/1	FE80: :2	2001:DB8:BBBB:1: :8/64

Task 33.16 Enable OSPFv3 area 88 on the connection between R1 and R8.

Task 33.17 On R8, configure the following loopback IPv6 addresses:

	Global Unicast
R8 interface lo8	2001:DB8:F:F:8000: :8 /80
R8 interface lo9	2001:DB8:F:F:9000: :8/80
R8 interface lo10	2001:DB8:F:F:A000: :8/80
R8 interface lo11	2001:DB8:F:F:B000: :8/80

Task 33.18 On R8, enable OSPFv3 on loopback8, loopback9, loopback10, and loopback11, and on R1 advertise a single summary network encompassing all the 4 loopbacks.

Task 33.19 Configure the following IPv6 addresses:

	Link Local Unicast	Global Unicast
R8 interface E0/0	FE80: :1	2001:DB8:4444:1: :8/64
R8 interface S3/0	FE80: :1	2001:DB8:5555:1: :8/64
R7 interface S3/0	FE80: :2	2001:DB8:5555:1: :7/64
R7 interface E0/1	FE80: :1	2001:DB8:7777:1: :7/64
R4 interface E0/1	FE80: :2	2001:DB8:4444:1: :4/64

Task 33.20 On R4 and on R8, configure RIPng with an ID of 48 on the connection between R4 and R8.

Task 33.21 On R8 and on R7, configure EIGRPv6 in AS 78 on the connection between R8 and R7.

Task 33.22 EIGRPv6 in AS 78 should also running on the interface E0/1 of R7.

Task 33.23 Ensure IPv6 connectivity between the RIPng routing domain, the OSPFv3 routing domain, and the EIGRPv6 routing domain. In particular, you should be able to IPv6 ping lo0 of R2 from the router R7, you should be able to ping the IP address 2001:DB8:4444:1: :8/64 from the router R3, and you should be able to ping the IP address 2001:DB8:4444:1: :8/64 from the router R7.

Task 33.24 The IPv4 protocol is running on the LAN between R5 and R6. Configure the following IP addresses:

R5 E0/0	10.1.56.5/24
R6 E0/0	10.1.56.6/24

Task 33.25 R3 should be able to ping 10.1.56.6 by using the IPv6 address 2001:DB8:6666:1: :6. You are allowed to configure a static route on R3. The rest of the configuration should be performed on R5.

Task 33.26 Make sure that you can ping IPv6 2001:DB8:6666:1: :6 from all the loopbacks 0 in the routing domain.

You have completed Lab 33

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.

Lab 34: Configure and Troubleshoot Quality of Service Mechanisms (Part 2)

Technologies covered

- Classification and marking
- Bandwidth percent
- LLQ
- WRED
- Dynamic flows
- ECNs

Overview

Voice over IP will be deployed in your network and you have been tasked to configure QOS in your network.

The topology used in the lab will be the following:

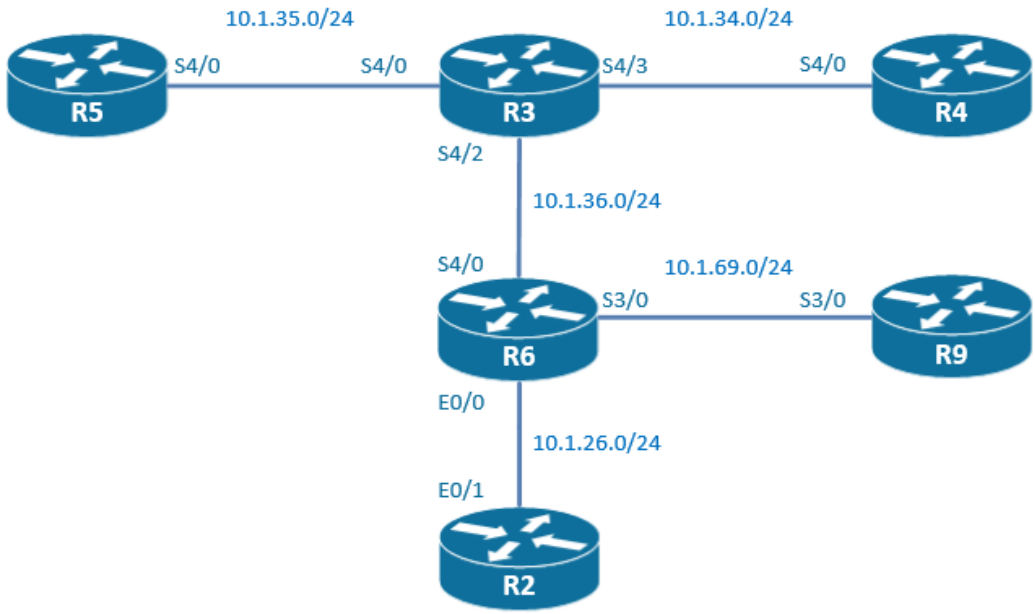

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 34.1 R2 is a customer managed CE and R6 is the entry point to the service provider. The traffic received on the E0/0 is untrusted and should be re-marked when entering the service provider network. A class called VOICE should be created for traffic with destination ports in the RTP range 32512 32768, a class called SQL should be created for traffic with destination ports in the TCP range 1433 1434 and a class called OFFICE_BOSS should be created for traffic originated from the LAN 10.1.222.0/24.

Task 34.2 On R6, configure a policy-map called TRAFFIC_COLOURING. This policy-map should mark the VOICE traffic with the DSCP EF, the SQL traffic with the DSCP AF31, and the OFFICE_BOSS with the DSCP AF21. The remaining unclassified traffic should have the DSCP field reset to 0.

Task 34.3 On the WAN link between R3 and R6, a QOS policy will be enforced. The Voice traffic should be prioritized before any other traffic in case of congestion. 10% of the bandwidth is allocated to VOICE traffic.

Task 34.4 In case of congestion, the SQL traffic should have 30% of the bandwidth reserved and the OFFICE_BOSS traffic should have 20% of the bandwidth reserved.

Task 34.5 In order to slow-down TCP traffic in case of congestion, some packets in the default queue should be randomly dropped before the queue is getting full and tail-dropping.

Task 34.6 On the interface S3/0 of R6, enable WRED to begin to randomly drop packets with the IP precedence of 3 when the queue contains 20 packets and to tail-drop when the number of packets in the queue reaches more than 30 packets. 1 out of 5 packets should be randomly dropped.

Task 34.7 On the interface S3/0 of R6, configure the minimum possible queue size.

Task 34.8 On the interface S4/0 of R4, configure a hold queue of 200 packets.

Task 34.9 On the interface S4/0 of R3, ensure that packets with a DSCP of AF21 begin to be randomly dropped when the queue contains 100 packets and to tail-drop when the number of packets in the queue reach more than 200 packets. 1 out of 10 packets should be randomly dropped.

Task 34.10 The TCP hosts that are transiting on the connection between R3 and R4 are supporting ECN. Enable WRED to take into account the DSCP field. Instead of randomly beginning to drop packets, WRED should be configured to mark the packet that was supposed to be dropped. The goal of this marking is to trigger the receiver to suggest the source to decrease the TCP window size.

~~**Task 34.11** On the connection between R6 and R9, configure WRED to create a queue for each flow. 128 should be allowed to be created and WRED should begin to random-drop in a flow when the queue of this flow contains more than 8 packets.~~

You have completed Lab 34

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.

Lab 35: Configure and Troubleshoot Quality of Service Mechanisms (Part 3)

Technologies covered

- Traffic shaping
- Policing
- Hierarchical policers
- Percent-based policers
- Header compression
- NBAR

Overview

You have been tasked to configure QOS in your network.

The topology used in the lab will be the following:

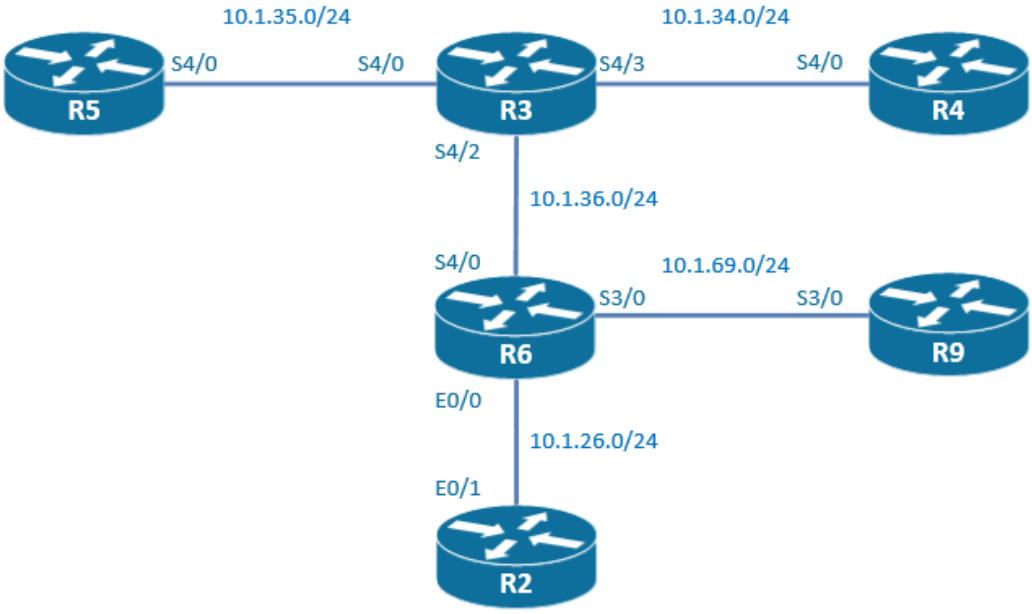

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 35.1 On the WAN link between R3 and R6, enforce a QOS policy using a policy-map called Serial_Policy1. This QOS policy has 3 classes of service. Under congestion, a class called BRONZE matching DSCP AF21 has 256 Kbits/s reserved, a class called SILVER matching DSCP AF31 has 256 kbits/s reserved, and a class called GOLD matching DSCP EF has 512 kbits/s reserved.

