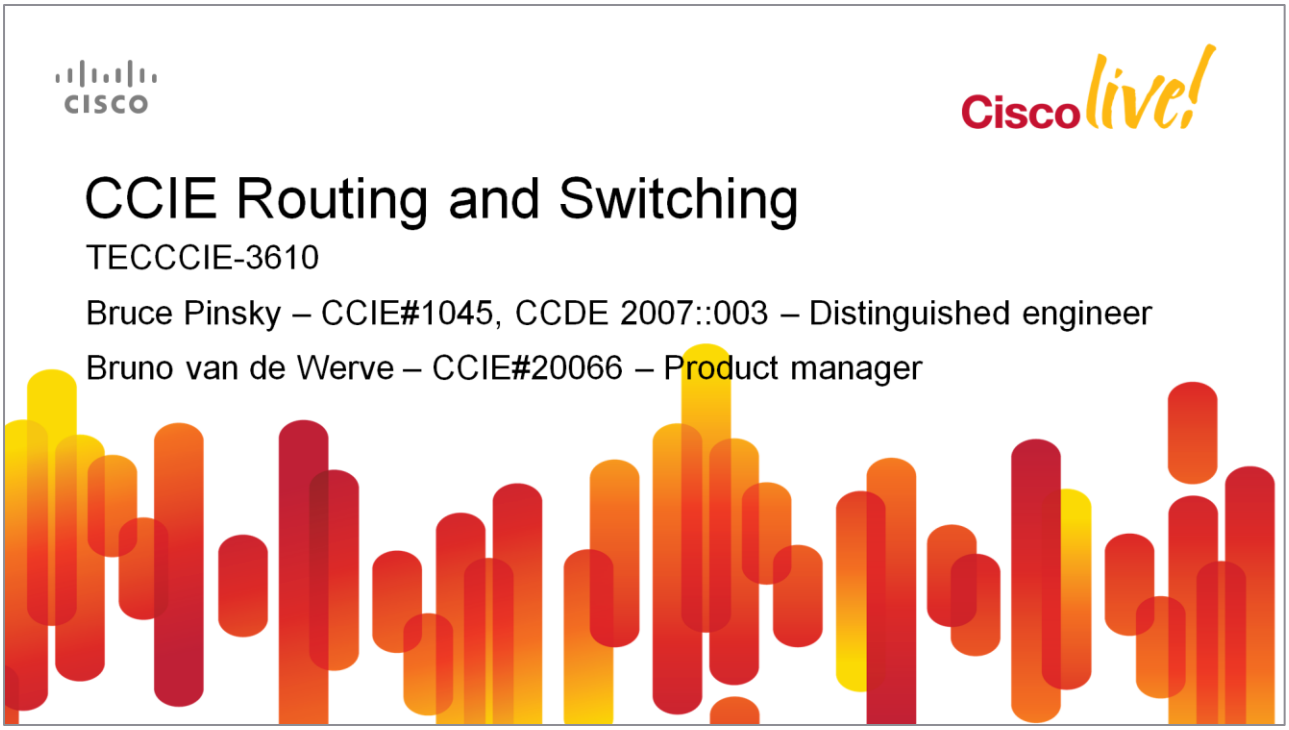

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

Disclaimer

- Not all the topics discussed today appear on every exam
- For time reasons, we're unable to discuss every feature and topic possible on the exam

- "Know what you don't know" is a short sentence that I personally kept in my mind when preparing for my CCIE, it was a very good way of reminding myself to always question my understanding of technical topics and look back for what i didn't fully understand in order to dig into it until I understood it from A to Z.

- Obviously, it is not very realistic to know everything about all topics covered by the exams. So what is the key here?

- Know where to look for things you know that you don't know!! Know where to find the information, and know how to find it quickly! Know the Cisco documentation! You are not expected to remember everything from the top of your head, but you are expected to be able to resolve issues quickly with the available resources provided to you, Cisco Documentation.

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

CCIEs Worldwide

- Most highly respected IT certification for more than 15 years
- Industry standard for validating expert skills and experience
- More than 22,000 CCIEs worldwide—less than 3% of all professionals certified by Cisco
- Demonstrate strong commitment and investment to networking career, life-long learning, and dedication to remaining an active CCIE

Cisco CCIE Certification

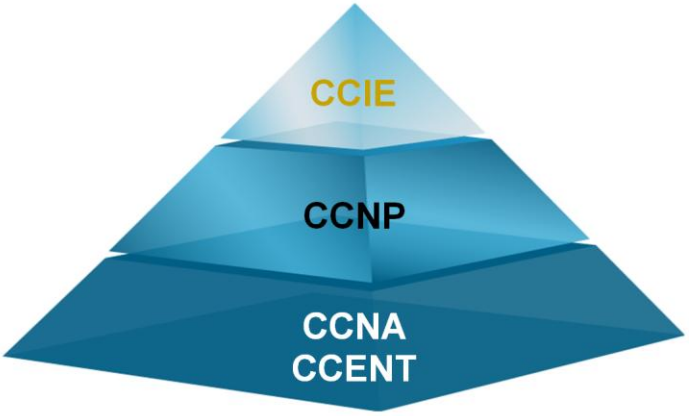

learningnetwork.cisco.com

Cisco CCIE Certifications

learningnetwork.cisco.com

CCIE Track	Major skills
R&S	Configure and troubleshoot complex converged networks
Security	Configure complex, end-to-end secure networks, troubleshoot environments, and anticipate and respond to network attacks
Service Provider	Configure and troubleshoot advanced technologies to support service provider networks
Storage	Configure and troubleshoot storage area networks over a variety of interfaces
Voice	Configure complex, end-to-end telephony, as well as network, troubleshoot, and resolve VoIP-related problems
Wireless	Plan, design, implement, operate, and troubleshoot wireless network and mobility infrastructure
SP Operation	Troubleshoot network performance problems, identify implementation plans problems, manage network fault management system and operation processes

- Refer to CLN for specific info about each CCIE track
- Along with CCIE, there is the CCDE which is the expert-level certification in Network Design (not listed in this slide)

CCIE Information World Wide

* Updated on 02/17/2011

Total	22,741*
RS	16,647
Security	2,537
SP	1,951
Voice	1,414
SAN	145
Wireless	42
CCDE	39
SP Ops	5

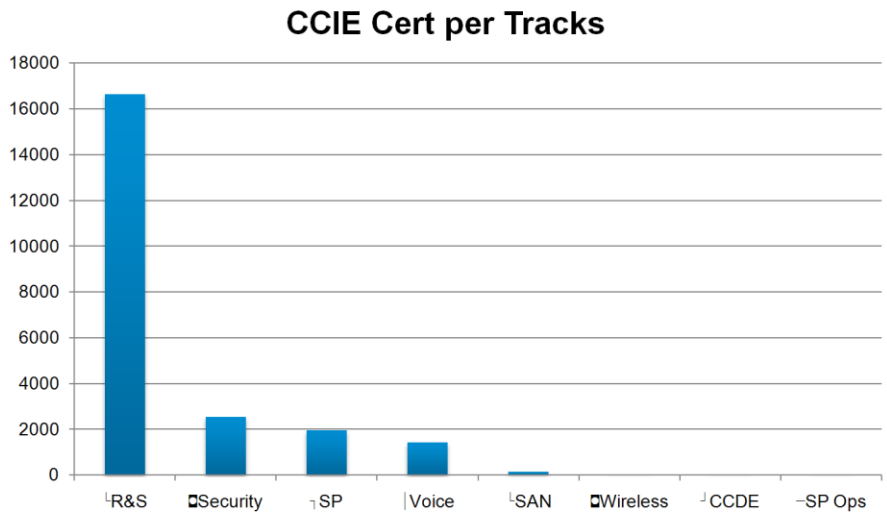

CCIE Information World Wide

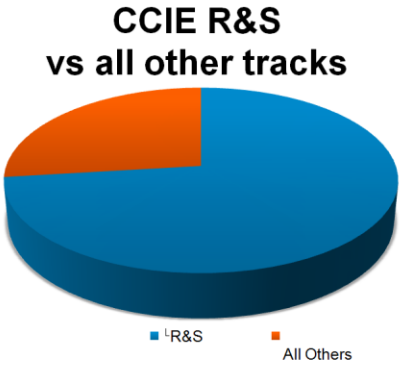

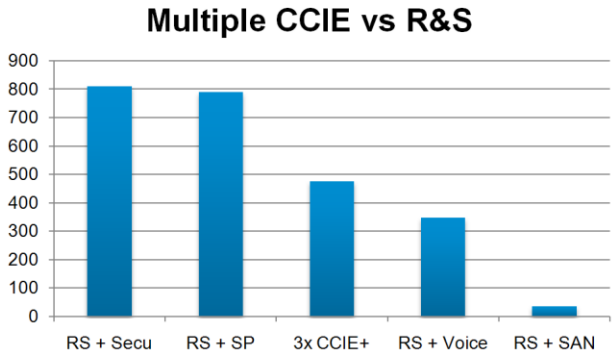

Updated on 02/17/2011

CCIE Information World Wide

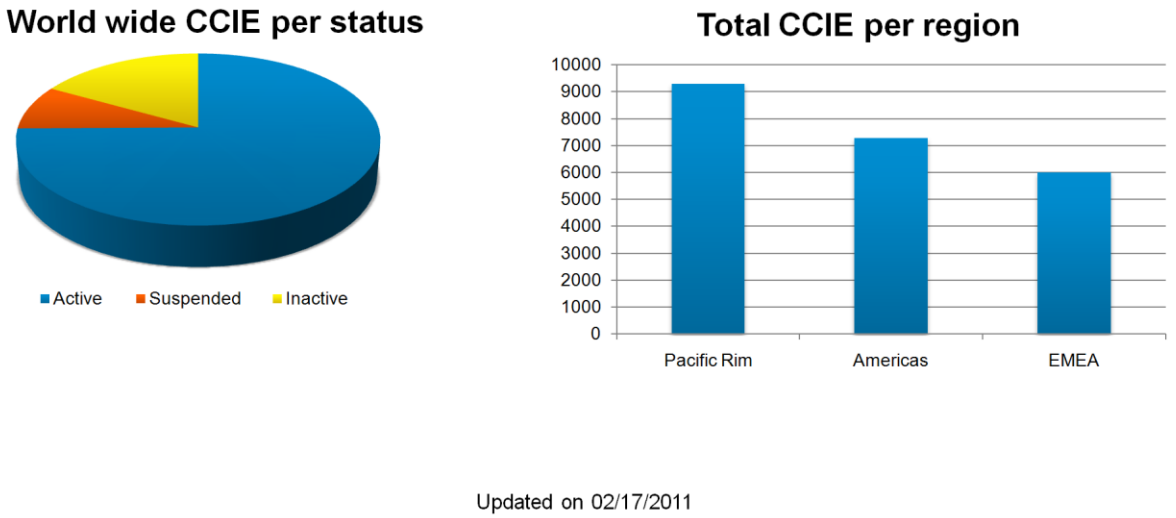

- Active
- Suspended
- Inactive
- Total = Active + Suspended + Inactive

CCIE Information World Wide

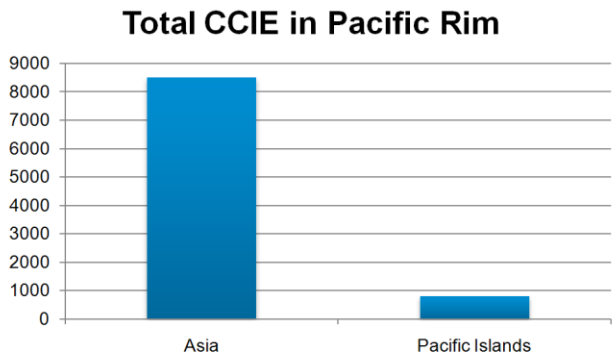

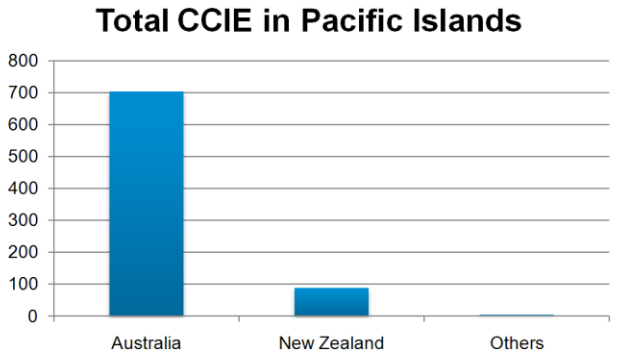

Updated on 02/17/2011

- Total = Active + Suspended + Inactive

Certification Process

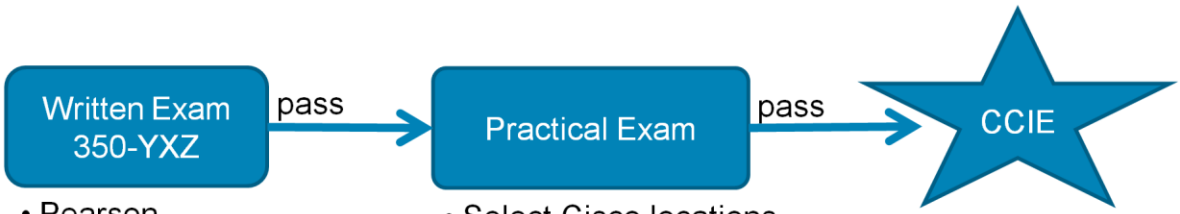

- Pearson
- 2 hours
- Multiple choices
- Simulations
- No documentation
- Immediately scored

- Select Cisco locations
- 8 hours
- Configurations
- Troubleshooting
- Cisco documentation
- Scored within 48h

CCIE Exam Development Process

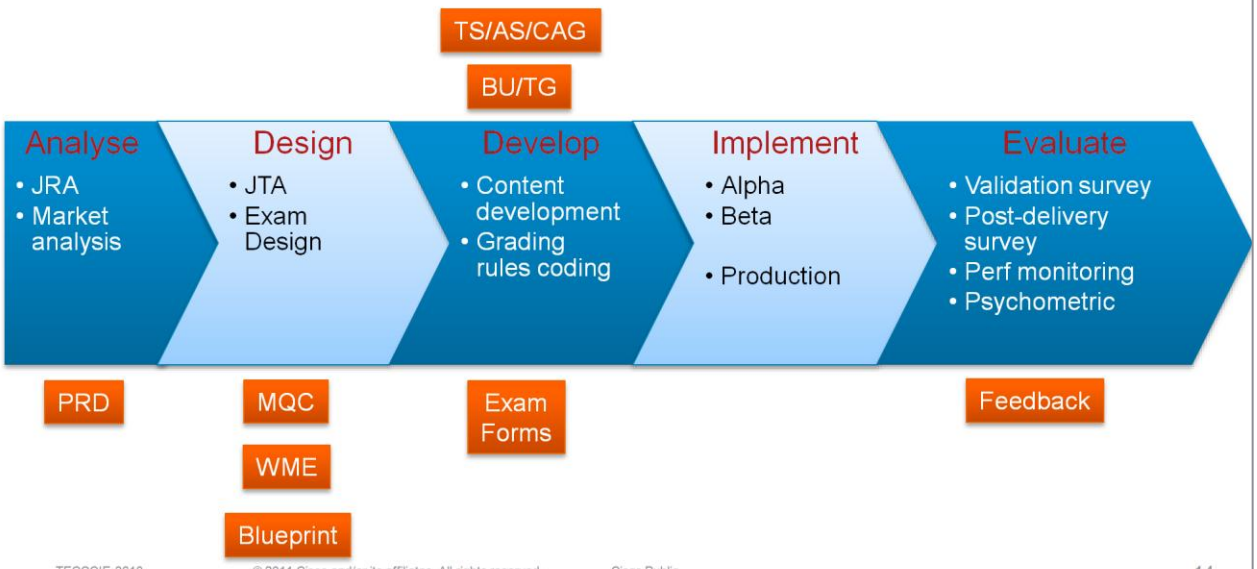

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

CCIE Permanent Lab Locations: R&S

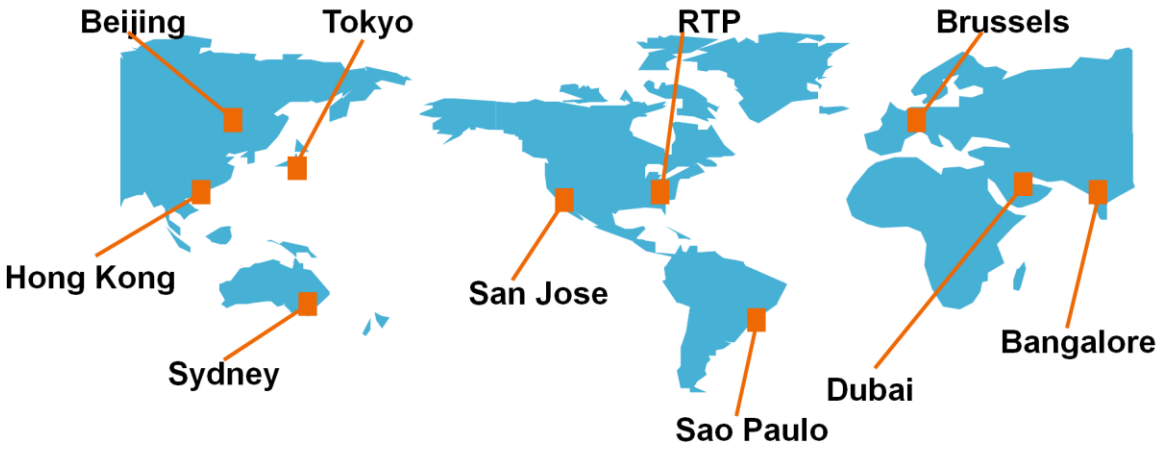

Permanent CCIE R&S Lab Locations

CCIE Routing and Switching Version 4

- Most popular track, over 80% of CCIE candidates attempt R&S first
- Expert-level knowledge of LAN and WAN interfaces, Routing Protocols, and variety of routers and switches
- Expert-level in troubleshooting to solve complex connectivity problems and apply solutions to increase bandwidth, improve response times, maximize performance, improve security, and support global applications