Task 35.2 Class SILVER has to be shaped to 512 kbits/s with a normal burst size of 2048 bits.

Task 35.3 Class BRONZE can obtain throughput up to a peak of 512 kbps if enough bandwidth is available.

Task 35.4 On the interface S3/0 of R6, configure traffic-shaping. Limit the egress traffic to 512 kbps. When a BECN is received on this interface, the traffic should be shaped to a minimum of 32 kbps. Make sure that R9 reflects back to R6 the FECNs that he received.

Task 35.5 On the interface E0/1.101 of R2, configure traffic-shaping. Limit the egress TCP traffic for destination port 80 to 1 kbps and the egress TCP traffic for destination port 443 to 300 kbps. Traffic not matching any access-list should be shaped to 100 kbps.

Task 35.6 On R3 and R6, in the policy-map called Serial_policy1, add the following classes: the class called CUSTOMER1 is matching IP DSCP CS4 and the class called CUSTOMER2 is matching IP traffic with a destination TCP port of 69.

Task 35.7 On R3 and R6, in the class called CUSTOMER1, police the traffic to a CIR of 128 kbps with a Bc of 1500 bytes and a PIR of 256 kbps with a Be of 4500 bytes. Packets that conform are sent, packets that exceed are re-marked with a COS of 0 and transmitted, and packets that violate are dropped.

Task 35.8 On R3 and R6, in the class called CUSTOMER2, police the traffic to a CIR of 64 kbps with a Bc of 1500 bytes and a PIR of 128 kbps with a Be of 3000 bytes. Packets marked with a DSCP of AF32 and AF33 that conform are sent, packets with a DSCP of AF32 and AF33 that exceed are re-marked with DSCP of AF11 and transmitted, and packets that violate are dropped. Packets that belong to neither AF32 nor AF33 are re-marked with a DSCP of AF12. Create a class-map called AF3233.

Task 35.9 On the WAN link between R3 and R4, enforce a QOS policy using a policy-map called Serial_Policy_Parent. This QOS policy has only the class default. This policy-map is used to police the traffic to 100 kbps.

Task 35.10 Create a policy-map called Serial_Policy_Child and enforce this QOS policy on the traffic that has already been policed in the previous question. The service-policy Serial_Policy_Child has two classes called CLASS1 and CLASS2. CLASS1 is matching UDP traffic and CLASS2 is matching TCP traffic. CLASS1 should be policed to 20 kbps and CLASS 2 should be policed to 50 kbps.

Task 35.11 On the WAN link between R3 and R5, enforce a QOS policy using a policy-map called Serial_Policy_Percentage. This QOS policy has only the class default. This policy-map is used to police the traffic to a CIR of 60% of the available bandwidth and to a PIR of 90% of the available bandwidth.

Task 35.12 On the WAN link between R3 and R5, configure PPP encapsulation and enable RTP enhanced header compression.

Task 35.13 Consider that the connection between R3 and R4 is a satellite link. Enable RTP header compression on this connection.

Task 35.14 On the link between R6 and R2, enforce a QOS policy using a policy-map called Serial_Policy_NBAR on R6. This QOS policy has 2 classes called LOTUS and URL. LOTUS class is matching Lotus notes traffic and is shaped to 512 kbps. URL class is matching HTTP traffic that contains a URL of /iPexpert is policed to 512 kbps.

You have completed Lab 35

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.

Lab 36: Security Part I

Overview

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 concept is very important when you take the CCIE lab administered by Cisco.

It is recommended to create your own diagram at the beginning of each lab so any potential information you find useful during your preparations can be reflected on this drawing, making it much easier when you step into the real lab.

Multiple topology drawings are available for this chapter.

General Rules

You will need to pre-configure the network with the base configuration files.

NOTE: Static/default routes are NOT allowed unless otherwise stated in the task.

NOTE: You can use "cisco" for any password if other password was not explicitly mentioned in the question.

Estimated Time to Complete: 3-4 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Device	Port	VLAN	IP Address
R1	G0/0 Loop0	41	172.41.41.1/24 2172:41:41::1/64 1.1.1.1/24 1::1/64
R3	F0/1 Loop0	115	10.0.115.3/24 2010:0:115::3/64 3.3.3.3/24 3::3/64
R2	F0/0 F0/1 Loop0	70 117	10.70.70.2/24 2010:70:70::2/64 10.0.117.2/24 2010:0:117::2/64 2.2.2.2/24 2::2/64

R9	F0/1 F0/0 Loop0	41 101	172.41.41.9/24 2172:41:41::9/64 10.10.11.9/24 2010:10:11::9/64 9.9.9.9/24 9::9/64
R8	G0/1 G0/0.115 G0/0.117 Loop0	101 115 117	10.10.11.8/24 2010:10:11::8/64 10.0.115.8/24 2010:0:115::8/64 10.0.117.8/24 2010:0:117::8/64 8.8.8.8/24 8::8/64
CAT4	SVI70	70	10.70.70.140/24

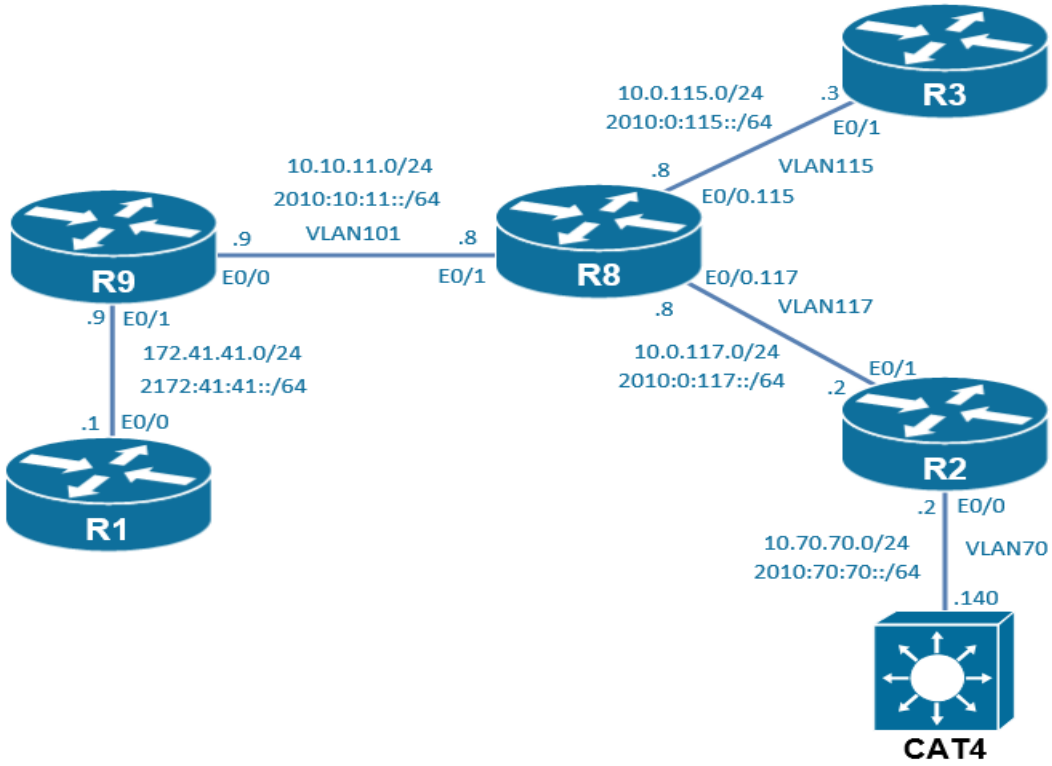

Task 36.1 AAA

- Configure R1 for AAA.
- Users who telnet to this device should be authenticated by the default method list using a line password ("iPexpert"). Console line should not be affected.
- PPP authentication requests should be authenticated using RADIUS server (10.10.11.90).
- Protect RADIUS communication using key "iPexpert". RADIUS traffic should be sent using new port numbers.

- Network access should be authorized – if RADIUS is down authorization should succeed for authenticated users.
- Enable accounting for network traffic – records should be kept for when a session initiates and when it terminates.

Task 36.2 Local Authentication & Authorization

- Enable SSH on R3. Use domain-name “ipexpert.com”.
- Create two local user accounts – “admin” and “secops”.
- When “admin” connects to R3 remotely via SSH it should be automatically placed at level 15 after successful authentication.
- When someone authenticates as “secops” he/she should be placed at level 8.
- Anyone who knows enable password (“cisco”) should be able to access Privilege Level.
- Make sure that enable password is MD5-encrypted.
- Don’t use AAA to accomplish this task.

Task 36.3 AAA EXEC Authorization

- Remove local authentication on R3. Enable AAA.
- Users “admin” and “secops” should be still assigned to privilege levels 15 and 8, respectively, after successful authentication.
- User “secops” should be able to access the following commands:
  - show running-config
  - configure terminal
  - ip routing
  - ip route
- User “admin” should have access to all commands.
- When “secops” issues the “enable” command he should be automatically given Privilege Level access without prompting for password.
- Don’t use any default method lists in this task.

Task 36.4 AAA with CLI Views

- Configure R2 for CLI Views using AAA.
- Create a local user account “administrator” who should be given access to all commands.
- Create a local user account “netops” who should be able to do the following:
  - Access all “show” commands except for any “show crypto” command.
  - Issue “ping” and “telnet”.
  - Configure any dynamic routing protocol.
- Create a local user account “secops” who should be able to do the following:
  - Access all “show crypto” commands.
  - Configure any “crypto” command in the global config mode.
- Create another user account - “ops”. This person should be always able to do what “netops” and “secops” can do.
- Use “iPexpert” as a password for all views.

Task 36.5 Traffic Filtering – Standard ACLs

- R8 is configured with the following loopback networks:

    - 111.111.111.2/32
    - 111.111.111.4/32
    - 111.111.111.6/32
- R1 should be configured to drop & log packets sourced from those addresses using a Standard ACL. This ACL should have as few entries as possible with a minimum overlap.
- All routers should be able to reach R3 only from interfaces configured with odd IPv4 addresses.
- Traffic sourced from other IPv4 addresses should be dropped.
- Implement this using a Standard ACL with a single "deny" entry.

Task 36.6 Traffic Filtering – Extended ACLs

- Configure an IPv4 ACL on R9's F0/0 inbound. Allow the following traffic:
    - OSPFv2 – be very specific here.
    - R1 acts as a Telnet, Web and SQLNET (TCP 1521) server – permit this traffic only to its loopback0 in a single ACL line.
    - UDP-based traceroute (IOS) to any destination – use a single ACL line.
    - All TCP segments destined to R1's Loopback 44 but only with SYN and ACK bits set and FIN bit being not set.
    - All IP packets with any source and destination with a TTL 0-253 and 255 (in a single ACL line).
    - Routers R2, R9, and R8 should be able to ping all interfaces of R1 (regardless of the TTL in the packets). R1 should be able to ping all routers except R3 as well.
- Configure an IPv6 ACL on R9's F0/0 inbound in the following way:
    - Allow Telnet to R1's Loopback 0.
    - Deny all IPv6 packets with a missing or unknown L4 information.
    - Deny all IPv6 packets with Routing Extension Header.
    - Make sure OSPFv3 adjacencies are not affected, same as all ICMPv6 packets.
- Deny and log all other IPv4 & IPv6 traffic. Make sure you see a log message for every packet dropped by this entry.

Task 36.7 Traffic Filtering – Time Ranges & Object-Groups

- All web traffic destined to R8's Loopback 12,14, and 16 interfaces should be denied during business hours Mon-Fri 9am-5pm. This Includes encrypted traffic.
- November 11, 2014 has been declared a no-work day. Ensure that no traffic is allowed to the above mentioned loopbacks for the entire day.
- Permit and log all IPv4 DNS traffic (TCP and UDP) to R8's Loopback0 and 12. Include source MAC address in the logs. Use a single ACL entry to configure this.
- All other traffic should not be affected.

Task 36.8 Traffic Filtering – IP Fragments

- Modify an ACL from the previous task to block all IPv4 fragments regardless of the time/date.
- Block all IPv4 and IPv6 fragments coming to F0/1 on R2 – don't use an access-list to accomplish that.

Task 36.9 Traffic Filtering – Reflexive Access-Lists

- Users at VLAN 70 should be allowed through R2 to any destination when using WWW, Telnet, and SSH.
- Return traffic should be allowed dynamically. Dynamic entries should timeout after a minute.
- Only allow OSPF, ICMP, and Telnet inbound on F0/1.
- Use Reflexive Access-Lists.

Task 36.10 Dynamic (Lock & Key) Access-Lists

- You decided that traffic originating in VLAN 70 should be allowed through R2 only for authenticated users.
- Users will be authenticating using Telnet to 2.2.2.2 over port 3023.
- Sessions should not be idle for more than 2 minutes.
- Sessions longer than 30 minutes require re-authentication.
- A valid local user account for this task is “intuser” with password “cisco”.
- AAA should be already enabled on this device (from one of the previous tasks).

Task 36.11 Policy-Based Routing

- Telnet traffic sourced from R2’s loopback0 destined to 3.3.3.3 should be blackholed on R8.
- Use PBR to accomplish that.

Task 36.12 Unicast Reverse Path Forwarding (URPF)

- Enable Loose Mode uRPF on R8.
- Packets received with unknown sources should be dropped.
- Don’t use a default route when uRPF decisions are made.
- An exception to this policy is packets coming from 192.168.1.0/24 – they should be allowed and logged.

You have completed Lab 36

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.

Lab 37: Security Part II

Overview

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 concept is very important when you take the CCIE lab administered by Cisco.

It is recommended to create your own diagram at the beginning of each lab so any potential information you find useful during your preparations can be reflected on this drawing, making it much easier when you step into the real lab.

Multiple topology drawings are available for this chapter.

General Rules

You will need to pre-configure the network with the base configuration files.

NOTE: Static/default routes are NOT allowed unless otherwise stated in the task.

NOTE: You can use "cisco" for any password if other password was not explicitly mentioned in the question.

Estimated Time to Complete: 2-3 hours

Pre-Lab Setup

Please login to your Security vRack at ProctorLabs.com and load the initial Configuration,
Verify basic L2/L3 connectivity. Use IP Addressing Table, Lab Diagram, and the Physical Topology.

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 configuration tasks as detailed below

Prerequisites

Load the initial configuration files before starting to work on the tasks

Device	Port	VLAN	IP Address
R1	G0/0 Loop0	41	172.41.41.1/24 2172:41:41::1/64 1.1.1.1/24 1::1/64
R3	F0/1 Loop0	115	10.0.115.3/24 2010:0:115::3/64 3.3.3.3/24 3::3/64
R2	F0/0 F0/1 Loop0	70 117	10.70.70.2/24 2010:70:70::2/64 10.0.117.2/24 2010:0:117::2/64 2.2.2.2/24

			2::2/64
R9	F0/1 F0/0 Loop0	41 101	172.41.41.9/24 2172:41:41::9/64 10.10.11.9/24 2010:10:11::9/64 9.9.9.9/24 9::9/64
R8	G0/1 G0/0.115 E0/0.117 Loop0	101 115 117	10.10.11.8/24 2010:10:8::8/64 10.0.115.8/24 2010:0:115::8/64 10.0.117.8/24 2010:0:117::8/64 8.8.8.8/24 8::8/64
CAT4	SVI70	70	10.70.70.140/24

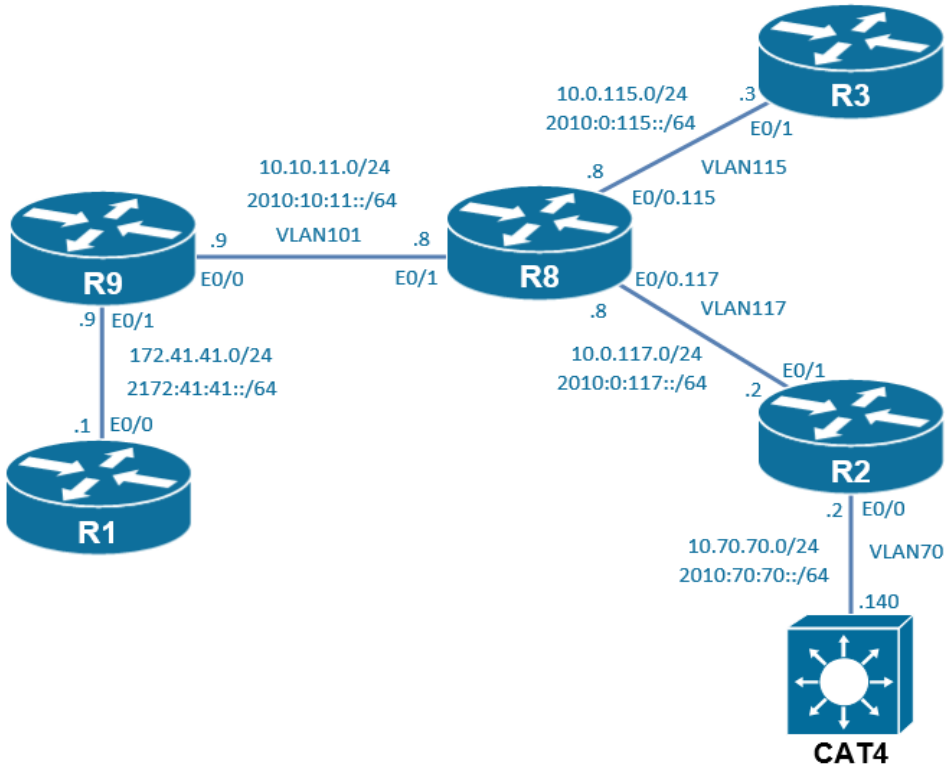

Task 37.1 NBAR

- Using NBAR create and apply a policy outbound on R2's F0/1 to drop the Slammer worm traffic.
- The Slammer worm propagates over UDP port 1434 and its packets are exactly 404B long.
- In the same policy all HTTP packets with string "attack" in the URL should be dropped but only when traffic is going to a WWW server 8.8.8.8 (R8).
- The string should be case insensitive.