Recent Changes to CCIE R&S

- New certification standards released on May 5, 2009
- Reflect growth of network as a service platform
- Align to employer expectations and job tasks of expert level network engineers
- No topics removed; less coverage of lower level skills

Recent Changes to CCIE R&S

- New areas include:
  - Troubleshoot networking issues
  - MPLS and VPN networking
  - Implementing performance routing and optimised edge routing
  - EIGRPv6
  - IPv6 multicast
  - PPPoE
  - …

CCIE R&S v4.0 Certification

- V4.0 exams released October 18, 2009
- Written and lab exams refreshed with new questions
- Added coverage of MPLS and VPN
- Written exam added simulations (flash based items)
- Lab exam added hands-on troubleshooting
- Exam durations and pricing remained same

CCIE R&S Mobile Labs: 2011

Note: Dates with an "R" Status are *confirmed events with location information. These events are available for open registration.*

Lab Locations	2011 Dates	Status	Lab
Shanghai, China	Jan 10 - 14	C	Routing & Switching, Security
Lagos, Nigeria	Jan 24-28	R	Routing & Switching, Security
Kiev, Ukraine	Jan 24 - 28	C	Routing & Switching, Security
Cairo, Egypt	Feb 6 - 10	C +	Routing & Switching, Security
Riyadh, Saudi Arabia	Feb 19 - 24	R	Routing & Switching, Security
Nairobi, Kenya	Mar 7 - 12	R	Routing & Switching, Security
Moscow, Russia	Mar 21 - 25	R	Routing & Switching, Security
Mexico City D.F., Mexico	Apr 4 - 9	R	Routing & Switching, Security
Istanbul, Turkey	Apr 18 - 22	R	Routing & Switching, Security
Riyadh, Saudi Arabia	Apr 30 - May 4	R	Routing & Switching, Security
Moscow, Russia	May 16 - 20	R	Routing & Switching, Security
Johannesburg, South Africa	May 30 - Jun 3	R	Routing & Switching, Security
Lagos, Nigeria	Jun 13 - 18	R	Routing & Switching, Security
Dalian, China	Jun 27 - Jul 1	C	Routing & Switching, Security
Bangkok, Thailand	Jun 27 - Jul 1	R	Routing & Switching, Security
Istanbul, Turkey	Jul 25 - 29	R	Routing & Switching, Security

Dates for 201[illegible] available at: [illegible]

- R = Confirmed events and are available for open registration, it does not indicate availability. You must Login to the CCIE Database to view availability and register for lab exams.
- C = Canceled, due to low enrollment
- + = For questions related to this cancellation please open a case with www.cisco.com/go/certsupport.

CCIE R&S Exam Philosophy

- The routing and switching exam tests your ability to apply configuration knowledge, troubleshoot skills and skill to new situations; **it is not a design test, nor is it always a test of "best practices" for use in the field**

CCIE R&Sv4: Minimally Qualified Candidate

"The minimally qualified Routing and Switching CCIE can
abstract functional elements of a complex network environment,
understand how infrastructure components interoperate,
grasp subtle issues,
perceive problem areas,
and quickly resolve problems.

The expert's fluency makes them ideally suited for
configuring and validating implementations,
troubleshooting critical network issues,
and participating in network design teams. "

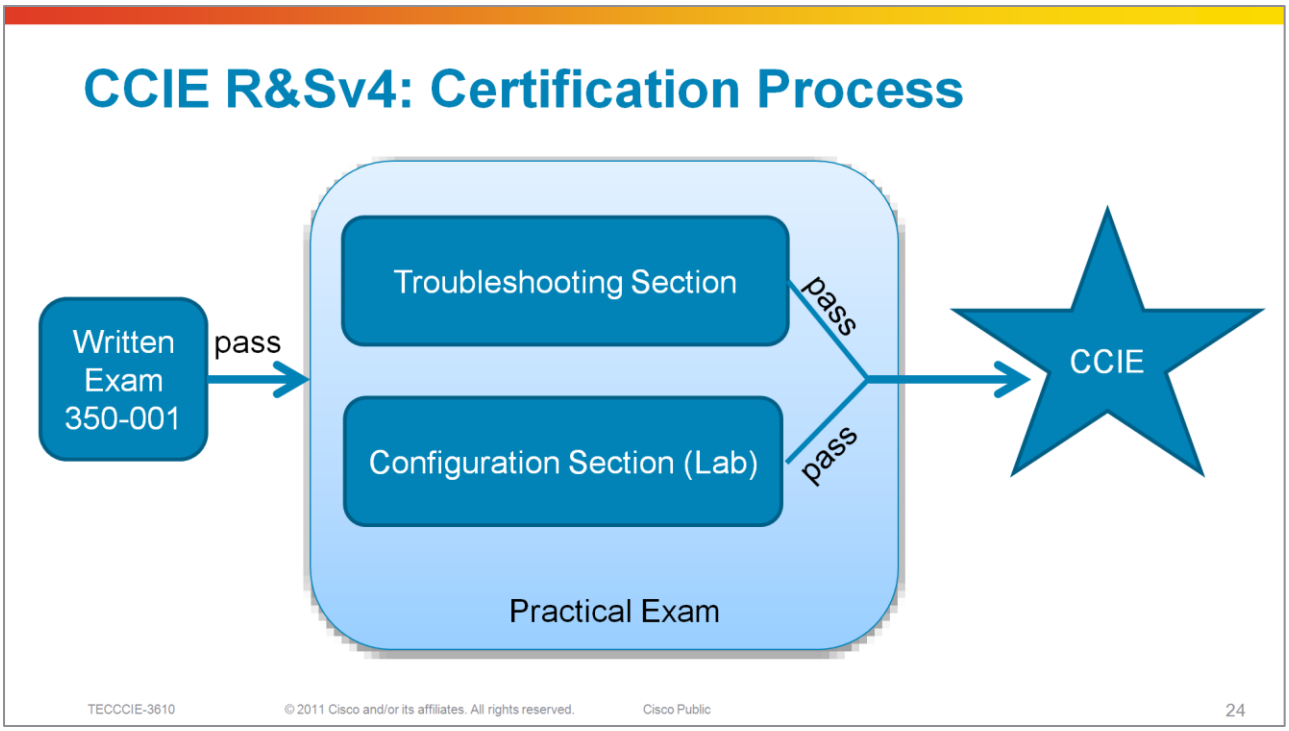

- Candidates will need to achieve minimum score on each section to pass the Practical exam.
- Candidates take both sections, independent of passing on the first section or not.
- Re-reads are available for the R&S track

CCIE R&Sv4: Written Exam Topics

- Implement Layer 2 Technologies
- Implement IPv4 – includes BGP
- Implement IPv6
- Implement MPLS Layer 3 VPNs
- Implement IP Multicast
- Implement Network Services
- Implement Quality of Service
- Implement Network Security
- Troubleshoot a Network
- Optimise the Network
- Evaluate proposed changes to a Network

Same as for the Lab exam!

- Candidates should check the Lab Blueprint regularly prior to taking the exam

- **Written lays foundation to the Lab Exam**

Most topics are identical, only the last one is specific to the Written exam:

- **Evaluate proposed changes to a Network**

  11.01 Evaluate interoperability of proposed technologies against deployed technologies

  (a) Changes to routing protocol parameters

  (b) Migrate parts of a network to IPv6

  (c) Routing Protocol migration

  (d) Adding multicast support

  (e) Migrate spanning tree protocol

CCIE R&Sv4: Lab Exam Structure

Modular structure

- Independent modules:

  2h-Troubleshooting

  6h-Configuration

- Composite cutscore:

  Must pass **both** modules, each with min 80%

TECCCIE-3610 © 2011 Cisco and/or its affiliates. All rights reserved. Cisco Public 26

- Benefits of the Modular structure

- “Minimize content exposure to maximize content security”
- Increase content production pace
- quicker to develop and test smaller modules than 8h lab!
- Increase the amount of exam versions with ~same development effort thanks to combinations of modules:
- Ex: 3x TS + 3x CFG = 9x 8h exams
- Vs 3x 8h exam

CCIE R&Sv4: TS Module

2h Troubleshooting

- Separate topology of ~30 virtual routers and switches*
- Scenario is fully preconfigured but contains faults
- Composed of ~10 independent incidents, worth 22/100 points
- 2h maximum (visible countdown timer + 15min warning)
- Questions are “symptom-based”
- Verifications are “result-based” + constraints

- IOU and L2IOU

CCIE R&Sv4: TS Module

Virtual infrastructure

- Virtual router = IOU: separate platforms for IOS (Solaris)
- Virtual switch = L2IOU: based on IOU + Switching code
- Mixing L2 and L3 links, like in real
- No Ethernet controller (yet)

  interfaces always up/up, never up/down
- L2IOU: More features under development

- Ethernet controller is going to be simulated in the very next release of L2IOU, so this limitation will not apply anymore very soon

CCIE R&Sv4: CFG Module

6h Configuration

- Separate topology of 9 hardware devices (5x ISR routers + 4x 3560 switches)
- Unique 6h-scenario based on a static rack (physical cabling is the same for all content)
- Composed of ~28 questions* worth 78/100 points
- Any time gained in TS is granted to CFG

- Questions are ‘opportunities to score’, it doesn t really mean anything as per the quantity of content being assessed in the exam.

  Ex OSPF config with multiple Areas might be asked in a single question or in multiple, content stays the same but the opportunities to score is different.

  Objective is to keep the exam forms fair by maintaining the same opportunities to score and as much as possible same content quantity (evaluated by testing phases)

CCIE R&Sv4: CFG Module

Hardware infrastructure

- Almost same rack as R&Sv3:
  - No Frame-relay switch (only back to back serials)
  - 5 ISR routers compared to 6
  - All identical 3560 switches (no mix of 3550 and 3560)
- 2x c1841 + 3x c3825 running 12.4T
- 4x c3560 running 12.2S

- The lab exam tests any feature that can be configured on the equipment and the IOS versions indicated on the CCIE web page. You may see more recent IOS versions installed in the lab, but you will not be tested on the new features of a release unless indicated below.

R&S Lab Exam Overview

- Candidates build a network to a series of supplied specifications
- The point values for each question are shown on the exam
- Some questions depend upon completion of previous parts of the network
- Each candidate has his/her own PC and dedicated rack of equipment, which may or may not be located in lab

R&S Lab Exam Overview

- Equipment should not require hardware configuration or cabling. Contact proctor if you discover a problem requiring intervention of this kind
- Report any suspected equipment issues to the proctor as soon as possible; adjustments cannot be made once the exam is over

Rack Access

Rack Connection Method

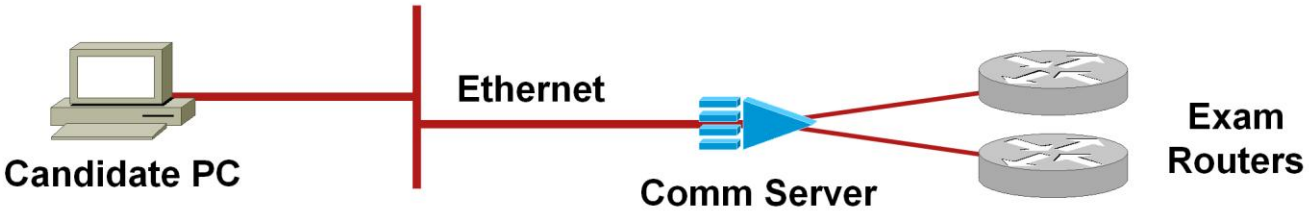

- The Comm Server is pre-configured
- The Candidate PC has the terminal emulator pre-configured to access all routers and switches, browsers and any other needed application (calc, notepad, …)

Passwords

- All routers and switches have a startup configuration: hostnames, passwords, line setup, and IP addresses for primary interfaces are already configured; since all tests require the router to be accessible via the VTY and AUX ports, do not change these established configurations

Standard Restrictions

Unless specified within the exam you are **NOT** allowed to use

- Static routes (of any kind)
- Default routes
- **Dynamic routes to null are permitted

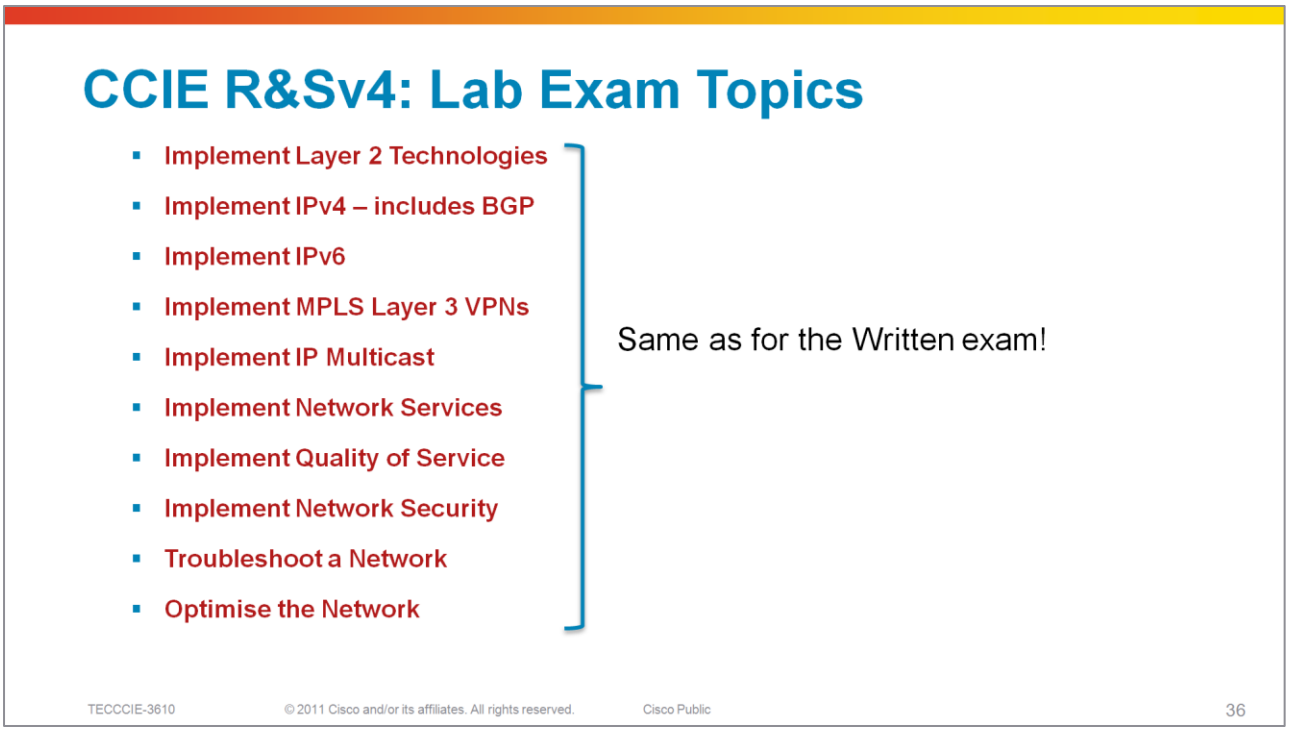

- Candidates should check the Lab Blueprint regularly prior to taking the exam
- Changes to the blueprint will be announced 6months in advance

CCIE R&Sv4: Lab Delivery System

LDS

- Web GUI delivering all the exam content
  (questions, diagrams, guidelines, access to devices, documentation …)
- Transparent servers backend managing the user sessions and holding the IOU instances for each users
- Benefits:
  - Flexibility in content delivery format
  - Content updates effective globally immediately
- No printed exam workbook

CCIE R&Sv4: Lab Delivery System

LDS: Example of TS session

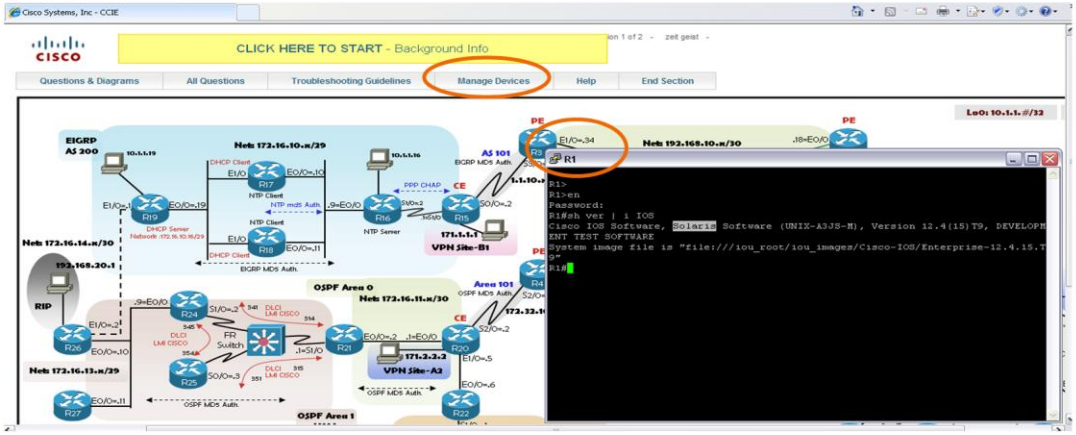

CCIE R&Sv4: Lab Delivery System

LDS: Example of CFG session

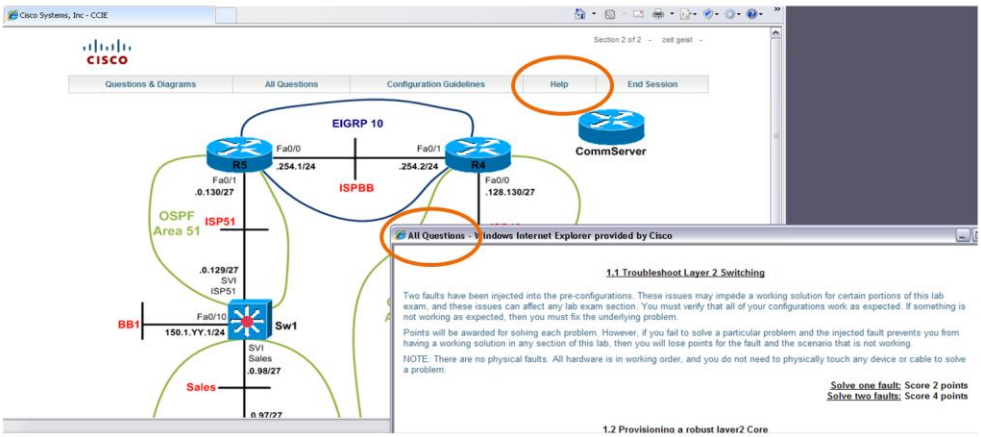

CCIE Lab Exam: Grading

- Proctors grade all lab exams
- Automatic tools aid proctors with simple grading tasks
- Automatic tools are never solely responsible for lab exam grading
- Completed and submitted within 48 hours
- Partial credit is not awarded on questions
- Points are awarded for working solutions only
- Some questions have multiple solutions

CCIE R&Sv4: Automated Grading tool

- Data collection engine :
  - Collects data from any devices in any order (any scenario)
- Grading rules engine :
  - Run grading scripts against collected output (regex)
  - Maintain a detailed report for each candidate attempt
- Benefits :
  - Faster and more consistent grading due to automation
  - Scalable and supporting expansion objectives
- Critical point in Content quality

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

CCIE R&Sv4: Lab exam's General Guidelines

Read them at least once!