Task 37.2 NBAR Next-Gen (NBAR2)

- Configure R8 to drop all terminal-related traffic except PCANYWHERE.
- Also implement a policy for peer-to-peer traffic :
  - All clear-text packets should be rate-limited to 200kbps.
  - All encrypted traffic should be dropped.
- Enable classification of IPv6 traffic that is carried over Teredo tunnels.
- Use a technology that examines IPv4 and IPv6 packets.
- Apply the policy outbound on G0/1.

Task 37.3 NBAR Protocol Discovery

- Enable NBAR Protocol Discovery on R9's F0/0.
- Make sure statistics are obtained for IPv4 and IPv6 traffic.

Task 37.4 TCP Intercept

- There are multiple servers in VLAN 70 hosting various TCP-based applications.
- Several DoS attacks took place recently targeted at those devices.
- Configure R2 to intercept TCP connection requests to this segment.
- If the total number of half-open connections reaches 400, R2 should start randomly dropping them.
- This should cease if the number of half-open sessions falls below 200.
- Make sure router stops managing the sessions after 40 minutes of inactivity.

Task 37.5 TCP Intercept Passive Mode

- There are some other TCP servers that were recently attacked with large amount of spoofed SYN requests (3.3.3.0/24 segment).
- R3 should be configured to send a reset to the server under attack but it should not participate in the handshake.
- The reset segment should be sent if a session does not establish within 20 seconds.
- If a number of connection attempts within the last minute exceed 100, or when a total number of half-open sessions exceed 300, the sessions should be reset faster - after 10 seconds.
- If a FIN exchange or RST packet was seen for a session it should be dropped after 7 seconds.

Task 37.6 Packet Logging

- Configure R1 to send all logged messages to a Syslog server located at 10.70.70.100.
- Use facility type local1.
- Use detailed time stamps for log and debugs including local time zone, and the time of day.
- Logs should be also sent to a buffer – allocate 16384B of memory for this purpose.
- Log messages should be sent with source of 1.1.1.1 and they should be rate-limited to 200 per second except for Sev 1 messages.

Task 37.7 VLAN Filtering

- Configure a VACL on CAT4 to deny the following traffic within VLAN 117 :
  - TCP packets destined to 2.2.2.2 over port 3023.
  - All DNS traffic.
  - Non-IP frames destined to 00:04:cc:1e:12:34.
- Log dropped IP packets – set the log table size to 300 flows.
- Ensure that a log message is seen for every dropped packet.

Task 37.8 Port Security

- Enable Port Security on CAT2.
- Make sure that port connected to R1 will accept frames with R1's MAC, but don't configure address statically.
- On the same interface also allow frames coming from 0000.2222.3333.
- If a violation occurs frames should be dropped, and a Syslog and SNMP traps should be generated. The switch should try to automatically recover from a violation every 50 seconds.
- Anytime the switch reboots it should not affect the Port Security table.

You have completed Lab 37

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.

Lab 38: Security Part III

Overview

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 concept is very important when you take the CCIE lab administered by Cisco.

It is recommended to create your own diagram at the beginning of each lab so any potential information you find useful during your preparations can be reflected on this drawing, making it much easier when you step into the real lab.

Multiple topology drawings are available for this chapter.

General Rules

You will need to pre-configure the network with the base configuration files.

NOTE: Static/default routes are NOT allowed unless otherwise stated in the task.

NOTE: You can use "cisco" for any password if other password was not explicitly mentioned in the question.

Estimated Time to Complete: 2-3 hours

Pre-Lab Setup

Please login to your Security vRack at ProctorLabs.com and load the initial Configuration.
Verify basic L2/L3 connectivity. Use IP Addressing Table, Lab Diagram, and the Physical Topology.

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 configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Device	Port	VLAN	IP Address
R1	G0/0	101	10.10.11.1/24 2010:10:11::1/64
	Loop0		1.1.1.1/24 1::1/64
R3	F0/1	115	10.0.115.3/24 2010:0:115::3/64
	Loop0		3.3.3.3/24 3::3/64
R2	F0/0	70	10.70.70.2/24 2010:70:70::2/64
	F0/1	117	10.0.117.2/24 2010:0:117::2/64
	Loop0		2.2.2.2/24 2::2/64
R9	F0/0	101	10.10.11.9/24 2010:10:11::9/64
	Loop0		9.9.9.9/24 9::9/64
R8	G0/1	101	10.10.11.8/24 2010:10:11::8/64
	G0/0.115	115	10.0.115.8/24 2010:0:115::8/64
	G0/0.117	117	10.0.117.8/24 2010:0:117::8/64
	Loop0		8.8.8.8/24 8::8/64
CAT4	SVI70	70	10.70.70.140/24

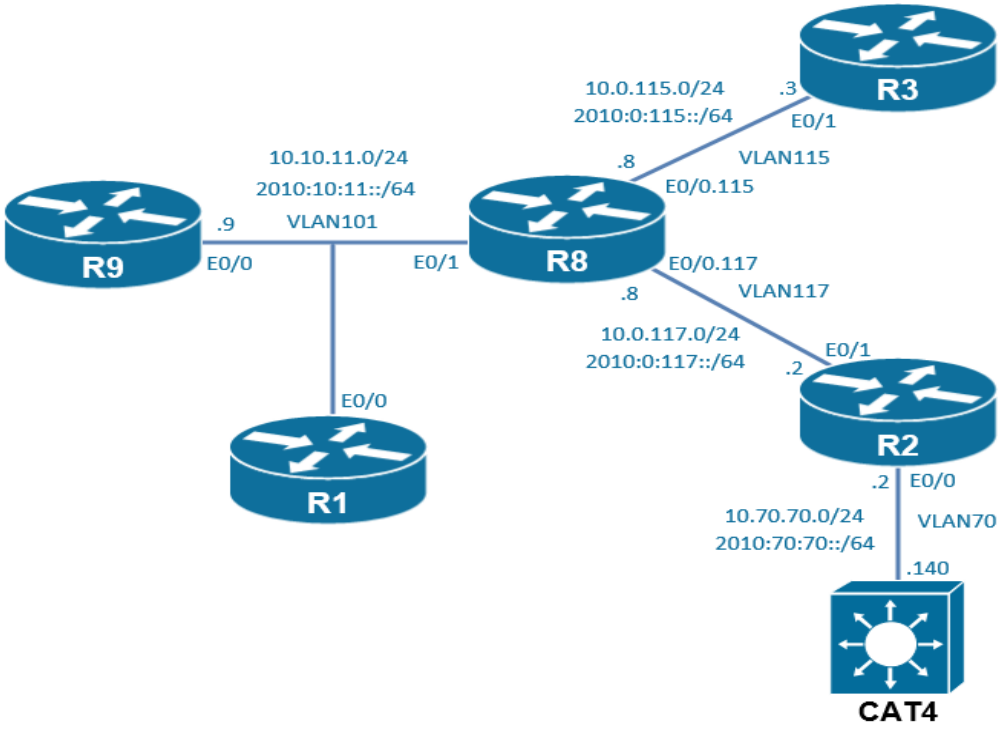

Task 38.1 DHCP Snooping

- Secure DHCP communication in VLAN 101 using DHCP Snooping.
- Configure R9 to act as a DHCP Server in this VLAN.
- Make sure R1 and R8 obtain their address dynamically.
- Rate-limit client DHCP traffic to 15pps.
- Ensure that snooping bindings don't disappear after a reload. The lease times should be accurate - configure & use R9 as a NTP server.

Task 38.2 Dynamic ARP Inspection

- Prevent ARP Man-in-the-Middle attacks in VLAN 101.
- Routers R1, R9, and R8 should be able to successfully communicate.
- ARP packets generated by those devices should be logged.
- Enable source and destination MAC address validation.
- Rate-limit ARP traffic on port connected to R1 to 10 pps. Set the burst interval to 3 seconds.
- Disable ARP Inspection on trunks to other switches.
- You are not allowed to modify the DHCP Snooping database in this task.

Task 38.3 IP Source Guard

- Configure Cat2 to prevent against IPv4 spoofing attacks in VLAN101.
- Not only IP addresses should be validated but also MACs.
- Enable IP Source Guard on F0/1 and F0/8.
- Also configure a static source binding for R9.

Task 38.4 Catalyst Ingress Access-lists

- Configure Port ACLs on CAT4.
- ICMP Echos and Telnet packets received on Fa0/8 should be dropped and logged.

- On the same interface block AppleTalk and ARP frames coming from 0000.cc1e.cc1e.
- Other traffic should not be affected.

Task 38.5 Controlling Terminal Line Access

- Secure VTY lines on R9 and R1.
- Management traffic should be allowed from the following subnets:
  - 10.0.115.0/24
  - 10.0.117.0/24
  - 2010:0:117::/64
- R1 should only accept Telnet.
- R9 should only allow SSH access (user: cisco, pw: cisco).
- You will have to disable IP Source Guard to test this configuration.

Task 38.6 Control Plane Policing

- R8 should be configured to protect its CPU using CoPP.
- Rate-limit all ICMP packets to 15 per second.
- Rate-limit all ICMPv6 packets to 70000bps.
- All HTTP packets originating from 3.3.3.3 should be dropped.
- Outbound telnet packets destined to 1.1.1.1 should be dropped and logged. Log messages should be generated every 2 seconds and they should include TTL and length of dropped packets.
- OSPFv2 and OSPFv3 packets should not be affected by this configuration.

Task 38.7 Control Plane Protection

- Enable Control Plane Protection on R9.
- Packets destined to non-listening ports should be silently dropped.
- Telnet connections over port 3020 should be unaffected.
- Input queue of R9 should not be overwhelmed by any single protocol traffic.
- No more than 100 BGP and 4 SSH packets should be queued.
- No more than 30 packets for all other TCP/UDP protocols enabled on the router should be seen in the queue.
- All IPv4 transit traffic punted to the CPU should be policed to 512kbps.

Task 38.8 Control Plane Protection – Logging

- All malformed & allowed packets received on Host subinterface should be logged.
- Rate-limit those log messages to one every 5 seconds.
- Log all dropped Transit packets that entered R9 through interface F0/0.
- Allowed and over the Input Queue limit SSH traffic should be logged as well.

Task 38.9 Flexible Packet Matching

- Use Flexible Packet Matching to drop & log malicious traffic going through R2.
- The offending packets are sourced in VLAN 101 and they contain string “xExe” within the first 200B from the beginning of TCP Payload.
- Those packets are destined to TCP port 8013.
- Other traffic flowing over the same port number should not be affected.

You have completed Lab 38

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.

Lab 39: Configure and Troubleshoot IP/IOS Services (Part 1)

Technologies covered

- Syslog logging
- Logging timestamps
- Logging to flash
- Configuration change notification
- Configuration archive and rollback
- Conditional debugging

Overview

You have been tasked to configure management services in your network.

The topology used in the lab will be the following:

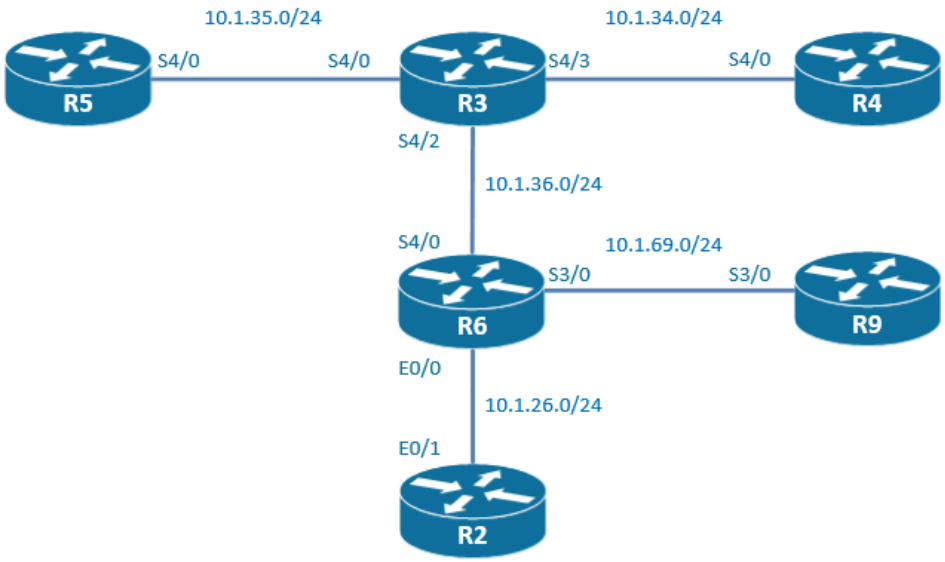

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 39.1 Configure R2 to log system messages to a syslog server with the IP address 10.2.2.2. Send only emergencies, alerts, and critical messages.

Task 39.2 Configure R2 to log all messages with a severity from 1 to 7 in an internal buffer. The size of this buffer should be 20000.

Task 39.3 Make sure that any type of log messages has the exact date and time stamp (and not the uptime).

Task 39.4 If two system messages arrive with the same timestamps, make sure (with sequence numbers) that you still know which one was generated first.

Task 39.5 Configure R2 to log only emergencies, alerts, critical, and error messages to the console.

Task 39.6 Ensure that the router does keep a history file of 10 logged messages prepared to be sent as SNMP traps.

Task 39.7 Limit the rate of logging messages to 70 per second for all logging messages, except for those with a severity level between 5 and 7.

Task 39.8 On R6, write the syslog messages into a file on the flash memory in a directory called "syslog". Once the size of the sum of all the logging files is reaching 64000 bytes, the oldest file is deleted. Each file should have a maximum size of 10000 bytes.

Task 39.9 Log every configuration command entered on R9. Log the last 500 configuration command messages locally. Make sure that the passwords and SNMP community strings are replaced by ****asterisks****. Log, also the configuration command messages on a syslog server.

Task 39.10 On R3, enable the archive feature to store the configuration files on the flash. The maximum number of archive saved should be 10.

Task 39.11 Save the configuration on R3. Change the hostname of R3 to R3-TEST, don't confirm and make sure that the configuration is automatically rolled back to the hostname R3 after 1 minute.

You have completed Lab 39

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.

Lab 40: Configure and Troubleshoot IP/IOS Services (Part 2)

Technologies covered

- SNMP v2
- SNMP v3
- NTP

Overview

You have been tasked to configure management services in your network.

The topology used in the lab will be the following:

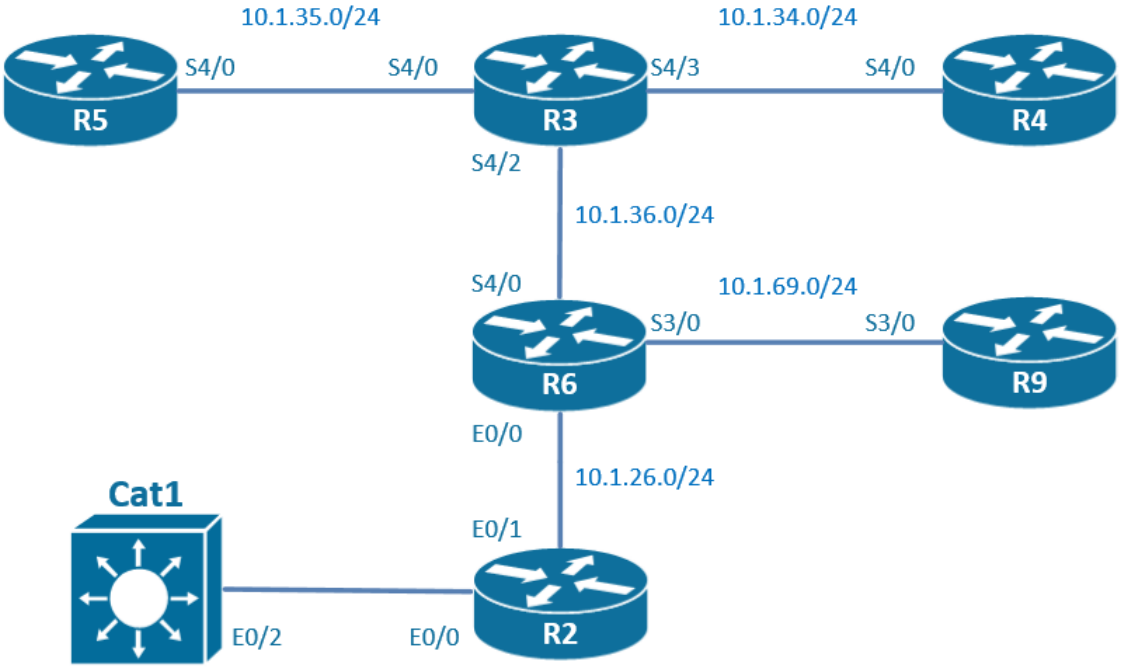

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 40.1 On R2, permit any SNMP server to poll the router with read-only permission using the community string iPexpert.