- Configuration exam's principles
  - Embedded troubleshooting
  - Incidents' or questions' dependency
- Rules of exam's format
  - Variables due to POD number
- Configuration restrictions
  - Console line and passwords

- CFG Guidelines:
  - READ ALL QUESTIONS in each section before proceeding with any configuration.
  - BEFORE STARTING THE EXAM, confirm that all devices in your rack are in working order. During the exam, if any device is locked or inaccessible for any reason, you must recover it. When you complete this exam, ensure that all devices are accessible to the grading proctor. A device that is not accessible for grading cannot be marked and may cause you to lose substantial points.
  - Knowledge of implementation and troubleshooting techniques is part of the skills tested in the configuration section of the lab exam.
  - If you suspect that there may be hardware problems with your equipment, contact the lab proctor immediately.
  - Points are awarded for working configurations only. Test the functionality of all of the requirements before you complete your exam. As you configure one part of the exam, you may break a previous requirement or configuration.
  - No partial points can be granted for any question, all requirements need to be fulfilled in order to receive the points for the question. Some requirements depend on other question's requirements, either before or after the current question.
  - You will be presented with preconfigured routers and switches. DO NOT change the following configuration on the devices:
    - Hostname
    - Enable password "cisco"
    - Console line configuration
  - Throughout the exam, assume these values for variables:
    - YY is your 2-digit rack number. For example, Rack number 1 will use "01", Rack number 11 will use "11".
    - B is the backbone number. For example, the value for Backbone 2 is 2. Backbone subnets use the following address convention: 150.B.YY.0/24. Do NOT change backbone addresses unless you are instructed to do so.
    - X is the router number. For example, the value of X for Router 1 is 1.
    - Z is any number.
  - In any configuration where additional addressing may be necessary, use only the major network as displayed in Diagram 1. Ensure that it does not conflict with a network that is already used in your network.
  - Unicast or multicast static and default routes are NOT permitted unless permission to use them is directly stated in a specific question. This restriction includes floating static routes and those routes that were generated by a routing protocol. Routes to null 0 that are generated as a result of a dynamic routing protocol solution are permitted.
  - Save your configurations frequently.

CCIE R&Sv4: Lab exam's General Guidelines

Read them at least once!

- Troubleshooting exam's principles

  Do not remove any feature preconfigured!

  Resolve the misconfiguration!!

  Ex. ACL, PBR, CoPP, MQC, etc

- TS Guidelines:
  - You will be presented with preconfigured routers and Frame Relay switches in the topology. DO NOT change the following configuration on the devices:
    - Hostname
    - Enable password "cisco"
    - Console line configuration
  - For all the authentication configurations in the lab, password is "cisco" unless changed to introduce a break. Do NOT change the AAA configuration unless explicitly stated in a question.
  - Points are awarded for finding AND fixing inserted faults in the presented fully configured topology. An inserted fault is an introduced break for a scenario that was previously working. Depending on the scenario, fixing inserted faults could require one or multiple command lines on the same or multiple devices.
  - The resolution of one incident may depend on the resolution of previous incident(s). The dependency will not be visible if incidents are resolved in sequence.
  - There are NO physical faults introduced in the presented topology.
  - DO NOT change any routing protocol boundaries. Refer to the provided diagrams.
  - DO NOT REMOVE ANY FEATURE CONFIGURED IN ORDER TO RESOLVE AN INCIDENT, YOU MUST RESOLVE THE MISCONFIGURATION RATHER THAN REMOVING IT ALL.
  - Static and default routes are NOT permitted unless preconfigured. These restrictions include floating static and those generated by a routing protocol. Routes to Null 0 that are generated as a result of a dynamic routing protocol solution are permitted.
  - Tunneling and policy-routing are NOT permitted unless preconfigured.
  - Dynamic Frame Relay mappings are NOT permitted.
  - Points will be deducted for every incident in which candidate uses a prohibited solution.
  - Candidates have control of all required devices in the topology.
  - If required to verify the reachability from a host machine during the lab exam, use the ping command with source option on the router that is shown connected to the subjected host in the diagram.

Study plan: Content topics review

Approach

- Quick review of the essential concepts for each main topics
  - Organised as per the Configuration exam' sections
- Sample question per main topic using a sample topology
  - Discussion on question's reading and interpretation
  - Review of Solution and Verification steps
- Frequently missed topic
- References

https://learningnetwork.cisco.com/docs/DOC-4375

Practice Lab: Sample topology

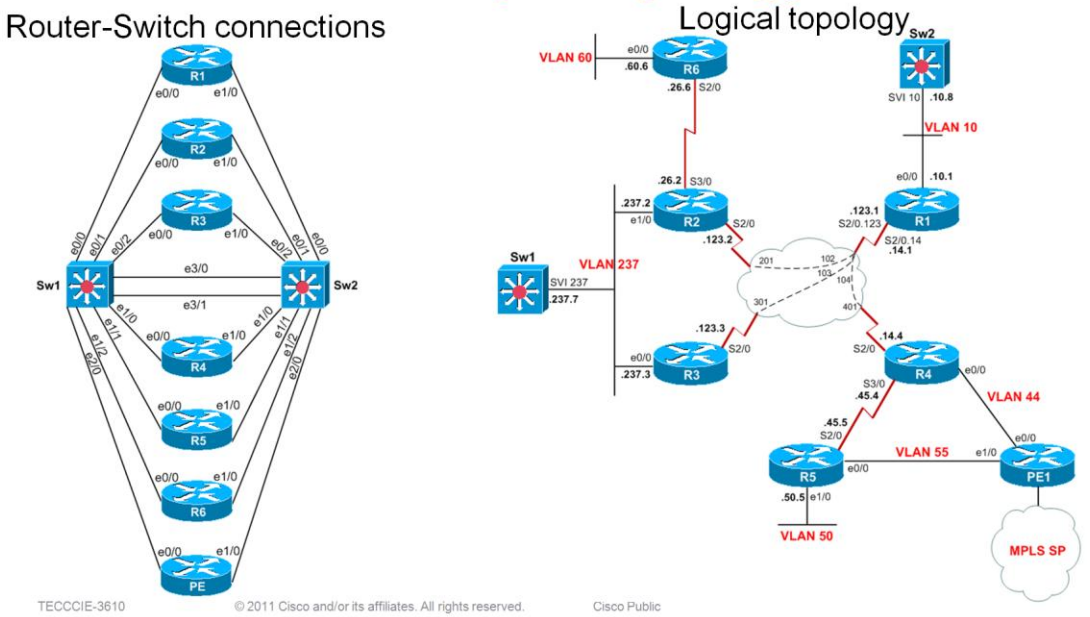

TECCCIE-3610 © 2011 Cisco and/or its affiliates. All rights reserved. Cisco Public 46

- This topology will be used throughout the presentation and sample questions will refer to it
- This is a pretty good example of the type of diagram that will be available to you during the actual exam

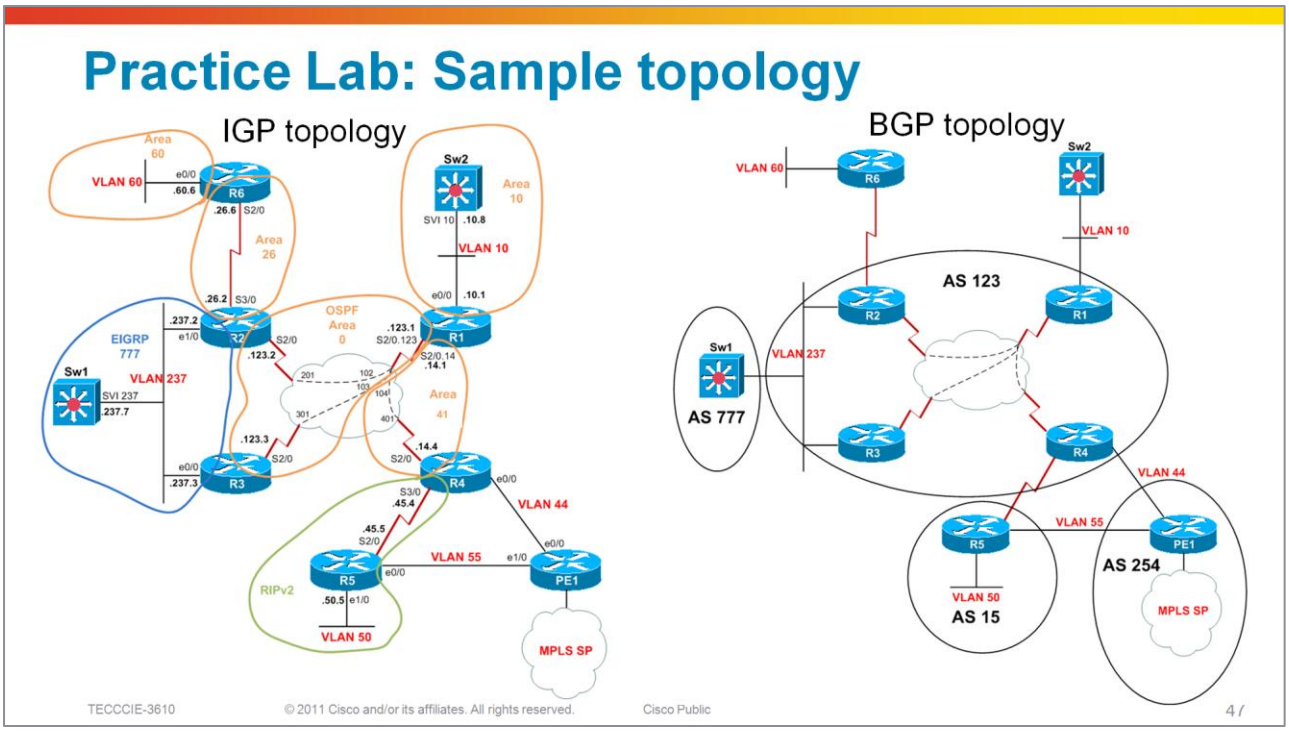

- You will typically have multiple diagrams: the logical diag with L3 links only
- The IGP diag with RP domains
- The BGP diag with AS domains
- Eventually more diags as appropriate

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

Section 3: Study plan: Content topics

Part 1	Switching:
2.1	Frame-Relay
2.2	LAN Switching
Part 2	Routing
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

Frame Relay Concepts

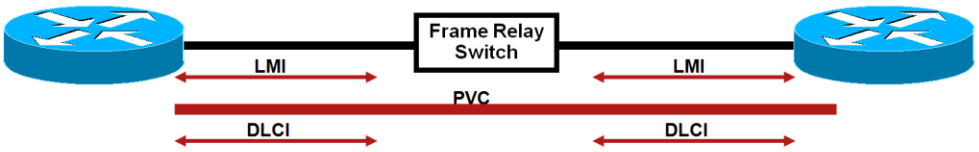

- DLCI Data-link connection identifier
- LMI Local Management Interface
- PVC Permanent Virtual Circuit
- CIR Committed Information Rate
- Bc Committed Burst
- Be Excess Burst
- DE Discard Eligible
- BECN Backward Explicit Congestion Notification
- FECN Forward Explicit Congestion Notification

Frame Relay Concepts

- Layer 2 to Layer 3 mapping
  - Inverse ARP
  - Static Mapping
- Topologies
  - Hub and Spoke
  - Point-to-point
  - Multipoint
- QoS

- FR QoS: LFI/MLPPP

Sample Lab Question – Frame-Relay

- **Frame Relay Configuration**

  Refer to the network diagram for frame-relay DLCIs and PVC information. Subinterfaces are pre-configured and numbered according to the third octet of the IP address configured on that interface.

  - Configure the frame-relay network between R1, R2 and R3 so that you can ping between all routers over the R1-R2-R3 frame-relay cloud.
  - Configure the point-to-point link between R1 and R4. Do not create a subinterface on R4. Ping between R1 and R4 when completed.

  Score: 2 Points

- Take few seconds to carefully read the question and think about the requirements and constraints and possible solutions

Practice Lab: Sample topology

Focus on Frame-Relay cloud

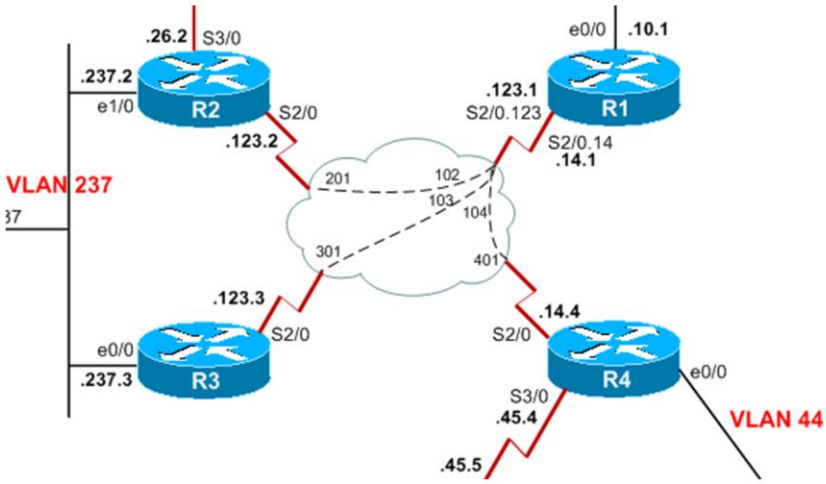

Sample Lab Question – Frame-Relay

- **Hints, requirements and constraints in the stem:**

  Refer to the network diagram for frame-relay DLCIs and PVC information. (…)

- Do not change the DLCIs!

Sample Lab Question – Frame-Relay

- **Hints, requirements and constraints in the stem:**

  (…) Subinterfaces are pre-configured and numbered according to the third octet of the IP address configured on that interface.(…)

- Do not change the pre-configured IP addresses and sub-interface's numbers!

Sample Lab Question – Frame-Relay

- **Hints, requirements and constraints in the stem:**
  - Configure the frame-relay network between R1, R2 and R3 so that you can ping between all routers over the R1-R2-R3 frame-relay cloud.(…)
- Only requirement is to be able to ping between the 3 devices!
- Can use inverse-arp or static mapping
- No words about mapping local IP… up to you.

- R1

```
interface Serial2/0
 encapsulation frame-relay
!
No int s2/0.123
interface Serial2/0.123 multipoint
 frame-relay map ip 172.16.123.3 103 broadcast
 frame-relay map ip 172.16.123.2 102 broadcast
 ip address 172.16.123.1 255.255.255.0
```

- R2:

```
interface Serial2/0
 encapsulation frame-relay
```

Sample Lab Question – Frame-Relay

- **Hints, requirements and constraints in the stem:**
  - Configure the point-to-point link between R1 and R4. Do not create a subinterface on R4. Ping between R1 and R4 when completed.
- Notice the constraint about R4's subinterface!
- Can use inverse-arp or static mapping on R4

- R1

```
interface Serial2/0.14 point-to-point
 frame-relay interface-dlci 104
 ip address 172.16.14.1 255.255.255.0
```

- R4

```
interface Serial2/0
 ip address 172.16.14.4 255.255.255.0
 encapsulation frame-relay
 frame-relay interface-dlci 401
```

Sample Lab Question – Frame-Relay

- **Verification steps:**

1. Successful ping from/to all devices
2. FR map on R1 and R4: sub-if, IP, DLCI and PVC active

```
R1#sh fram map
Serial2/0.123 (up): ip 172.16.123.2 dlci 102(0x66,0x1860), static,
              broadcast,
              CISCO, status defined, active
Serial2/0.123 (up): ip 172.16.123.3 dlci 103(0x67,0x1870), static,
              broadcast,
              CISCO, status defined, active
Serial2/0.14 (up): point-to-point dlci, dlci 104(0x68,0x1880), broadcast
        status defined, active
R1#

R4#sh fram map
Serial2/0 (up): ip 172.16.14.1 dlci 401(0x191,0x6410), dynamic,
              broadcast,
              CISCO, status defined, active
R4#
```

Sample Lab Question – Frame-Relay

- **Frequently missed points:**
  - Broadcast statement missing, wrong DLCI
  - Requirement missed
  - Constraints not respected
  - Main guidelines not respected

References

- CiscoPress®

  Cisco Frame Relay Solutions Guide, Chin

- Technology Handbook:

  http://docwiki.cisco.com/wiki/Frame_Relay

- CCO Documentation:

  http://www.cisco.com/en/US/tech/tk713/tk237/tsd_technology_support_protocol_home.html

  http://www.cisco.com/en/US/docs/ios/wan/configuration/guide/wan_cfg_frm_rly_ps6441_TSD_Products_Configuration_Guide_Chapter.html

Section 3: Study plan: Content topics

Part 1	Switching:
2.1	Frame-Relay
2.2	LAN Switching
Part 2	Routing
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

LAN Switching Concepts & Protocols

- Layer 2 VLAN's (Normal, Extended, Native)
- Layer 2 Ethernet Trunk
  IEEE 802.1q, ISL (Cisco Proprietary)
- CDP
- DTP (Dynamic Trunking Protocol)
- VLAN Trunk Protocol (VTP)
- EtherChannels
  Access Ports, L2 Trunks or L3 Links
  PAgP, IEEE 802.3ad LACP

LAN Switching Features

- Spanning Tree
  - PVST, MSTP, RSTP
- Spanning Tree Features
  - Portfast
  - Bpduguard, Bpdufilter
  - Uplinkfast, Backbonefast
  - Etherchannelguard
  - Rootguard, Loopguard
- SVI (Switched Virtual Interface)

* Portfast: Bypass learning / listening state in access or trunk ports (BLK->FWD).