Task 40.2 R2 should send IPSEC traps to the server 10.4.4.4 using SNMPv2c. The community iPexpert is included in the traps.

Task 40.3 On R6, permit only hosts 10.4.4.4 and 10.4.4.3 to poll the router with read-only permission using the community string iPexpert. Use access-list number 6.

Task 40.4 R2 should send all syslog messages as SNMP ACKed traps to the server 10.4.4.4 using SNMPv2c. ACKed trap means that an ACK packets should be sent by the server back to R2 to confirm that he received the trap. The community iPexpert is included in the traps.

Task 40.5 R3 is going to be polled by a NMS with an IP address of 10.5.5.5. This polling should be configured according to the AuthPriv security model. Create two views, a RO view called ROVIEW and a RW view called RWVIEW. Make the MIB-2 objects accessible for both views.

Task 40.6 On R3, define a RO group called ROGROUP. Associate to this group the following user:

- username: Username1
- password: Password1
- encryption password: iPexpert
- Use the SHA authentication method and the 3-DES encryption method.

Task 40.7 On R3, define a RW group called RWGROUP. Associate to this group the following user:

- username: Username2
- password: Password2
- encryption password: iPexpert
- Use the MD5 authentication method and the AES-256 encryption method.

Task 40.8 On R3, enable traps and informs to be sent to 10.5.5.5 using payload encryption. The user Username1 generates the traps and informs.

Task 40.9 Configure Cat1 to send an SNMP version 2C trap with a community of “iPexpert” to the NMS 10.5.5.5 whenever the switch learns or time-outs a MAC address.

Task 40.10 On Cat1, enable the MAC address notification feature. Store the MAC address notification traps and send them to the NMS every 30 seconds. Keep a historical table of the 10 last MAC address notification messages locally on the switches.

Task 40.11 Configure R5 as a stratum 5 NTP master.

Task 40.12 NTP server on R5 should source from interface S4/0.

Task 40.13 Configure R3 as client from NTP server R5. Configure NTP authentication between R3 and R5 with a key number of 1 and a password of “ iPexpert”.

Task 40.14 On R5, make sure that the only NTP client that can synchronized with R5 is the client with the IP address 10.1.35.3. Use an access-list called NTPCLIENT.

Task 40.15 Make sure that only 10.1.35.5 can be the NTP server for R3. Configure on R3 an access-list called NTPSERVER.

You have completed Lab 40

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.

Lab 41: Configure and Troubleshoot IP/IOS Services (Part 3)

Technologies covered

- EEM
- Proxy ARP
- Local Proxy ARP
- DHCP

Overview

You have been tasked to configure management services in your network.

The topology used in the lab will be the following:

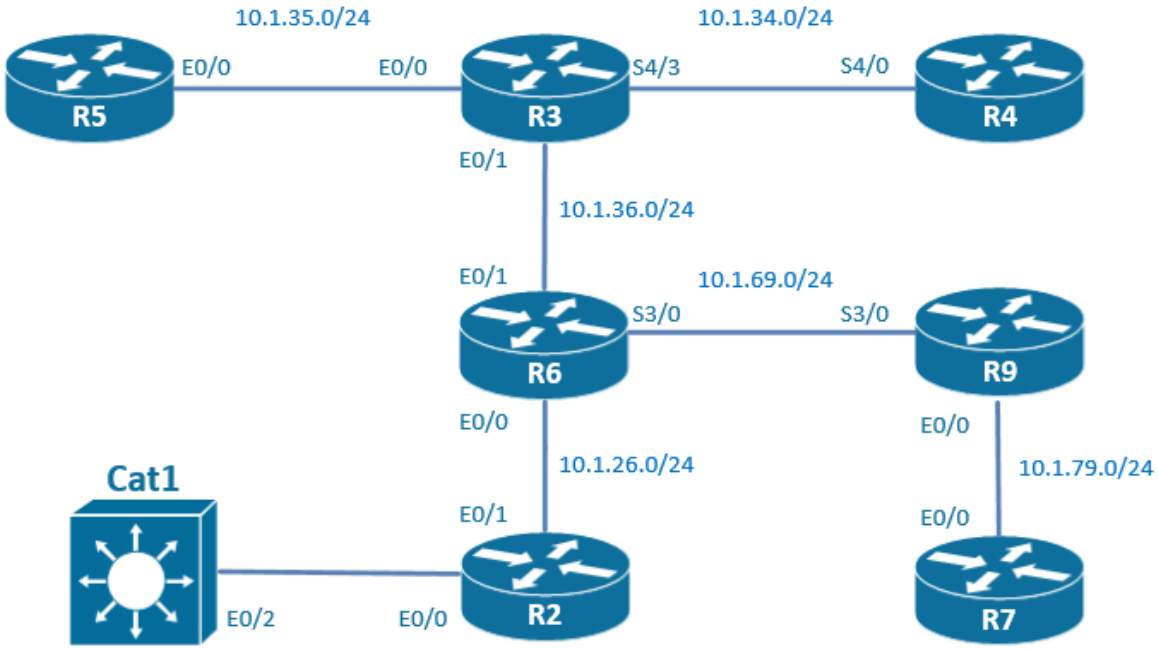

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 41.1 On R2, when the interface E0/1 is going down and the router is generating a syslog message regarding this event, create an EEM applet that will perform a show int E0/1 and send to the email noc@ipexpert.com the output of the

command in the body of the mail. The mail server is 10.3.3.3, the originator of the mail is R2@ipexpert.com, the subject of the mail is ALERT_R2_E0_1_DOWN.

Task 41.2 On R2, when someone is trying to reload the router, the reload command should have no effect. It should trigger an EEM applet to check who is currently logged in and store the output of this command in the system flash in a file called reload_user. The EEM applet should also send the following syslog message: "Someone tried to reload the router R2".

Task 41.3 On R6, when E0/1 is up, S3/0 has to be administratively shut down. When E0/1 is in a down state, S3/0 has to be enabled. Use 2 EEM applets to achieve this.

Task 41.4 On R6, configure an EEM applet that is saving the configuration to NVRAM every hour. Each time the script is run; generate a syslog message stating "Configuration saved by EEM applet."

Task 41.5 Configure the IP address 10.1.36.6 with a mask 255.255.0.0 on the interface E0/1 of R6. Do not modify this mask on the other side of the connection between R6 and R3. In the routing table of R6, there are only the connected networks. However, R6 is able to ping 10.1.35.5 with the ping sourced from IP address 10.1.36.6. On the interfaces of R3, disable the mechanism that makes this IP connectivity possible.

Task 41.6 On R2, make sure that the interface E0/1 is replying to all the ARP requests sent on the network 10.1.26.0/24.

Task 41.7 Configure R3 as a DHCP server for the network 10.1.35.0/24 and 10.1.36.0/24. Default gateways are 10.1.35.1 and 10.1.36.1 respectively. The DNS server IP address is 10.2.2.2.

Task 41.8 The IP address range 10.1.35.1-10.1.35.11 should be excluded from the IP addresses allocated to the clients by the server.

Task 41.9 The IP address range 10.1.36.1-10.1.36.11 should be excluded from the IP addresses allocated to the clients by the server.

Task 41.10 R3 will also be DHCP servers for the network 10.1.26.0/24. Default gateway is 10.1.26.1. The DNS server IP address is 10.2.2.2. Use static routing in order to enable routing between R2 and R3.

Task 41.11 The IP address 10.1.35.100 should always be assigned to the server with the mac address aaaa.bbbb.cccc.

Task 41.12 Configure R9 as a DHCP server for the network 10.1.79.0/24. Default gateway is 10.1.79.1. The DNS server IP address is 10.2.2.2. Exclude 10.1.79.1-11 from the DHCP range.

Task 41.13 The interface EG0/0 of R7 should retrieve an IP address from the DHCP pool configured earlier.

Task 41.14 On R9, configure AAA and Radius for DHCP accounting. The RADIUS server has IP address 10.2.2.2.

You have completed Lab 41

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.

Lab 42: Configure and Troubleshoot IP/IOS Services (Part 4)

Technologies covered

- IP SLA
- HSRP
- VRRP
- GLBP

Overview

You have been tasked to configure management services in your network.