- (config)# spanning-tree portfast default <! Enable portfast on all non trunk ports.
- (config-if)# spanning-tree portfast
- (config-if)# spanning-tree portfast trunk

- * Bpduguard: Block portfast interfaces if bpdus received.
- (config)# spanning-tree portfast bpduguard default <! Only for port fast interfaces.
- (config-if)# spanning-tree bpduguard enable

- * Bpdufilter: Stops sending and receiving BPDUs in portfast interfaces.
- (config)# spanning-tree portfast bpdufilter default <! Disable portfast and bdpufilter if bpdus.
- (config-if)# spanning-tree bpdufilter enable <! Disable STP on the interface.

- * Uplinkfast: Quickly switch root port in the event of an uplink failure (BLK->FWD).
- - Flood the network with multicast from learnt addresses at 150 pps (or max-update-rate)
- - UplinkFast cannot be enabled on VLANs configured with switch priority.
- (config)#spanning-tree uplinkfast

LAN Switching Features

- Private VLAN
- (R)SPAN
- SVI (Switched Virtual Interface)

Sample Written Question - DTP

- What trunk mode combination would not produce an operational ISL trunk?

A. Local: auto	Remote: auto
B. Local: on	Remote: auto
C. Local: nonegociate	Remote: on
D. Local: nonegociate	Remote: nonegociate
E. Local: auto	Remote: desirable

- Solution: A
  If both sides are set to Auto, trunk will never come up
- If set to ‘nonegociate’, it implies that an encaps is configured (config not accepted when no encaps configured first)

Sample Lab Question – LAN Switching

- **VTP Configuration**

  Configure VLAN Trunking Protocol (VTP) as per the following requirements:

  - The VTP domain name is ”CCIE”.
  - Secure the VTP advertisements to be sent. Use ”CCIE" as your key.
  - SW1 will propagate all VLAN configuration changes to SW2.
  - In the future, these four switches will be configured as VTP transparent hosts. They should NOT inspect the VTP domain name and version, and they should support unrecognised type, length, value (TLV).
  - DO NOT Use VTP v3.

  Score: 2 Points

Practice Lab: Sample topology

Router-Switch connections

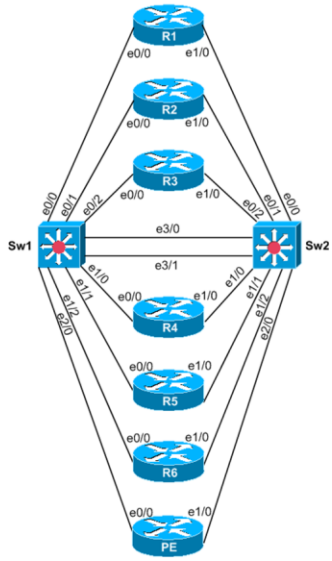

Sample Lab Question – LAN Switching

- **Hints, requirements and constraints in the stem:**
  - The VTP domain name is ”CCIE”.
  - Secure the VTP advertisements to be sent. Use ”CCIE" as your key.
  - SW1 will propagate all VLAN configuration changes to SW2.(…)
- Very clear requirements about VTP config !
- SW1 must be VTP server and SW2 VTP client or server

- SW1

vtp domain CCIE

vtp mode server

vtp password CCIE

- SW2

vtp domain CCIE

vtp mode client

vtp password CCIE

Sample Lab Question – LAN Switching

- **Hints, requirements and constraints in the stem:**
  (…)
  - In the future, these four switches will be configured as VTP transparent hosts. They should NOT inspect the VTP domain name and version, and they should support unrecognised type, length, value (TLV).
  - DO NOT Use VTP v3.
- "In the future" means not yet! This would violate the previous requirement…
- Requirement describes a specificity of VTP version 2

- SW1 and SW2

vtp version 2

Sample Lab Question – LAN Switching

- **Hidden/Implicit requirements in the stem:**
  (…)
  - SW1 must propagate all VLAN configuration changes to SW2
- VTP relies on trunks!
- Configure trunks, either with ISL or DOT1q
- Trigger VTP update with any change in SW1

- SW1 and SW2

int range e3/0-1

switchport trunk encapsulation dot1q

switchport mode trunk

- SW1

vlan 100

name test

exit

no vlan 100

Sample Lab Question – LAN Switching

- **Verification steps: SW1**

Check VTP config: domain, version, password and mode

```
SW1#sh vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 6
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 11
VTP Operating Mode              : Server
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Enabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xEC 0x05 0x97 0x80 0x3C 0x74 0xF7 0xA1
Configuration last modified by 0.0.0.0 at 2-24-11 14:44:31
Local updater ID is 0.0.0.0 (no valid interface found)
VTP version running             : 2
SW1#
SW1#sh vtp password
VTP Password: CCIE
SW1#
```

Sample Lab Question – LAN Switching

- **Verification steps: SW2**

Check if VTP config is identical as SW1 (mode can be client)

```
SW2#sh vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 6
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 11
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Enabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xEC 0x05 0x97 0x80 0x3C 0x74 0xF7 0xA1
Configuration last modified by 0.0.0.0 at 2-24-11 14:44:31
Local updater ID is 0.0.0.0 (no valid interface found)
VTP version running             : 2
SW2#
SW2#sh vtp password
VTP Password: CCIE
SW2#
```

Sample Lab Question – LAN Switching

- **Spanning Tree Protocol**

  Configure Spanning-tree protocol on SW1 and SW2 as per the following requirements:

  - Configure SW1 and SW2 so that SW1 is the root for all even* VLANs, and SW2 is the root for all odd* VLANs.
  - SW1 and SW2 should also serve as backup root switches for their respective VLANs.
  - Ensure that both switches are running one STP instance per VLAN

  NOTE:

  –Your solution must affect only the VLANs used in this topology.

  –Odd numbers are not divisible by 2, for example: 1, 3, 5, 7, and so on.

  –Even numbers are divisible by 2, for example: 2, 4, 6, 8, and so on.

Score: 2 Points

Practice Lab: Sample topology

Router-Switch connections

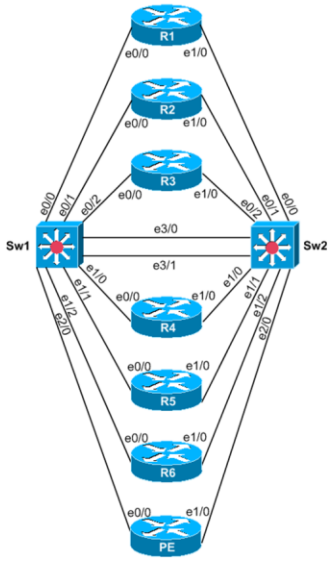

Sample Lab Question – LAN Switching

- **Hints, requirements and constraints in the stem:**
  - Configure SW1 and SW2 so that SW1 is the root for all even* VLANs, and SW2 is the root for all odd* VLANs.
  - SW1 and SW2 should also serve as backup root switches for their respective VLANs.
  - Ensure that both switches are running one STP instance per VLAN
- Very clear requirements about STP config

```
spanning-tree vlan [VLAN LIST] root primary
spanning-tree vlan [VLAN LIST] root secondary
```

- MST not valid solution!

- SW1

spanning-tree vlan 1,55,237 root secondary

spanning-tree vlan 10,44,50,60 root primary

- SW2

spanning-tree vlan 1,55,237 root primary

spanning-tree vlan 10,44,50,60 root secondary

Sample Lab Question – LAN Switching

- **Hints, requirements and constraints in the stem:**

  (…)

  - Your solution must affect only the VLANs used in this topology.

- Check which VLANs are used

  => VLAN 1, 10, 44, 50, 55, 60, 237

  SW1 = root for 10, 44, 50, 60

  SW2 = root for 1, 55, 237

  DO NOT FORGET VLAN 1 !

- By default SW1 will be root for all Vlans, if VLAN 1 is forgotten, it will cost the points for the question.

Sample Lab Question – LAN Switching

- **Verification steps: SW1**

Check current STP Priority for ALL vlans:

```
SW1#sh span s | i Root
Root bridge for: VLAN0010, VLAN0044, VLAN0050, VLAN0060
SW1#
SW1#sh spanning-tree vlan 55 | i Priority
  Root ID     Priority    24631
  Bridge ID   Priority    28727  (priority 28672 sys-id-ext 55)
SW1#sh spanning-tree vlan 44 | i Priority
  Root ID     Priority    24620
  Bridge ID   Priority    24620  (priority 24576 sys-id-ext 44)
SW1#
```

Both must not match for **odd** VLANs and Bridge ID Priority must be <32768!

Both must match for **even** VLANs

TECCCIE-3610 Cisco Public 77

- Checking that the Switch is secondary root means that its Priority must be set lower than the default 32768 but higher than the Root's priority
- That's what the grading script will be looking for, not the configuration

Sample Lab Question – LAN Switching

- **Verification steps: SW2**

Check current STP Priority for ALL vlans:

```
SW2#sh span s | i Root
Root bridge for: VLAN0001, VLAN0055, VLAN0237
SW2#
SW2#sh spanning-tree vlan 55 | i Priority
  Root ID     Priority    24631
  Bridge ID   Priority    24631  (priority 24576 sys-id-ext 55)
SW2#
SW2#sh spanning-tree vlan 44 | i Priority
  Root ID     Priority    24620
  Bridge ID   Priority    28716  (priority 28672 sys-id-ext 44)
SW2#
```

Both must match for **odd** VLANs

Both must not match for **even** VLANs and Bridge ID Priority must be <32768!

- Checking that the Switch is secondary root means that its Priority must be set lower than the default 32768 but higher than the Root's priority

Sample Lab Question – LAN Switching

- **Frequently missed points:**
  - Requirement missed or not fully configured (ex. VLAN 1!)
  - Constraints not respected
  - Main guidelines not respected
  - VTP Configuration revision not matching as expected
  - Access ports not configured in proper VLAN
  - Ports shut down by software (err-disabled)

Errdisable Port State Recovery on the Cisco IOS Platforms

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00806cd87b.shtml

sh interface status > errdisabled

Means that the software detected an error situation on the port and must shut down the port.

There are various reasons for the interface to go into errdisable. The reason can be:

- Duplex mismatch
- Port channel misconfiguration <
- BPDU guard violation <
- UniDirectional Link Detection (UDLD) condition <
- Late-collision detection
- Link-flap detection
- Security violation
- Port Aggregation Protocol (PAgP) flap
- Layer 2 Tunneling Protocol (L2TP) guard
- DHCP snooping rate-limit
- Incorrect GBIC / Small Form-Factor Pluggable (SFP) module or cable
- Address Resolution Protocol (ARP) inspection
- Inline power

- After you fix the root problem, the ports are still disabled if you have not configured errdisable recovery on the switch. In this case, you must reenable the ports manually. Issue the shutdown command, and then the no shutdown interface mode command on the associated interface in order to manually reenable the ports.

References

- Cisco Press®

  Cisco LAN Switching, Kennedy Clark

  Cisco Field Manual: Catalyst Switch Configuration, David Hucaby, Stephen McQuerry

- Interconnections, 2nd edition, Radia Perlman

- CCO Documentation

  Cisco Catalyst 3560 configuration guide CCO

  http://www.cisco.com/en/US/products/hw/switches/ps5528/tsd_products_support_series_home.html

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	IGP Redistribution
2.4	Routing Protocols: EGP (BGP)
2.5	Multicast Routing
2.6	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

IP Routing Concepts

- Classless and Classful
- Static and Dynamic routing
- Policy-based Routing
  Forwarding decisions based on criteria other than destination IP address
- Administrative Distance
  Determining the priority of different sources of routing updates
- Passive Interfaces
  Interfaces that participate in routing to learn but not advertise routes

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	IGP Redistribution
2.4	Routing Protocols: EGP (BGP)
2.5	Multicast Routing
2.6	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

Disclaimer—Reminder

- With the time allocated, we can only review the cornerstones of the most important IGPs

  EIGRP and OSPF

EIGRP Topics

- Neighbour Relationships
  - Packet Types
  - Establishment Process
  - Instability
- DUAL (Route Computation)
  - Reported and Feasible Distance
  - (Feasible) Successor
  - Internal vs External Distance

- EIGRP Neighbors:

  Packet types

  Hello – multicast - Not acknowledged

  Must be in same AS and K-values must match

  Source of hello is primary subnet on interface

  Passive-int stops hellos on specified interfaces

  Neighbor cmd sends hellos as unicast and avoid mcsat hello > required on adjacent neighbors

EIGRP Topics

- Summarisation
  - EIGRP Stub
- Load Balancing
  - Equal Cost
  - Unequal Cost

- http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/eigrpstb.html

Sample Lab Question – EIGRP

- **EIGRP AS 777**

  Configure EIGRP Autonomous System 777 on R2, R3 and SW1 for VLAN 237 as per the following requirements:

  - Do not activate EIGRP on the serial interfaces.
  - Do not use the passive- interface command on any EIGRP router.
  - Disable automatic summarisation for EIGRP.
  - Without using the network command advertise Loopback0 on SW1 into EIGRP.
  - Configure R3 so that it will only accept EIGRP updates from SW1

Practice Lab: Sample topology

Focus on EIGRP AS777

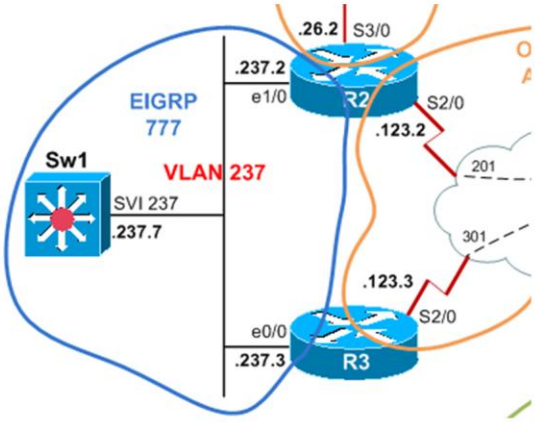

Sample Lab Question – EIGRP

- **Hints, requirements and constraints in the stem:**
  - (…) Disable automatic summarisation for EIGRP.
  - Do not use the passive-interface command on any EIGRP router.
  - Do not activate EIGRP on the serial interfaces. (…)
- These requirements imply that only the VLAN 237 interfaces are running EIGRP

```
router eigrp 777
 no auto-summary
 network 172.16.237.0 0.0.0.255
```


- R2

router eigrp 777

no auto-summary

network 172.16.237.0 0.0.0.255

- R3

router eigrp 777

no auto-summary

network 172.16.237.0 0.0.0.255

- SW1

router eigrp 777

no auto-summary

network 172.16.237.0 0.0.0.255

Sample Lab Question – EIGRP

- **Hints, requirements and constraints in the stem:**

  (…)

  - Without using the network command advertise Loopback0 on SW1 into EIGRP.

- This requirement implies to redistribute the prefix

```
access-list 7 permit 7.7.7.7 0.0.0.0
!
route-map L00 permit 10
 match ip address 7
router eigrp 777
 redistribute connected route-map L00
```

Sample Lab Question – EIGRP

- **Hints, requirements and constraints in the stem:**

  (…)

  - Configure R3 so that it will only accept EIGRP updates from SW1

- This requirement points to using the gateway option:

```
ip prefix-list ALLOWED-NEIGH seq 5 permit 172.16.237.7/32
ip prefix-list ALLOWED-NEIGH seq 10 deny 0.0.0.0/0 le 32
ip prefix-list ALLOWED-PREF seq 5 permit 0.0.0.0/0 le 32
!
router eigrp 777
 distribute-list prefix ALLOWED-PREF gateway ALLOWED-NEIGH in
```


- The filters may be specific or general but need to use the gateway option with the ingress distribute list.

Sample Lab Question – EIGRP

- **Verification steps:**

1. Check EIGRP 777 neighbors and interfaces (No Serial seen)

```
R3#sh ip protocol | s eigrp
  Automatic network summarization is not in effect
  => NO passive interface section seen at all
```

2.

```
R2#sh ip ro 7.7.7.7 | i via
  Known via "eigrp 777", distance 170, metric 409600, type external
  Redistributing via eigrp 777
  * 172.16.237.7, from 172.16.237.7, 00:06:39 ago, via Ethernet1/0
```

3.

```
R3#sh ip protocol | I gateway
 Incoming update filter list for all interfaces is (prefix-list) ALLOWED-PREF gateway ALLOWED-NEIGH
```


- Green code in verif means that the script will have to recursively check for the ‘prefix-list’ or ‘access-list’ named “ALLOWED-PREF” and check for the expected values…

Sample Lab Question – EIGRP

- **Frequently missed points:**
  - Prefix missing
  - Requirement not fully configured (ex. Passive-int)
  - Requirement missing (ex. Filter with Gateway option)
  - Constraints not respected
  - Main guidelines not respected

References

- CiscoPress®
  - Routing TCP/IP Volume I & II, Jeff Doyle
- CCO Documentation
- CCO Whitepapers
  - Introduction to EIGRP
  - Enhanced Interior Gateway Routing Protocol

- Technology handbook

  http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol

- Introduction to EIGRP

  http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800 93f07.shtml

- Enhanced Interior Gateway Routing Protocol

  http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a008 0094cb7.shtml

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	IGP Redistribution
2.4	Routing Protocols: EGP (BGP)
2.5	Multicast Routing
2.6	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

OSPF Topics

- OSPF Areas
  - Area 0
  - Normal
  - NSSA, Stubby, Totally Stubby
- OSPF Router Types
  - ABR, ASBR
- OSPF Link States
  - Router
  - Network
  - Externals (Type 1 and 2)

OSPF Topics

- OSPF over NBMA
  - Point-to-point, point-to-multipoint, broadcast, non-broadcast
- OSPF over broadcast
  - DR & BDR, Wildcard masks
- Virtual link
- OSPF Route Selection
- Summarization
- Authentication
  - Peer and area

Sample Lab Questions – OSPF

…Aggregated…

- **OSPF Areas configuration**

  Configure the OSPF domain as per the IGP topology diagram and the following requirements:

  - In Area 0, use the OSPF point-to-multipoint network type on the serial interfaces of each router.
  - In Area 41, do not change the OSPF network type on R1's S2/0.14 subinterface.
  - Change the OSPF network type on R4's S2/0 interface to point-to-multipoint.
  - Permit OSPF inter area routing information into Area 41
  - Permit external type 7 link state advertisements into Area 41
  - Permit OSPF inter area routing information into Area 10 and ensure that external routing information is not flooded into Area 10
  - R1 must inject a default route into area 10. Do not use the 'default-information-originate' command.

- This sample question is only listing some interesting requirements from a typical set of OSPF questions: in the real lab, all these requirements would be distributed into multiple sub-questions, usually separated per area

Practice Lab: Sample topology

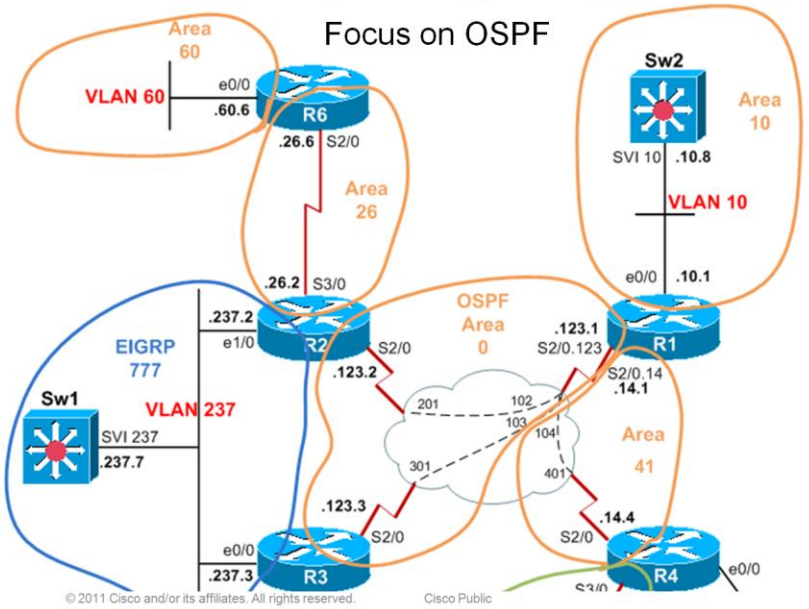

Sample Lab Question – OSPF

- **Hints, requirements and constraints in the stem:**
  - (…) In Area 0, use the OSPF point-to-multipoint network type on the serial interfaces of each router.(…)
- This requirements is very explicit, no tricks

```
interface Serial2/0
 ip ospf network point-to-multipoint
!
router ospf 1
 network 172.16.123.0 0.0.0.255 area 0
```

Sample Lab Question – OSPF

- **Hints, requirements and constraints in the stem:**
  - (…) In Area 41, do not change the OSPF network type on R1's S2/0.14 subinterface.
  - Change the OSPF network type on R4's S2/0 interface to point-to-multipoint. (…)
- These requirements imply to tune the timer (R1 or R4)

```
R1: interface Serial2/0.14 point-to-point
 ip ospf hello-interval 30
                          OR
R4: interface Serial2/0
 ip ospf hello-interval 10
```


- Timer may be tuned in either R1 or R4, the point is to get the OSPF neighborship up while fullfilling the requirement of network-type mismatch
- This is an example of multiple solutions and the verification will not account for any of them, just checking the requirement is respected and the OSPF neighorship up as expected
- Check in either side what the default is "sh ip ospf interface s2/0" and set the hello time identical on the other side…

Sample Lab Question – OSPF

- **Hints, requirements and constraints in the stem:**
  - (…) Permit OSPF inter area routing information into Area 41
  - Permit external type 7 link state advertisements into Area 41
  - Inject a default route into Area 41
- These requirements imply that Area 41 is NSSA

```
router ospf 1
 area 41 nssa default-information-originate
```

Sample Lab Question – OSPF

- **Hints, requirements and constraints in the stem:**
  - (…) Permit OSPF inter area routing information into Area 10 and ensure that external routing information is not flooded into Area 10
  - R1 must inject a default route into area 10. Do not use the 'default-information-originate' command.
- These requirements imply that Area 10 is stub

```
router ospf 1
 area 10 stub
```

Sample Lab Question – OSPF

- **Hidden/Implicit requirements in the stem:**
  - Area 60 is not connected to Area 0

    This implies a Virtual Link via Area 26

```
router ospf 1
 area 26 virtual-link <router-id>
```

  - No explicit requirements about Loopback interfaces nor router-id

    Up to you where to put loopbacks and which router-id to use.

    (Usually, this would be specified in the guidelines or in a full reachability statement.)

Sample Lab Question – OSPF

- **Verification steps:**

1. Check OSPF neighbors, interfaces and area

```
R1#sh ip ospf int bri
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/32         1     LOOP  0/0
Se2/0.123    1     0               172.16.123.1/24    64    P2MP  2/2
Et0/0        1     10              172.16.10.1/24     10    BDR   1/1
Se2/0.14     1     41              172.16.14.1/24     64    P2P   1/1
R1#
```

2. Check Area types

```
R1#sh run | s ospf
router ospf 1
 log-adjacency-changes
 area 10 stub
 area 41 nssa default-information-originate
 network 172.16.10.0 0.0.0.255 area 10
 network 172.16.14.0 0.0.0.255 area 41
 network 172.16.123.0 0.0.0.255 area 0
```


- Checking neighborship, interfaces, state and area can all be done at once, more details can be seen individually per interface with ‘sh ip ospf int br’
- Checking area types would also be done via “sh ip protocol” : ex:

R1#sh ip proto | s ospf

Routing Protocol is "ospf 1"

Outgoing update filter list for all interfaces is not set

Incoming update filter list for all interfaces is not set

Router ID 180.1.1.1

It is an area border and autonomous system boundary router

Redistributing External Routes from,

Number of areas in this router is 3. 1 normal 1 stub

Sample Lab Question – OSPF

- **Verification steps:**

3. Check reachability between Area 0 and Area 60

```
R2#ping 172.16.60.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.60.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/60/60 ms
R2#
```


- Checking reachability implies that the Virtual-link is up and functional, whatever config was used for it
- Checking from Area 0's R2 ensures that no other issue may impact on this question (ie. redistribution between EIGRP/OSPF could eventually alter reachability to Area 60 while the OSPF config is correct, see later).

Sample Lab Questions – OSPF

- **OSPF Route Control**
  - Create interface loopback10 on R1 and assign it the address 180.1.1.1/24.
  - Advertise the loopback10 interface into OSPF.
  - DO NOT use the network command.
  - Ensure this network appears in the routing tables of all OSPF routers except R4 and Sw2.
  - Do not use any filters on any router to complete this question.

Practice Lab: Sample topology

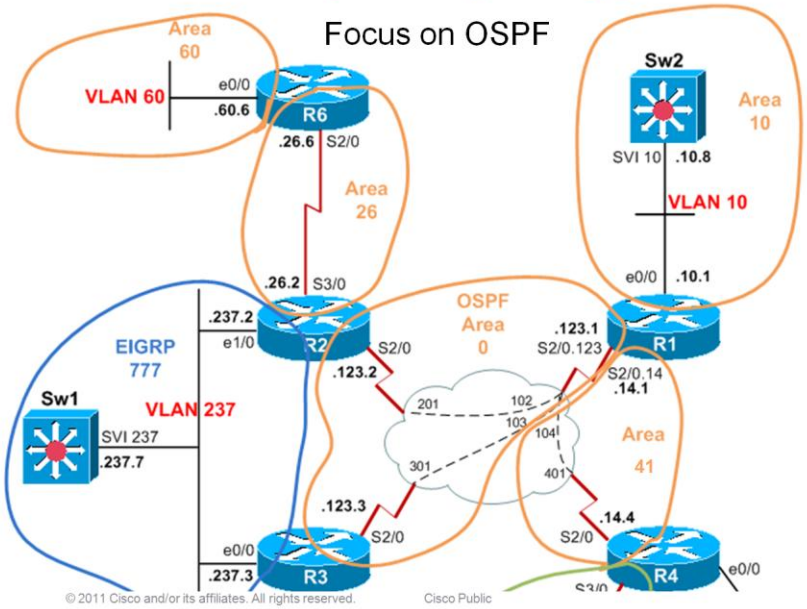

Sample Lab Question – OSPF

- **Hints, requirements and constraints in the stem:**
  - (…) Ensure this network appears in the routing tables of all OSPF routers except R4 and Sw2.(…)
- This requirement implies to use the NSSA no-redistribution:

```
interface Loopback10
 ip address 180.1.1.1 255.255.255.0
access-list 1 permit 180.1.1.0 0.0.0.255
route-map LOOPBACK permit 10
 match ip address 1
router ospf 1
 area 41 nssa no-redistribution default-information-originate
 redistribute connected subnets route-map LOOPBACK
```

Sample Lab Question – OSPF

- **Verification steps:**

1. Check if R1.Lo10 is seen as external in all Areas but 41 and 10:

```
R3#sh ip route 180.1.1.0 | I ospf|via
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 64
  * 172.16.123.1, from 1.1.1.1, 00:05:07 ago, via Serial2/0
R3#
```

```
R6#sh ip route 180.1.1.0 | I ospf|via
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 128
  * 172.16.26.2, from 1.1.1.1, 00:04:11 ago, via Serial2/0
R6#
```

```
R4#sh ip route 180.1.1.0 | I ospf|via
R4#
>> CHECK FOR NO OUTPUT
```

```
SW2#sh ip route 180.1.1.0 | I ospf|via
SW2#
>> CHECK FOR NO OUTPUT
```

Sample Lab Question – OSPF

- **Verification steps:**

2. Check if the nssa no-redistribution was used for area 41

```
R1#sh run | i nssa
 area 41 nssa no-redistribution default-information-originate
R1#
```

Sample Lab Question – OSPF

- **Frequently missed points:**
  - Prefix or interface missing (loopbacks, stub network)
  - Requirement not fully configured
  - Constraints not respected
  - Main guidelines not respected

- Missing a prefix in the routing table may be very expensive as all questions requiring reachability to it will lose points, even if the configuration is fine
- This is very frequent and candidates who make this mistake often complain about their score report because it doesn't match their expectation as all their configurations are functional as per the requirements. They just forgot something that at the end doesn't change much on the functional config.

References

- CiscoPress®
  - Routing TCP/IP Volume I & II, Jeff Doyle
  - Cisco OSPF Command and Configuration Handbook, William R. Parkhurst
  - OSPF Network Design Solutions, Thomas M. Thomas
- CCO Whitepapers
- CCO Documentation

- OSPF Design guide:

  http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml

- Troubleshooting OSPF

  http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800949f7.shtml

- OSPF FAQ

  http://www.cisco.com/en/US/tech/tk365/technologies_q_and_a_item09186a0080094704.shtml

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	IGP Redistribution
2.4	Routing Protocols: EGP (BGP)
2.5	Multicast Routing
2.6	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

IPv6 Addressing Topics

- Addressing
  - Structure
  - Types
    - •Link Local
    - •Global Unicast
    - •Unique Local Unicast
    - •Multicast/Anycast
    - •IPv4 Compatible/Mapped

IPv6 Addressing Topics

- Neighbour Discovery Protocol
  - Router Solicitation/Advertisement
  - Neighbour Solicitation/Advertisment
  - Link layer mapping
  - Duplicate Address Detection
- Address Assignment
  - Static
  - Stateless Autoconfig

IPv6 Routing Topics

- OSPFv3
- EIGRPv6
- Filtering and Route redistribution

OSPFv3 Topics

- Adjacency Formation
  - Per-link vs per-subnet
  - Link-local address
  - Multi-topology via instance-id
- LSA Flooding
  - IPv6 specific multicast addresses
  - New LSA types
  - Renamed LSA types
- Configuration
  - Explicit router-id
  - Per-interface vs per-process
  - IPv6 security replaces OSPF security

- Two New LSA Types

  Link-LSA (LSA Type 0x2008)

  There is one Link-LSA per link. This LSA advertises the router's link-local address, list of all IPv6 prefixes and options associated with the link to all other routers attached to the link

  Intra-Area-Prefix-LSA (LSA Type 0x2009)

  Carries all IPv6 prefix information that in IPv4 is included in Router-LSAs and Network-LSAs

- Two LSAs are Renamed

  Type-3 summary-LSAs, renamed to “Inter-Area-Prefix-LSAs”

  Type-4 summary LSAs, renamed to “Inter-Area-Router-LSAs”

Sample Lab Questions – IPv6

- **IPv6 Addressing**

  Configure IPv6 on R1, R2, R3 and SW1 using the following addressing structure:

  - Assign an IPv6 address to all IPv4 enabled interfaces except additional loopbacks created in the lab. (enable IPv6 on loopback 0 interfaces and omit the additional loopbacks)
  - Use the assigned prefix of 2001:ABC:123::/48 on all interfaces.
  - All subnets are 64 bits.
  - Use EUI-64 to generate the host portion of the address.
  - Complete the network portion of the address from the third octet in the IPv4 address for that interface. For example, use 237 for SW1 E0/0.

Preconfigured

Sample Lab Questions – IPv6

- **OSPFv3**

  Configure OSPFv3 on the frame-relay network (R1, R2, R3) for area 0.

  - Do not enable OSPFv3 on VLAN 237 but ensure it is advertised as an IPv6 prefix to R1 as an external route
  - Ensure that you can IPv6 ping all interfaces and OSPF adjacencies are established.

Preconfigured

- In this case, SW1/VLAN237 will not have any routes to the OSPF domain. Since there is no requirements about this, do not spend time resolving reachability to it.
- Full fill the requirements only
- Again, CCIE is not a test on real-life scenario nor best practices

Sample Lab Question – OSPFv3

- **Frequently missed points:**
  - Typo in IPv6 address
  - EUI Addressing mismatch
  - Requirement not fully configured
  - Constraints not respected
  - Main guidelines not respected

References

- CiscoPress®

  Implementing Cisco IPv6 Networks, Regis Desmeules

  Deploying IPv6 Networks, Popoviciu, Levy-Abengnoli, Grossetete

- CCO Documentation

- www.cisco.com/ipv6
- IPv6 Addressing Architecture (RFC 4291): http://www.ietf.org/rfc/rfc4291.txt
- IPv6 Global Unicast Address Format (RFC 3587): http://www.ietf.org/rfc/rfc3587.txt
- Deprecating Site Local address (RFC 3879): http://www.ietf.org/rfc/rfc3879.txt
- Unique Local IPv6 Unicast Addresses (RFC 4193): http://www.ietf.org/rfc/rfc4193.txt

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	IGP Redistribution
2.4	Routing Protocols: EGP (BGP)
2.5	Multicast Routing
2.6	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

IGP Redistribution topics

- Protocol Metrics
  - Defaults
  - Redistribution
- Filtering
  - Route-maps
  - Tags

- Routes redistributed from one protocol (higher AD) into another protocol (lower AD) will NOT be in the routing table on redistributing router as originated by the second protocol, although AD is lower.
- Know the differences between each protocols' redistributions

  RIP auto-summarization

  EIGRP metric required

  OSPF subnet keyword

  Etc…

Sample Lab Question – IGP Redistribution

- **EIGRP AS 777 and OSPF**
  - Mutually redistribute between the RIP and OSPF domains.
  - Mutually redistribute EIGRP AS 777 into the OSPF backbone area on R2 and R3
  - Filter to avoid routing loops
  - Your solution must ensure that any future external EIGRP prefixes redistributed by SW1 will not require configuration changes in R2 and R3

Practice Lab: Sample topology

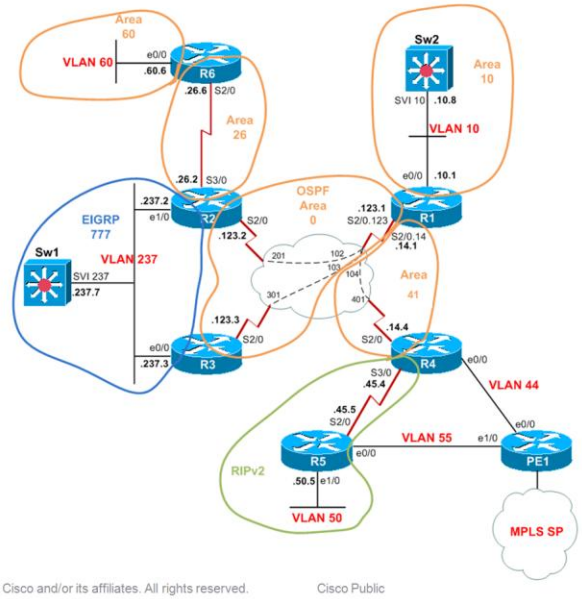

Sample Lab Question – IGP Redistribution

- **Hints, requirements and constraints in the stem:**
  - Mutually redistribute between the RIP and OSPF domains.(…)

- No tricks here, don't forget the 'subnet' and 'metric' keywords and VERIFY reachability!

```
router rip
 redistribute ospf 1 metric 1
router ospf 1
 redistribute rip subnet
```


- The RIP metric can be any number, the script will just check for RIP routes on R5 and ignore the metric.