The topology used in the lab will be the following:

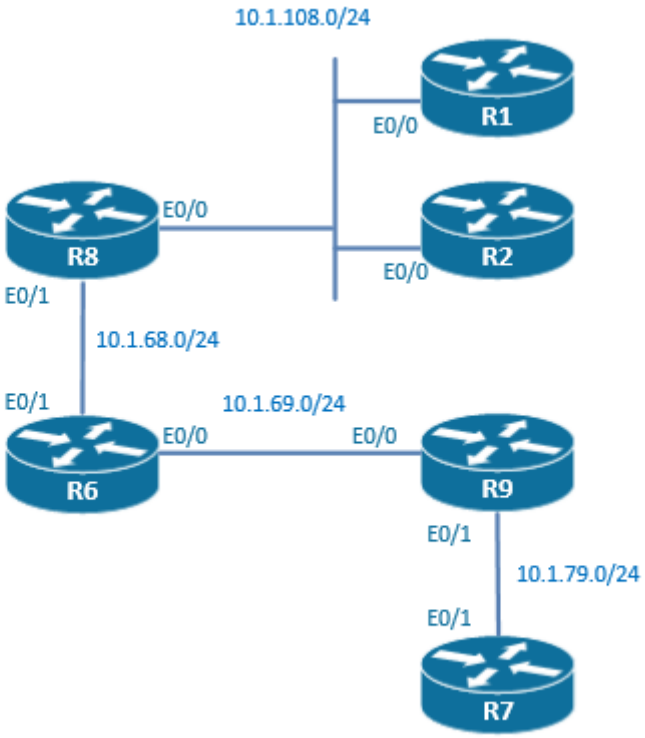

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 42.1 On the connection between R7 and R9, configure IP SLA on R7 to measure the UDP jitter. UDP packets should be sent to 10.1.79.9 port 3200 every 10 seconds with a DSCP marking of EF. This measurement should run indefinitely.

Task 42.2 When the connection between R7 and R9 is lost, R7 will send a trap and trigger a ping 10.1.79.9 every 3 seconds during 60 seconds. If connectivity is not re-established after 60 seconds, a second trap will again be sent. Enable R7 to send CISCO-SYSLOG-MIB traps to the SNMP server 10.1.222.200 with the community "iPexpert".

Task 42.3 Between R6 and R9, configure an IP SLA job on R6 that will generate an ICMP echo with a packet size of 1000 bytes every 10 seconds. Those packets have to be sent to 10.1.69.9.

Task 42.4 The IP SLA control messages between R6 and R9 have to be authenticated using a key-chain called "iPexpert". This key-chain should use key number 3 and a key string of "iPexpert".

Task 42.5 Between R6 and R8, configure on R6 a TCP operation to 10.1.68.8 on port 443 that doesn't require R8 to be configured as a responder.

Task 42.6 Between R8 and R2, configure on R8 a TCP operation to 10.1.108.2 on port 80 that requires R2 to be configured as a responder.

Task 42.7 Configure R8 to perform every 30 seconds a DNS lookup on the DNS server 10.1.222.222 for the website www.ipexpert.com.

Task 42.8 Configure GLBP between R8, R2, and R1 on the network 10.1.108.0/24. Virtual IP address is 10.1.108.133. 10% of the traffic should use R2 as a gateway and 10% of the traffic should use R11 as a gateway.

Task 42.9 Authenticate the GLBP routers with a MD5 hashed password of "iPexpert133".

Task 42.10 Configure VRRP between R2 and R1 on the network 10.1.108.0/24. Virtual IP address is 10.1.108.144. When R2 is up and running, it should always be the master.

Task 42.11 Authenticate the VRRP routers with a password of "iPexpert".

Task 42.12 Configure HSRP between R8 and R2 on the network 10.1.108.0/24. Virtual IP address is 10.1.108.155. As long as R8 is up and running, it should stay the master and when an outage occurs, it should recover this role 1 minute after coming back online.

Task 42.13 When the ICMP echo from R8 to R6 fails, the priority should be decreased the minimum in such a way that R2 takes over the primary role.

Task 42.14 Authenticate the HSRP routers with a clear text password of "iPexpert".

You have completed Lab 42

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.

Lab 43: Configure and Troubleshoot IP/IOS Services (Part 5)

Technologies covered

- NAT Overload
- NAT Route-maps
- Static NAT
- Static PAT
- NAT no alias
- NAT no payload
- Policy NAT

Overview

You have been tasked to configure management services in your network.

The topology used in the lab will be the following:

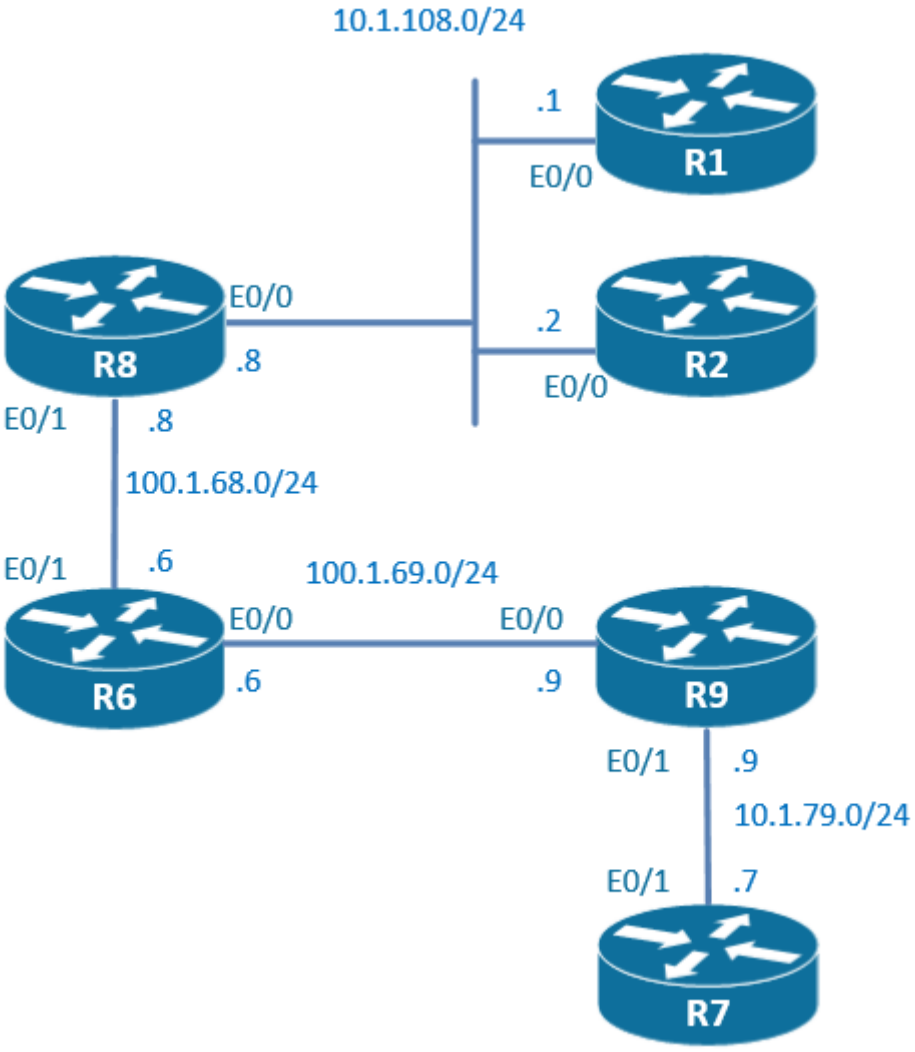

Estimated time to complete: 2 hours

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 43.1 On R7, configure a default route towards R9. 10.1.79.0/24 is the inside network, 100.1.69.0/24 is the outside network. Make sure that the ping from R7 to 100.1.69.6 is successful using a static NAT between 100.1.69.7 and 10.1.79.7.

Task 43.2 We don't want R9 to respond to the ARP request for 100.1.69.20. Clear the ARP cache and verify that the ping from R7 to 100.1.69.6 is unsuccessful.

Task 43.3 Ensure that the ping from R7 to 100.1.69.6 is again successful by configuring a static ARP entry on the router R6.

Task 43.4 Ensure that the payload will not be modified by the static NAT entry configured on R9.

Task 43.5 On R9, configure loopback0 with an IP address of 10.1.9.9/24. 10.1.9.0/24 is the inside network, 100.1.69.0/24 is the outside network.

Task 43.6 On R9, configure a dynamic NAT that maps the internal range 10.1.9.0/24 to the public address range 100.1.69.241-100.1.69.255. When no more address is available in the public range, a new connection will use a mapping of an already mapped public IP address with a different TCP port number.

Task 43.7 On R9, configure loopback1 with an IP address of 11.1.9.9/24. 11.1.9.0/24 is the inside network, 100.1.69.0/24 is the outside network.

Task 43.8 On R9, configure a dynamic PAT that maps the internal range 11.1.9.0/24 to the interface E0/0.

Task 43.9 On R9, enable the TCP small server service on TCP port 13 called "datetime".

Task 43.10 On R8, configure a default route towards R6. 100.1.68.0/24 is the inside network, 100.1.69.0/24 is the outside network. Make sure that the ping from R8 to 100.1.69.9 is unsuccessful and that the telnet 100.1.69.9 on port 4000 will return the daytime information.

Task 43.11 10.1.108.0/24 is the inside network, 100.1.68.0/24 is the outside network. Make sure that the ping from R1 to 100.1.68.6 is successful without configuring a default route pointing to R8 on R1. You have to use the ip nat outside command on R8.