Sample Lab Question – IGP Redistribution

- **Hints, requirements and constraints in the stem:**
  - (…) **Mutually** redistribute EIGRP AS 777 into the OSPF (…)
  - Filter to **avoid routing loops**
  - Your solution must ensure that any future external EIGRP prefixes redistributed by SW1 will **not** require **configuration changes** in R2 and R3
- Routing loops easily happen with two redistribution points between the same protocols due to route feedback

- Routing loop will likely happen when external routes are being redistributed mutually
- Be careful to not over-configure things to ensure optimal routing but introducing induced routing loops.

Sample Lab Question – IGP Redistribution

- **Route feedback filtering with tags account for any future EIGRP external prefixes:**
  - Mark prefixes per protocol before redistributing them
  - Filter prefixes marked by the other

R2 & R3

```
route-map EIGRP2OSPF deny 10
 match tag 88
route-map EIGRP2OSPF permit 20
 set tag 77
route-map OSPF2EIGRP deny 10
 match tag 77
route-map OSPF2EIGRP permit 20
 set tag 88
!
router eigrp 777
 redistribute ospf 1 metric 1 1 1 1 1 route-map OSPF2EIGRP
router ospf 1
 redistribute eigrp 777 subnet route-map EIGRP2OSPF
```


- Both R2 and R3 are marking prefixes per protocol (set tag 77 for prefixes originated in EIGRP and set tag 88 for prefixes originated into OSPF) then they redistribute them in both sides and both protocol filters the prefixes marked AND redistributed by the other side, which ensures no routing loop may ever happen

- However, this solution may cause some suboptimal routing

  R2 may point SW1.lo via R1 (AD of OSPF 110 is lower than EIGRP external 170)

  If EIGRP ext distance is lowered on R2/R3, it will cause R6.e0/0/VLAN60 to become unreachable due to a routing loop

```
R1#traceroute 172.16.60.6

Type escape sequence to abort.
Tracing the route to 172.16.60.6

 1 172.16.123.2 28 msec 40 msec 40 msec
 2 172.16.237.3 40 msec 40 msec 40 msec
 3 * * *
 4 172.16.123.2 68 msec 84 msec 76 msec
 5 172.16.237.3 76 msec 80 msec 72 msec
 6 * *
R1#
```

- This solution ensures that any new EIGRP external prefixes redistributed by SW1 will be marked, redistributed and filtered in both sides as appropriate

- Ex. http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009487e.shtml#ex2

Sample Lab Question – IGP Redistribution

- **Verification steps:**

1. Check if SW1.lo0 is seen tagged in R1

```
R1#sh ip ro 7.7.7.7 | i via|tag|ext
  Known via "ospf 1", distance 110, metric 20
  Tag 77, type extern 2, forward metric 64
  * 172.16.123.3, from 3.3.3.3, 00:15:07 ago, via Serial2/0.123
      Route tag 77
R1#
```

2. Check for reachability to VLAN60

```
R2#sh ip ro 172.16.60.6 | I via
Known via "ospf 1", distance 110, metric 74, type inter area
  Redistributing via eigrp 777
▌172.16.26.6, from 6.6.6.6, 00:00:27 ago, via Serial3/0
```

1. Check for keywords on R2 and R3

```
R2#sh run | i router |redist
router eigrp 777
 redistribute ospf 1 metric 1 1 1 1 1 route-map OSPF2EIGRP
router ospf 1
 redistribute eigrp 777 subnets route-map EIGRP2OSPF
R2#
```


- Checking reachability to VLAN60 from R1/SW1 ensures that the redistributed prefixes are not being messed up by admin distance tuning in EIGRP domain:

  If EIGRP ext distance is lowered, SW1.lo0 will appear to be pointing to SW1 from both R2 and R3 (AD lower than OSPF AD)

  But in that case, ALL OSPF redistributed prefixes will be preferred

R2#sh ip ro 172.16.60.6 | I via

Known via "ospf 1", distance 110, metric 74, type inter area

Redistributing via eigrp 777

- 172.16.26.6, from 6.6.6.6, 00:00:27 ago, via Serial3/0

With 'distance eigrp 70 109':

R2#sh ip ro 172.16.60.6

Routing entry for 172.16.60.0/24

Known via "ospf 1", distance 110, metric 74, type inter area

Redistributing via eigrp 777

Advertised by eigrp 777 metric 1 1 1 1 1 route-map OSPF2EIGRP

Last update from 172.16.26.6 on Serial3/0, 00:00:27 ago

Routing Descriptor Blocks:

* 172.16.26.6, from 6.6.6.6, 00:00:27 ago, via Serial3/0

Route metric is 74, traffic share count is 1

Sample Lab Question – IGP Redistribution

- **Frequently missed points:**
  - Too much configuration to ensure optimal routing! (ex. SW1 lo0 via R1 from R3)
  - Requirement missing (ex. Route feedback filter, routing loop)
  - Prefix not reachable
  - Constraints not respected
  - Main guidelines not respected

References

- CiscoPress®

  Routing TCP/IP Volume I & II, Jeff Doyle

- CCO Documentation

  Check protocol documentation

- Redistributing routing protocol http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009487e.shtml

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	Routing Protocols: EGP (BGP)
2.4	Multicast Routing
2.5	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

BGP topics

- Peering
  - iBGP
  - eBGP
  - Directly connected vs multi-hop
  - Neighbour Local-AS
- BGP Path Selection
  - IGP Synchronisation
  - Intra-AS attributes
  - Inter-AS attributes
  - Multipath

BGP topics

- BGP Attributes
  - Atomic Aggregate
  - Communities
- Filtering
  - Route-maps
  - Prefix-lists
  - AS Path Lists
  - Regular Expressions
- Scaling
  - Route Reflectors
  - Confederations

Sample Written Exam Question

- Which types of prefixes will a router running BGP advertise to an iBGP peer? Consider that this router is not configured as router reflector. Select the best option

A. Prefixes received from any other BGP peer and prefixes locally originated via network statements or redistributed

B. All prefixes of its routing table

C. Prefixes received from eBGP peers and prefixes locally originated via network statements or redistribute

D. Prefixes received from eBGP peers and prefixes received from Router Reflectors

E. Prefixes received from other iBGP peers, prefixes received from eBGP peers and prefixes redistributed to BGP

F. Prefixes received from other iBGP peers and prefixes received from Router Reflectors

- Option C is the only true statement

Sample Lab Question – BGP

- **IBGP**
  - Configure BGP on R2, R3 and R4 for AS 123.
  - Use a full mesh.
  - Establish all peering sessions using the loopback0 interfaces of each router.
  - Do not configure BGP on R1.
  - On R4 create a Loopback interface and assign it the IP address 99.99.99.1/24
  - Inject the 99.99.99.0/24 prefix into BGP using the network command.
  - Verify the route appears in the routing tables of all routers running BGP.

Preconfigured

- No tricks here, other than the fact the 99.99.99.0 prefix will not be reachable!
- R1 is not running BGP and is blackholing the traffic as it's the FR and OSPF Hub
- Typical example of things that are not best practices nor real life but still requires you to fill the requirement
- Do not spend time making this prefix reachable if not required. Just get it in BGP and RIB tables

Practice Lab: Sample topology

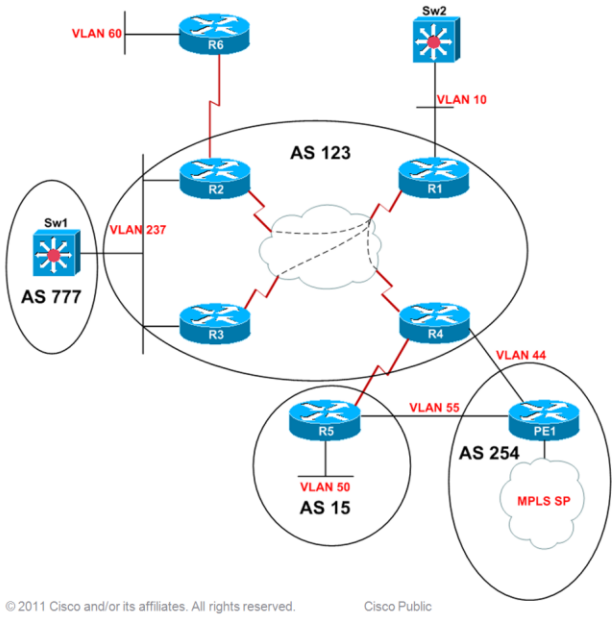

Sample Lab Question – BGP

- **EBGP**

AS 777

- Configure BGP AS 777 on SW1
- Configure SW1 to peer with R2 and R3 in AS 123.

AS 15

- Configure BGP AS 15 on R5.
- Configure R5 to peer with R4 in AS 123
- Advertise the loopback0 interface on R5 via BGP.
- Ensure that network 99.99.99.0/24 is in the BGP and routing tables of SW1 and R5

Preconfigured

- Same example as previously: R5.lo0 will not be reachable for other routers than R4 because it is known as a BGP prefix only and R1 is not running BGP

Practice Lab: Sample topology

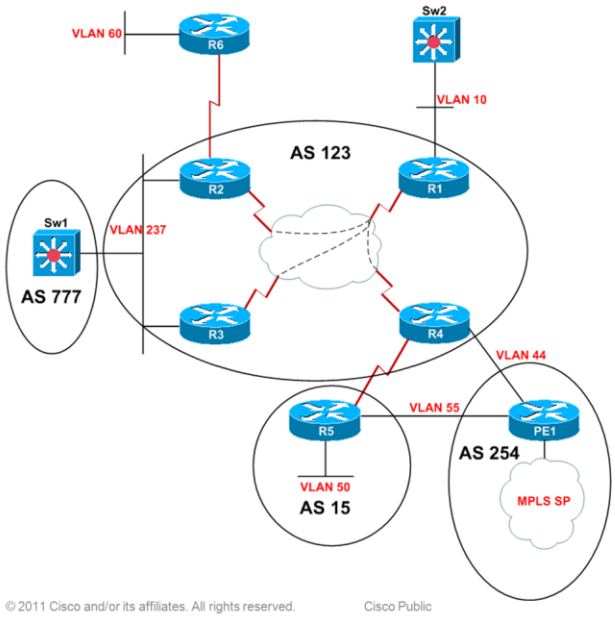

Sample Lab Question – BGP

- **BGP Policy**
  - The administrator of AS 123 wants to influence how traffic from AS 777 enters the network for certain prefixes.
  - Configure R2 such that the 99.99.99.0/24 prefix advertised to AS 777 has the AS path ^123 123 123$.
  - Ensure that other BGP routes are not affected by this policy and will be forwarded to AS 777 unmodified.

Lets focus on this BGP policy question…

Practice Lab: Sample topology

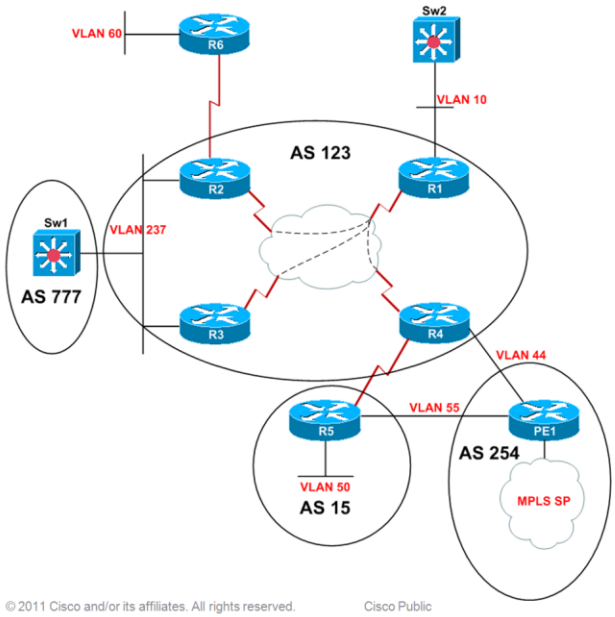

Sample Lab Question – BGP

- **Hints, requirements and constraints in the stem:**
  - Configure R2 such that the 99.99.99.0/24 prefix advertised to AS 777 has the AS path ^123 123 123$.
  - Ensure that **other BGP routes are not affected by this policy** and will be forwarded to AS 777 unmodified.

!! Watch out for hidden bombs !!

Sample Lab Question – BGP

- **Solution seems easy…**

```
access-list 100 permit ip 99.99.99.0 0.0.0.255 any
!
route-map BGP_policy permit 10
 match ip address 100
 set as-path prepend 123 123
route-map BGP_policy permit 20
!
router bgp 123
 neighbor 172.16.237.7 route-map BGP_policy out
end
clear ip bgp *
```


- Note the second permit statement in the route-map, it ensures that other prefixes are still forwarded by R2

Sample Lab Question – BGP

- **Before the policy is applied:**
  - SW1 points both prefixes to R2 (lowest router-id):

```
SW1#sh ip bgp
BGP table version is 21, local router ID is 77.77.77.77
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  5.5.5.5/32       172.16.237.3                           0 123 15 i
*>                  172.16.237.2                           0 123 15 i
*  99.99.99.0/24    172.16.237.3                           0 123 i
*>                  172.16.237.2                           0 123 i
SW1#
```


- Best practice would be to clear ip bgp * in all BGP speakers at the same time, once the iBGP and eBGP sessions are all configured then observe how the BGP table looks like

Sample Lab Question – BGP

- **After the policy is applied:**
  - SW1 **MUST** point 99.99.99.0 to R3 **and** 5.5.5.5 to R2!

```
SW1#sh ip bgp
BGP table version is 30, local router ID is 77.77.77.77
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  5.5.5.5/32       172.16.237.3                         0 123 15 i
*>                  172.16.237.2                         0 123 15 i
*> 99.99.99.0/24    172.16.237.3                         0 123 i
*                   172.16.237.2                         0 123 123 123 i
SW1#
```


- .

Sample Lab Question – BGP

- **IF only R2 BGP is cleared:**
  - SW1 **WILL** point 5.5.5.5 to R3 as the path is the oldest!

```
SW1#sh ip bgp
BGP table version is 31, local router ID is 77.77.77.77
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  5.5.5.5/32       172.16.237.2                           0 123 15 i
*>                  172.16.237.3                           0 123 15 i
*  99.99.99.0/24    172.16.237.2                           0 123 123 123 i
*>                  172.16.237.3                           0 123 i
SW1#
```


- **The algorithm of best path selection will first select the oldest path before the lowest router ID (if paths are external) !!**

- http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml

See step 10 and 11!

10. When both paths are external, prefer the path that was received first (the oldest one).

This step minimizes route-flap because a newer path does not displace an older one, even if the newer path would be the preferred route based on the next decision criteria (Steps 11, 12, and 13).

Skip this step if any of these items is true:

You have enabled the bgp best path compare-routerid command.

Note: Cisco IOS Software Releases 12.0.11S, 12.0.11SC, 12.0.11S3, 12.1.3, 12.1.3AA, 12.1.3.T, and 12.1.3.E introduced this command.

The router ID is the same for multiple paths because the routes were received from the same router.

There is no current best path.

The current best path can be lost when, for example, the neighbor that offers the path goes down.

11. Prefer the route that comes from the BGP router with the lowest router ID.

The router ID is the highest IP address on the router, with preference given to loopback addresses. Also, you can use the bgp router-id command to manually set the router ID.

Note: If a path contains route reflector (RR) attributes, the originator ID is substituted for the router ID in the path selection process.

Sample Lab Question – BGP

- **Solution:**

1. Configure R2
2. Clear BGP ‘simultaneously’ on both R2 and R3!
3. **Verify resulting state!**

- The algorithm of best path selection will first select the oldest path before the lowest router ID (if paths are external) !!

- http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml

  See step 10 and 11!

Sample Lab Question – BGP

- **Verification steps:**

1. Check best path on SW1's BGP table
2. Check R2 and R3 for expected config

```
R3#sh ip bgp nei 172.16.237.7 | i outgoing
R3#

R2#sh ip bgp nei 172.16.237.7 | i outgoing
  Route map for outgoing advertisements is BGP_policy
R2#sh route-map BGP_policy
route-map BGP_policy, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
  Set clauses:
    as-path prepend 123 123
  Policy routing matches: 0 packets, 0 bytes
route-map BGP_policy, permit, sequence 20
  Match clauses:
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
R2#
```


- R3 must not see any outgoing policies for SW1
- R2 must see an outgoing policy, check for the route-map with 2 permit statements, collect the ACL number, check for AS path prepend and check the ACL

Sample Lab Question – BGP

- **Frequently missed points:**
  - Requirement missed (ex. Prefix not affected by policy)
  - Password not in use (session not cleared after config)
  - Constraints not respected
  - Main guidelines not respected

References

- CiscoPress®

  Internet Routing Architectures, Bassam Halabi

  Cisco BGP-4 Command and Configuration Handbook, William Parkhurst

- CCO Documentation

- http://www.cisco.com/en/US/tech/tk365/tk80/tsd_technology_support_sub-protocol_home.html
- http://docwiki.cisco.com/wiki/Border_Gateway_Protocol
- BGP Case studies: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c95bb.shtml

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	Routing Protocols: EGP (BGP)
2.4	Multicast Routing
2.5	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

Multicast topics

- Multicast Forwarding
  - RPF
  - Shared Tree
  - Source Tree
  - IGMP
- Protocol Independent Multicast (PIM)
  - Messages
  - Dense Mode
  - Sparse Mode

- PIM messages:
  - Hello
  - Register
  - Register-Stop
  - Join/Prune
  - Bootstrap
  - Assert
  - Candidate RP advertisement
- PIM Sparse Mode
  - Any Source Multicast (ASM)
  - Source Specific Multicast (SSM)
  - Bidirectional (Bidir)
  - RP Configuration – Static, Auto-RP, BSR, Anycast
  - MSDP

Sample Written Question - Multicast

- Which of the following is NOT true of IP Multicast Addressing?
  1. Multicast Group addresses comprise the range 224.0.0.0–239.255.255.255
  2. The Link-Local Address Range is 224.0.0.0–224.0.0.255
  3. Administratively Scoped Addresses (239.0.0.0– 239.255.255.255) are assigned to user applications by IANA
  4. EIGRP Hello's to 224.0.0.10 have a TTL = 1
  5. Scope Relative Addresses are the top 256 addresses of a scoped address range

- IP Addresses use the Class D address space

  Class D addresses are denoted by the high 4 bits set to 1110.

- Link Local Address Range.

  Range: 224.0.0.0/24

  Used for local wire multicast by protocols such as PIM, OSPF, EIGRP, etc. These multicasts are always sent with a TTL of 1 and never leave the local subnet.

- Global Address Range.

  Range: 224.0.1.0 – 238.255.255.255

  In general, this address space can be thought of as reserved for global multicast to/from the Internet.

- **Option #3 is wrong: Administratively Scoped Address Range.**

  Range: 239.0.0.0/8

  This range is the private multicast address space similar to RFC 1918 unicast address space and should be used for multicast contained entirely within an Enterprise network.

  The Administratively Scoped address range is often subdivided into smaller ranges that define smaller geographical "scopes" within an Enterprise.

- Scope Relative Address Range.

  When the Administratively Scoped address range is subdivided into smaller scope ranges, the top 256 addresses of each scope is reserved for "relative" addressing.

Sample Lab Question – Multicast

- Auto RP
  - Configure PIM Sparse-mode on R1, R2, R3, SW1.
  - Multicast servers are located on VLAN 10.
  - Multicast receivers are located on VLAN 237.
  - Use auto RP and define loopback0 on R3 as the rendezvous point (RP) for all multicast groups.
  - Configure the VLAN10 interface on SW1 to join IGMP group 227.1.1.1
  - Ensure you can ping the group address from all PIM routers and VLAN 10.

Preconfigured

Practice Lab: Sample topology

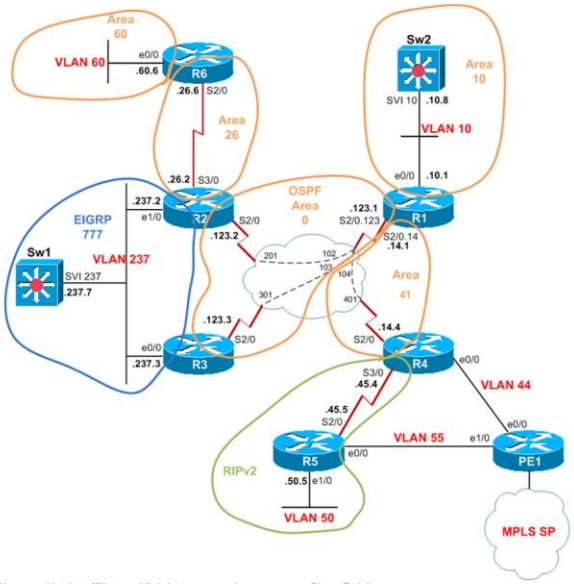

Sample Lab Question – Multicast

- **Hints, requirements and constraints in the stem:**
  - Configure **PIM Sparse-mode** on R1, R2, R3, SW1.
  - Multicast servers are located on VLAN 10, receivers are located on VLAN 237.
  - Use **auto RP** and define loopback0 on **R3 as the rendezvous point (RP)** for all multicast groups.
- PIM Sparse-mode and Auto RP are 'exclusive'…

```
ip multicast-routing
(R3)ip pim send-rp-announce Loopback0 scope 16
(Any)ip pim send-rp-discovery Loopback0 scope 16
ip pim autorp listener
interface X/Y
 ip pim sparse-mode
 (SW1)ip igmp join-group 227.1.1.1
```


- Auto RP requires sparse-dense mode
- If sparse-mode is required, then PIM speakers need to enable 'ip pim autorp listener'
- http://www.cisco.com/en/US/docs/ios/12_2/ipmulti/command/reference/1rfmult2.html#wp1090395

ip pim autorp listener

To cause IP multicast traffic for the two Auto-RP groups 224.0.1.39 and 224.0.1.40 to be Protocol Independent Multicast (PIM) dense mode flooded across interfaces operating in PIM sparse mode, use the ip pim autorp listener command in global configuration mode.

Sample Lab Question – Multicast

- **Verification steps:**

1. Check if R1 can ping the group

```
R1#ping 227.1.1.1 source 172.16.10.1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 227.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.10.1

Reply to request 0 from 172.16.237.7, 52 ms
R1#
```

1. Check if R3 is the RP for all groups

```
SW1#sh ip pim rp map | i Group|Auto-RP
PIM Group-to-RP Mappings
Group(s) 224.0.0.0/4
    Info source: 3.3.3.3 (?), elected via Auto-RP
SW1#
```

2. Check if SW1 has joined the group

Sample Lab Question – Multicast

- **Frequently missed points:**
  - Requirement missing
  - RPF Failure
  - Constraints not respected
  - Main guidelines not respected

References

- CiscoPress®

  Developing IP Multicast Networks, Beau Williamson

  Routing TCP/IP Volume II, Jeff Doyle

- ftp://ftpeng.cisco.com/ipmulticast/training/index.html
- CCO Documentation

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing:
2.1	Routing Concepts
2.2	Routing Protocols: IGP (EIGRP & OSPF)
2.3	Routing Protocols: EGP (BGP)
2.4	Multicast Routing
2.5	MPLS-VPN
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

MPLS/VPN topics

- Network devices
  - P (Provider) routers
  - PE (Provider Edge) routers
  - CE (Customer Edge) routers
- Protocols
  - IGP: core routing protocols: OSPF, EIGRP, IS-IS
  - Label Distribution Protocol (LDP)
  - Resource reservation (RSVP) protocol
- MPLS labels
  - Forwarding Equivalence Class (FEC)
  - MPLS label encapsulation
  - MPLS label stacking

MPLS/VPN topics

- MPLS planes
  - MPLS control planes
  - MPLS forwarding planes
- Layer 3 VPNs
  - VRFs
  - Route Distinguishers
  - Route Targets
  - Multiprotocol BGP
  - PE-CE routing protocols: eBGP, OSPF, EIGRP, RIPv2

- NO L2 VPN in R&S Lab exam

Sample Lab Question – MPLS/VPN

Configure the Provider Edge (PE) router and R4/R5 to connect to the MPLS network and receive VPN routes from the service provider (SP) network.

- LDP
  - Configure MPLS label switching on the PE router interface E2/0 and verify an LDP session is established into the SP core.
- Multi-protocol BGP
  - Configure multi-protocol BGP on the PE router to support MPLS Layer 3 VPNs.
  - The service provider is in AS 254 as is the PE router.
  - Peer with address 11.11.11.11
  - Peering source must be loopback 0

- It is implied that the PE has reachability to the remote address 11.11.11.11
- If nothing is explicitly mentioned about it, candidate will have to figure out how to resolve it.
- One way is to try OSPF with pinging all OSPF routers with 224.0.0.5… this would result in a reply from 192.168.10.1, the neighboring P router ☺
- NB to advertise the local loopback in OSPF as BGP is sourced from it!

```
ip cef

mpls ldp router-id Loopback0

interface Ethernet2/0

mpls ip
```

Practice Lab: Sample topology

Focus on the PE-CE and MPLS

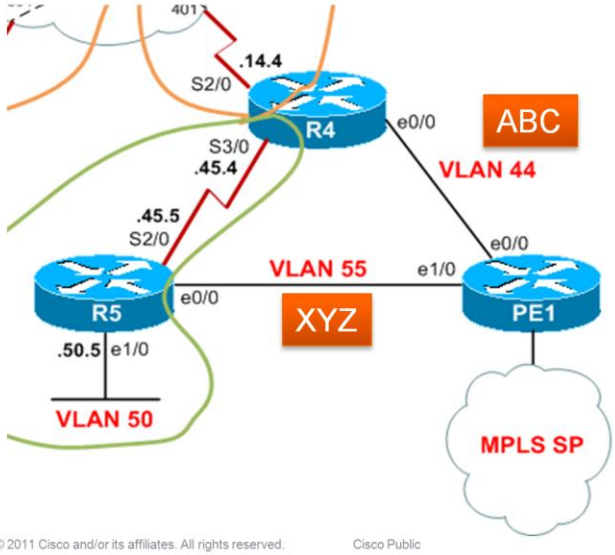

TECCCIE-3610 © 2011 Cisco and/or its affiliates. All rights reserved. Cisco Public 165

Practice Lab: Sample topology

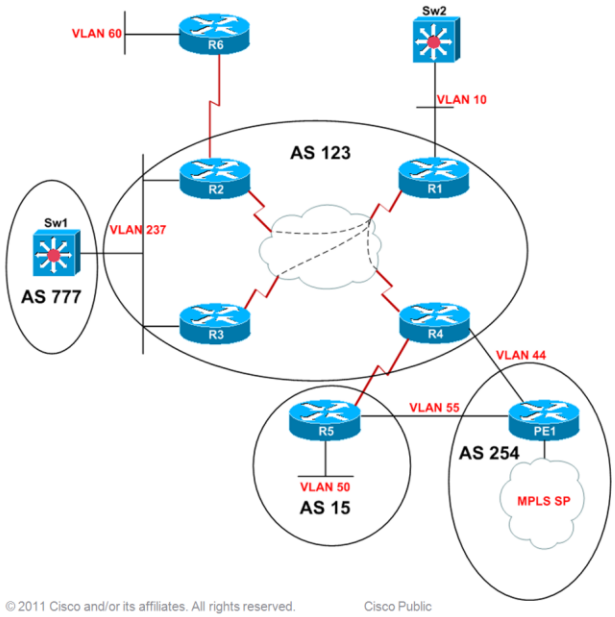

Sample Lab Question – MPLS/VPN

- MPLS Layer 3 VPN
  - Configure two VRFs on the PE router with the following parameters:
  - Create VRF ABC and accept routes with the RD 5.125.16.1:254 into the VPN
  - Create VRF XYZ and accept routes with the RD 254:254 into the VPN.
  - To ensure end-to-end connectivity export using the same route-targets listed above.


```
ip vrf ABC
rd 254:1
route-target export 5.125.16.1:254
route-target import 5.125.16.1:254
!
ip vrf XYZ
rd 254:2
route-target export 254:254
route-target import 254:254
```

Sample Lab Question – MPLS/VPN

- Provider Edge to Customer Edge
  - Advertise the VPN prefixes to R4 and R5 as follows:
  - Insert the link between PE and R4 into the ABC VPN and advertise the ABC VPN routes to R4 via eBGP
  - Insert the link between PE and R5 into the XYZ VPN and advertise the XYZ routes to R5 via eBGP.
  - Configure PE so that VPN routes do not leak between AS 15 and AS 123


```
interface Ethernet0/0
ip vrf forwarding ABC
ip address 192.168.4.254 255.255.255.0
interface Ethernet1/0
ip vrf forwarding XYZ
ip address 192.168.5.254 255.255.255.0
router bgp 254
address-family ipv4 vrf XYZ
neighbor 192.168.5.1 remote-as 15
neighbor 192.168.5.1 activate
no synchronization
exit-address-family
!
```

Sample Lab Question – MPLS/VPN

- **Hints, requirements and constraints in the stem:**
  - LDP
  - MP-BGP
  - MPLS L3VPN
- Not much tricks here…
- CEF is required but IOS reminds it when configuring VRF:

```
% Enable CEF globally before configuring VRF on any interface
```


- This whole part is really vanila MPLS config but requiring to configure only one PE and PE-CE

Sample Lab Question – MPLS/VPN

- **Hints, requirements and constraints in the stem:**
  - PE-CE:
  - **Configure PE** so that VPN routes do **not leak** between AS 15 and AS 123
- If nothing is done: ZYX routes are seen in R4 and vice versa!

```
R4#sh ip ro 10.125.16.0
Routing entry for 10.125.16.0/24
  Known via "bgp 123", distance 20, metric 0
  Tag 15, type external
  Last update from 172.16.45.5 00:46:04 ago
  Routing Descriptor Blocks:
  * 172.16.45.5, from 172.16.45.5, 00:46:04 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 15

R4#
```

```
R5#sh ip ro 10.125.16.0
Routing entry for 10.125.16.0/24
  Known via "bgp 15", distance 20, metric 0
  Tag 254, type external
  Last update from 192.168.5.254 00:43:53 ago
  Routing Descriptor Blocks:
  * 192.168.5.254, from 192.168.5.254, 00:43:53 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 254

R5#
```


PE#sh ip bgp vpn vrf ABC

BGP table version is 41, local router ID is 9.9.9.9

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

r RIB-failure, S Stale

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

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 254:1 (default for vrf ABC)					
*> 5.5.5.5/32	192.168.4.1			0	123 15 i
***>i5.125.16.0/24**	**11.11.11.11**	**0**	**100**	**0**	**i**
*> 99.99.99.0/24	192.168.4.1	0		0	123 i

Sample Lab Question – MPLS/VPN

- If nothing is done: ZYX routes are seen in R4 and vice versa!

```
R4#sh ip bgp
BGP table version is 13, local router ID is 99.99.99.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 5.5.5.5/32       172.16.45.5              0             0 15 i
*> 5.125.16.0/24    192.168.4.254                          0 254 i
*> 10.125.16.0/24   172.16.45.5                            0 15 254 i <<<<<
*> 99.99.99.0/24    0.0.0.0                  0         32768 i
*> 192.168.4.0      0.0.0.0                  0         32768 i
*> 192.168.5.0      172.16.45.5              0             0 15 i
R4#
R4#
```

Sample Lab Question – MPLS/VPN

- Configure PE to prevent eBGP propagation in downstream
- Set community no-export on PE!

```
route-map noexport permit 10
 set community no-export
!
router bgp 254
 address-family ipv4 vrf XYZ
  neighbor 192.168.5.1 route-map noexport out
!
 address-family ipv4 vrf ABC
  neighbor 192.168.4.1 route-map noexport out
```


- Don’t forget to enable ‘send-community’ for both neighbors ☺

Sample Lab Question – MPLS/VPN

- **Verification steps:**

  Check if **legitimate** VPN routes are seen in R4

```
R4#sh ip bgp 5.125.16.0
BGP routing table entry for 5.125.16.0/24, version 16
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to EBGP peer)
  Advertised to update-groups:
        2
  254
    192.168.4.254 from 192.168.4.254 (9.9.9.9)
      Origin IGP, localpref 100, valid, external, best
      Community: no-export
      Extended Community: RT:5.125.16.1:254
R4#
```

Sample Lab Question – MPLS/VPN

- **Verification steps:**

Check if **legitimate** VPN routes are seen in R5

```
R5#sh ip bgp 10.125.16.0
BGP routing table entry for 10.125.16.0/24, version 16
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to EBGP peer)
  Not advertised to any peer
  254
    192.168.5.254 from 192.168.5.254 (9.9.9.9)
      Origin IGP, localpref 100, valid, external, best
      Community: no-export
      Extended Community: RT:254:254
R5#
```

Sample Lab Question – MPLS/VPN

- **Verification steps:**

  Check if **non-legitimate** VPN routes are **not** seen

```
R4#sh ip ro 10.125.16.0
% Network not in table
R4#
R4#sh ip bgp 10.125.16.0
% Network not in table
R4#
```

```
R5#sh ip ro 5.125.16.0
% Subnet not in table
R5#
R5#sh ip bgp 5.125.16.0
% Network not in table
R5#
```

Sample Lab Question – MPLS/VPN

- **Frequently missed points:**
  - Implicit/hidden requirement missed (ex. Shamlink)
  - Constraints not respected
  - Main guidelines not respected

- Shamlinks are used to route traffic via the MPLS backbone when there is a backdoor link between two sites which by default would attract the traffic

References

- CiscoPress®

  MPLS and VPN Architectures – Jim Guichard, Ivan Papelnjak

  Traffic Engineering with MPLS – Eric Osborne, Ajay Simha

  Layer 2 VPN Architectures – Wei Luo, Carlos Pignataro, Dmitry Bokotey, Anthony Chan

  MPLS QoS – Santiago Alvarez

  MPLS Fundamentals – Luc DeGhein

- CCO Documentation:

  http://www.cisco.com/en/US/products/ps6557/products_ios_technology_home.html

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing
Part 3	Advanced Services:
3.1	Quality of Services
3.2	Security
3.3	Services
Part 4	Optimisation
Part 5	Troubleshooting

Quality of Service topics

- QoS Models
  - Diffserv
  - Intserv/RSVP
- QoS Operations
  - Classification/Marking
  - Queuing
  - Policing/Shaping
- Cisco Implementation
  - Modular QoS Command Line (MQC)
  - Catalyst Switch Specifics

Sample Lab Question – Quality of Service

- **Congestion Avoidance**
  - Configure the E0/0 interface on SW1 for congestion avoidance.
  - Ensure SW1 will drop TCP packets prior to periods of congestion.
  - Do not use the modular CLI.
  - Change the minimum queue depth for all IP precedence values to 100.
  - Change the maximum queue depth for all IP precedence values to 250.
  - Use the command show queuing to verify your solution


```
interface Ethernet0/0
random-detect
random-detect precedence 0 100 250
random-detect precedence 1 100 250
random-detect precedence 2 100 250
random-detect precedence 3 100 250
random-detect precedence 4 100 250
random-detect precedence 5 100 250
random-detect precedence 6 100 250
random-detect precedence 7 100 250
```

Sample Lab Question – Quality of Service

- **Frequently missed points:**
  - Policy configured but not applied or wrongly applied
  - Requirement missing
  - Prefix not reachable
  - Constraints not respected
  - Main guidelines not respected

References

- CiscoPress®

  End-to-End QoS Network Design Quality of Service in LANs, WANs, and VPNs, by Tim Szigeti, Christina Hattingh

- CCO Documentation:

  IOS 12.4

  http://www.cisco.com/en/US/products/ps6350/tsd_products_support_series_home.html

  Catalyst 3560

  http://www.cisco.com/en/US/products/hw/switches/ps5528/tsd_products_support_series_home.html

- http://www.cisco.com/en/US/technologies/tk543/tk766/technologies_white_paper09186a00800a3e2f.html

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing
Part 3	Advanced Services:
3.1	Quality of Services
3.2	Security
3.3	Services
Part 4	Optimisation
Part 5	Troubleshooting

Network Security topics

- Access-Control List (ACL)
- Unicast Reverse Path Forwarding (uRPF)
- IP Source Guard
- Authentication, Authorization, and Accounting (AAA)
- Control Plane Policing (CoPP)
- Context-Based Access Control (CBAC)
- Zone Based Firewall

Network Security topics

- Cisco IOS Intrusion Prevention System (IPS)
- Secure Shell (SSH)
- 802.1x
- Routing Protocol Authentication
- LAN Switching security
  - VLAN Access map
  - Private VLAN
  - DAI
  - DHCP snooping

Sample Lab Question – Network Security

- **Access-list**
  - Configure an out-going access list on the S2/0.123 interface of R1 such that:
  - Mail servers on VLAN 237 cannot connect to mail servers on VLAN 10 (SMTP) but mail servers on VLAN 10 can connect to mail servers on VLAN 237
  - Restrict UDP traffic from VLAN 10 such that only hosts using source ports in the range 6000 to 7000 (inclusive) can reach hosts on VLAN 237.
  - Prevent SW1 from successfully pinging the R4"s loopback0 interface. R4 should be able to ping SW1.