Task 43.12 Ensure that you can telnet from R1 to 10.1.68.6 by using the add-route keyword in a command.

Task 43.13 On R2, configure a default route towards R8. Traffic coming from R2 should be statically NATed to the IP address 100.1.68.20. Use a route-map to achieve this task. Verify that you can ping from R2 to 100.1.68.6.

You have completed Lab 43

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.

Lab 44: Configure and Troubleshoot IP/IOS Services (Part 6)

Technologies covered

- IP precedence accounting
- IP output packet accounting
- IP access violation accounting
- MAC address accounting
- TCP optimization

Overview

You have been tasked to configure management services in your network.

The topology used in the lab will be the following:

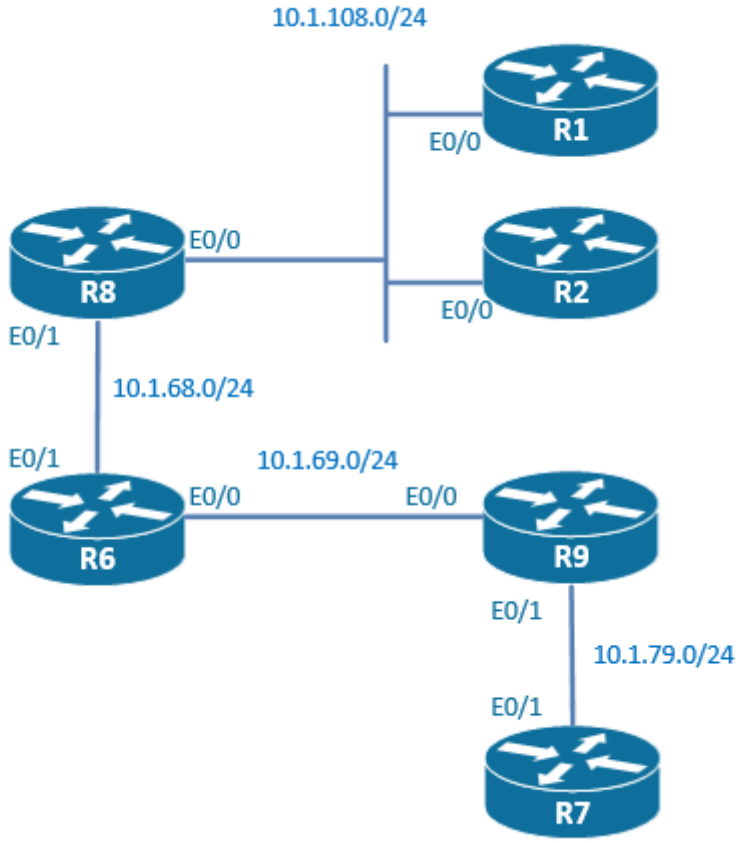

Estimated time to complete: 1 hour

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 44.1 On R7, perform on the E0/1 accounting based on IP precedence on received packets.

Task 44.2 Configure the following loopbacks:

R8 loopback0	10.1.8.8/32
R9 loopback0	10.1.9.9/32

Task 44.3 Enable OSPF area 0 on the path between R8 and R9, and advertise the loopback0 of R8 and R9 using network statements.

Task 44.4 On R6, create an access-list to block traffic going from loopback0 of R8 to the loopback0 of R9.

Task 44.5 Apply this access-list on the interface E0/0 and ensure that IP accounting displays the number of packets blocked by this access-list.

Task 44.6 On the interface E0/1 of R6, collect statistics about traffic per MAC address in the egress and ingress direction.

Task 44.7 On R8, activate high performance TCP options as described in RFC 1323.

Task 44.8 On R2, configure the outgoing TCP queue to contain a maximum of 10 packets.

Task 44.9 On R2, activate the TCP connection to discover of the minimum MTU size along the path of the TCP connection and therefore avoid fragmentation.

Task 44.10 R8 should wait for a maximum of 10 seconds to receive a TCP SYN.

Task 44.11 Make sure that R8 will not be affected by the “TCP silly window syndrome”.

You have completed Lab 44

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.

Lab 45: Configure and Troubleshoot IP/IOS Services (Part 7)

Technologies covered

- Netflow ingress and egress
- Netflow top talkers
- Netflow aggregation cache
- Netflow random sampling
- Netflow input filters

Overview

You have been tasked to configure management services in your network.

The topology used in the lab will be the following:

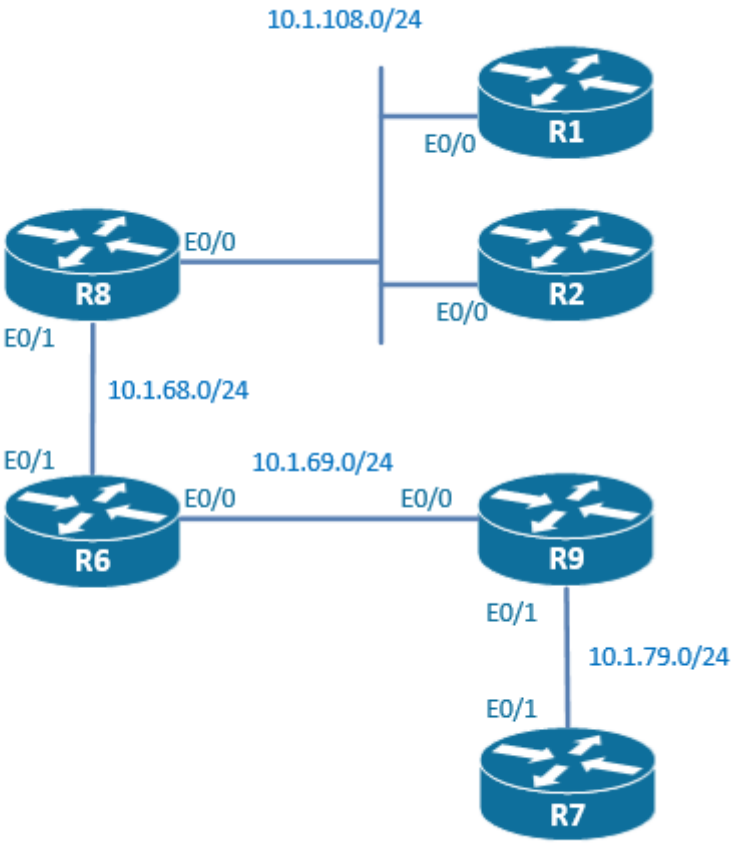

Estimated time to complete: 1 hour

Pre-Lab Setup

Logically connect and configure your network as displayed in the drawing below. You may also refer to the Diagram located within your configuration files for topology information.

This lab is intended to be used with online rack access provided by www.proctorlabs.com. Connect to the terminal server for the online rack, and complete the configuration tasks as detailed below.

Prerequisites

Load the initial configuration files before starting to work on the tasks.

Task 45.1 Setup R8 to collect Netflow version 9 statistics on E0/0 and E0/1, and to send them to server 10.1.33.33 on port 2333 in version 5 format. If R8 uses BGP, the peer AS should be included in exports. Make sure that the flows information is not duplicated.

Task 45.2 Configure R8 to export flow records every 2 minutes.

Task 45.3 On R8, ensure that a flow in the cache that was not refreshed during 10 seconds expires.

Task 45.4 Setup R9 to collect Netflow version 9 statistics on E0/0 and E0/1, and to send them to server 10.1.33.33 on port 2333 in version 5 format. Only 1 out of 50 packets should be captured by Netflow. Ensure that it is not 1 every 50 packets which is captured but randomly 1 out of 50 packets. Make sure that the flows information is not duplicated.

Task 45.5 On R6, configure Netflow on interface E0/1 and interface E0/0 to only capture traffic between 10.1.8.8 and 10.1.9.9. Only 1 out of 2 packets from this flow should be captured. Use a class-map called “NETFLOWCLASS” and a policy-map called “NETFLOWPOLICY”.

Task 45.6 On R1, configure Netflow version 9 on interface E0/0 to capture Netflow statistics in egress and ingress directions. The Netflow template should be sent every minute in version 9 to server 10.1.44.44.

Task 45.7 On R1, on the Netflow running on the E0/0, aggregate flow based of destination prefix present in the routing table. Never aggregate with a mask number lower than /24.

Task 45.8 On R2, setup Netflow to display in the command line the 20 top speakers going through interface E0/0. Sort the top speaker by bytes.

Task 45.9 On R2 interface E0/0, configure Netflow to capture the statistics for IPv6 packets.

Task 45.10 On R7, configure Flexible Netflow to collect the source and destination IP address, the flow direction, the next-hop IP address using a flow record called “IPEXPERTRECORD”.

Task 45.11 On R7, configure Flexible Netflow to export statistics to the server 10.1.55.55 on port 3444 every 30 seconds using a flow exporter called “IPEXPERTEXPORTER”.

Task 45.12 Apply a flow monitor called “IPEXPERTMONITOR” in the ingress and egress direction on interface E0/1.

You have completed Lab 45

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.