- access-list 100 deny tcp 172.16.10.0 0.0.0.255 eq smtp 172.16.237.0 0.0.0.255 eq smtp ack
- access-list 100 permit udp 172.16.10.0 0.0.0.255 range 6000 7000 172.16.237.0 0.0.0.255
- access-list 100 deny udp 172.16.10.0 0.0.0.255 172.16.237.0 0.0.0.255
- access-list 100 deny icmp host 172.16.14.4 host 172.16.237.7 echo-reply
- access-list 100 permit ip any any

- interface Serial2/0.123 multipoint
- ip access-group 100 out

References

- CiscoPress®

  Network Security Technologies and Solutions, Bhaiji

CCO Documentation:

http://www.cisco.com/en/US/tech/tk713/tk237/tsd_technology_support_protocol_home.html

http://www.cisco.com/en/US/docs/ios/wan/configuration/guide/wan_cfg_frm_rly_ps6441_TSD_Products_Configuration_Guide_Chapter.html

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing
Part 3	Advanced Services:
3.1	Quality of Services
3.2	Security
3.3	Services
Part 4	Optimisation
Part 5	Troubleshooting

Network services topics

- Services
  - HSRP
  - GLBP
  - VRRP
  - NTP
  - DHCP
  - WCCP

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing
Part 3	Advanced Services
Part 4	Optimization
Part 5	Troubleshooting

Network optimisation topics

- Network management
  - SNMP
  - EEM
  - (T)FTP, SCP, HTTP(S), Telnet
- Network monitoring
  - Syslog and logging
  - Netflow
  - (R)SPAN
  - IPSLA
  - RMON

“Know what you don’t know”

Section 3: Study plan: Content topics

Part 1	Switching
Part 2	Routing
Part 3	Advanced Services
Part 4	Optimisation
Part 5	Troubleshooting

Network Troubleshooting topics

- Layer 2 network issues
- Layer 3 network issues
- Application response issues
- Network services issues
- Network security issues

Sample Lab Question – Troubleshooting

- TS Exam Topology

See TS Case studies

Large network, allowing numerous independent incidents

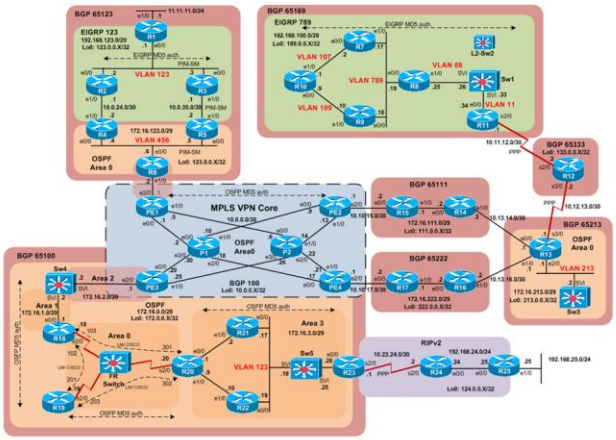

Sample Lab Question – Troubleshooting

- Per-question Topology

  Mini-diagram showing the scope of the incident

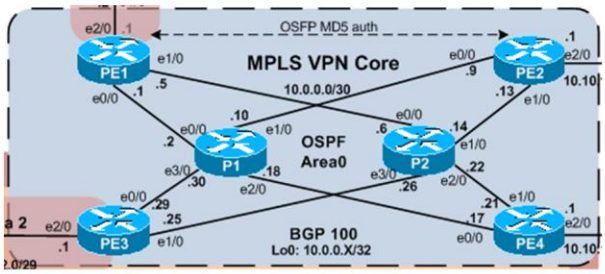

Speeds up finding where to look at in the larger picture

Sample Lab Question – Troubleshooting

- Incident 3:

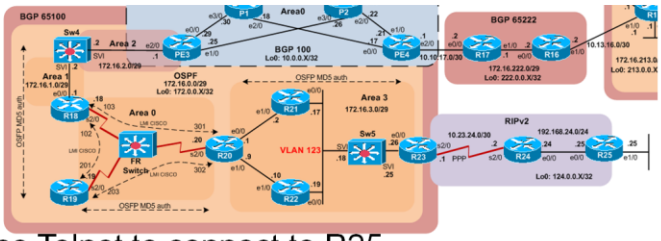

R20 can't use Telnet to connect to R25.

Fix the issue so that the following telnet connection establishes:

R20#telnet 10.1.1.25 /source-int lo0

While resolving the issue, you are not allowed to create any new interfaces.

Refer to the troubleshooting guidelines to determine if your solution is appropriate

Make sure to disconnect the telnet session after verification.

Sample Lab Question – Troubleshooting

- Symptom based question wording

  R20 can't use Telnet to connect to R29.

See TS Case studies

- Explicit validation test

  Fix the issue so that the following telnet connection establishes:

  R20#telnet 10.1.1.29 /source-int lo0

- Optional constraints

  While resolving the issue, you are not allowed to create any new interfaces.

- Pointer to the general guidelines

  Refer to the troubleshooting guidelines…

  Make sure to disconnect the telnet session after verification.

Sample Lab question - Troubleshooting

- Incident 1:

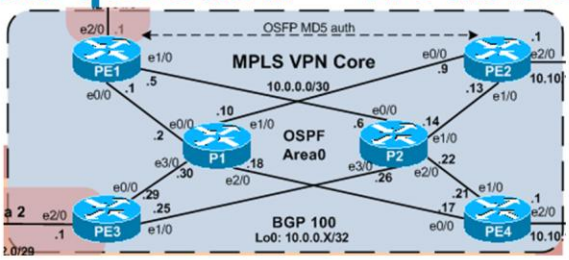

- PE1 is the NTP master for the whole MPLS VPN Core.
- Make sure that all devices are synchronized and authenticating the server.

```
<any MPLS router>#sh ntp assoc det | i auth
10.0.0.111 configured, authenticated, our_master, sane, valid, stratum 1
<any MPLS router>#
```

- Use “CCIE” as the MD5 authentication key (without quotes)
- Refer to the Troubleshooting guidelines to determine if your solution is appropriate.

Sample Lab Question – Troubleshooting

REMINDERS

- CCIE is not testing real-life scenario!
  - Ex. Additional constraints
    - Inserted faults are not always very realistic
- Report any issue to the proctor DURING the TS section! (impossible to grant extra time after its over!)

- Ex. Report slowness with the virtual device during the time it happens and not after the section is completed!
- Use pipe modifier to speed up outputs
- Think before

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

Study plan: Preparation materials

- Cisco 360 Learning Program for CCIE R&S
- Written exam
- Lab exam
- Troubleshooting

Cisco 360 for CCIE R&S

A 360-degree, blended learning program to accelerate expert-level competency and build the skills required for CCIE R&S certification

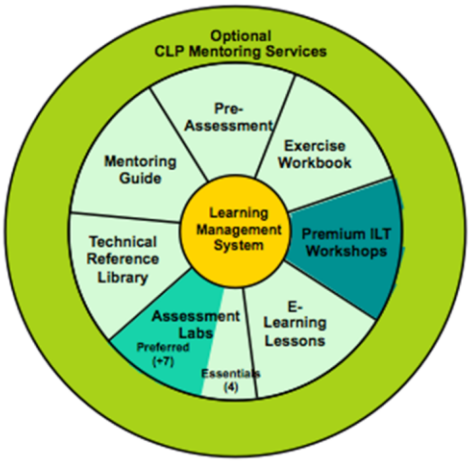

Learning Management System Available to all Cisco 360 students
Essentials Package: ▪Pre-Assessment ▪Exercise Workbook ▪E-Learning Lessons ▪4 Assessment Labs ▪Technical Reference Library ▪Mentoring Guide
Preferred Package: ▪Essentials Package, plus… ▪6 additional Assessment Labs
Premium Package ▪Preferred Package, plus… ▪2 Instructor Workshops
Mentoring Services available from authorised CLPs, to support any package

- Blended learning has proven to be the most effective way of acquiring technical knowledge and practical skills.
- Cisco 360 for CCIE is a flexible, 360-degree learning program with multiple purchase options.
- Training is delivered by specially trained and authorized expert-level instructors.
- In addition to the packages noted here, many of the components are sold individually by Cisco from the Cisco Learning Network. Components purchased online are supported by Cisco Learning Partners.
- Cisco 360 has been updated to reflect the recent changes to CCIE R&S certification through new workshop material, new e-learning lessons on MPLS and troubleshooting, an enhanced exercise workbook, and new graded assessment labs.

Written Exam Preparation

- Study the CCIE R&S Written Exam Topics posted on the Cisco Learning Network (CLN). Some topics-such as 'planning and evaluating network changes-will only appear on the written exam.
- Reading list is only suggested.
- Refer to online resources and Cisco documentation.

Written Exam Preparation

- Use Cisco 360 Learning Program or other training courses to fill holes in your knowledge.
- Written Exam stresses procedures and concepts more than configuration skills.

Lab Exam Preparation, General

- Study the CCIE R&S Lab Exam Topics posted on the Cisco Learning Network (CLN).
- Evaluate your skills against the exam requirements.
- For areas of strength—review and practice for speed. Speed and accuracy is vital on exam.
- For weaker areas—increase knowledge with training or books first, then practice with equipment.

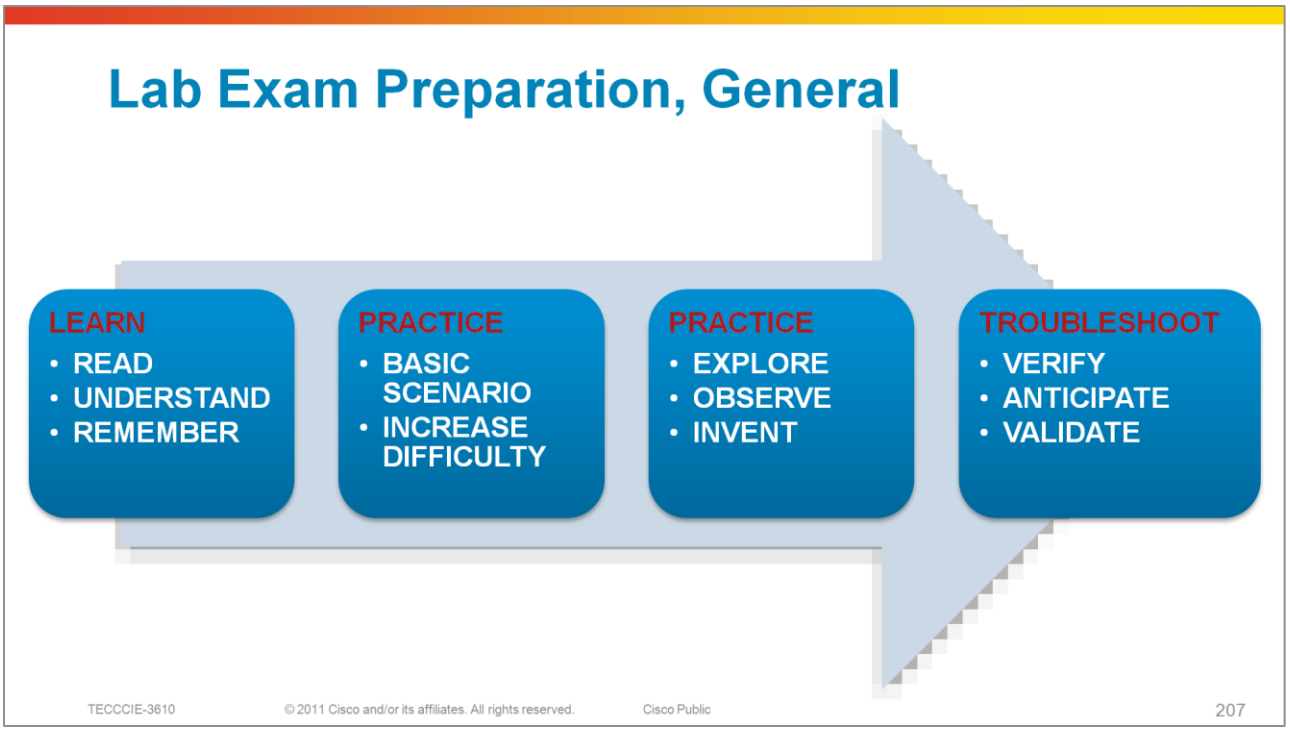

- Build strong foundation knowledge with reading reference materials (books, configuration guide, trainings)
- Practice each technologies individually first then mix then
- Explore the different options in configuration and observe their effects with debugs until you understand every steps
- Invent new scenario using the additional config options and understand their effect
- Use available verification commands, do not rely on running config!
- Anticipate the effect of changes you do and validate the effects with verifications

Troubleshooting Preparation

- Have a solid foundation knowledge first
- Build your own strategy
- Work with experienced peers
- Have them break an unknown topology for you to fix
- Practice for speed !!! Crucial for TS!

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

Lab Exam: D-Day Tips and Tricks

- Build your exam strategy…

…and stick to it!

Lab Exam: D-Day Tips and Tricks

- **Just before the exam:**
  - Reduce stress - arrive early (depends on lab location)
  - Sleep & eat well!
  - Listen to the proctor's briefing (avoid lose time later on)

Lab Exam: D-Day Tips and Tricks

- **When starting exam's section (TS then CFG):**
  - Carefully read the section guidelines!!!
  - Quickly read ALL questions to get an overview of the exam!
  - Prepare a checklist on the scratch paper
  - Manage your time
    - Cherry pick questions on which you can gain time
    - Go back to more challenging questions later on

Lab Exam: D-Day Tips and Tricks

- Make no assumptions
- Work questions as a unit
- Minimize last-minute changes

Lab Exam: Tips and Tricks

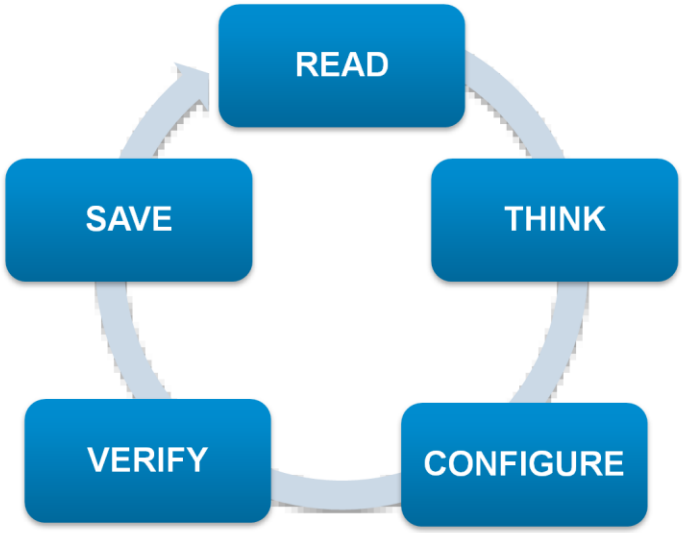

Lab Exam: Tips and Tricks

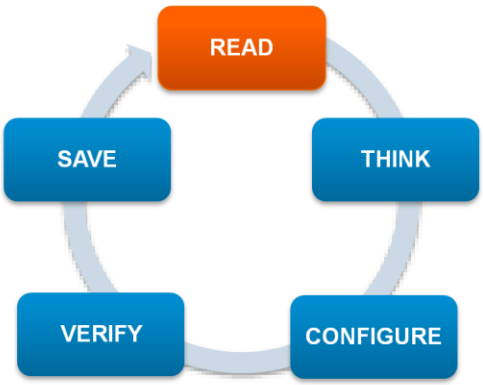

- **READ** the guidelines!
- **READ** the whole exam before starting anything!
- **READ** the question carefully
- **RE-READ** the question
- **ASK** the proctor for clarifications if needed

- Questions often contains many requirements. Its easy to miss one.

Lab Exam: Tips and Tricks

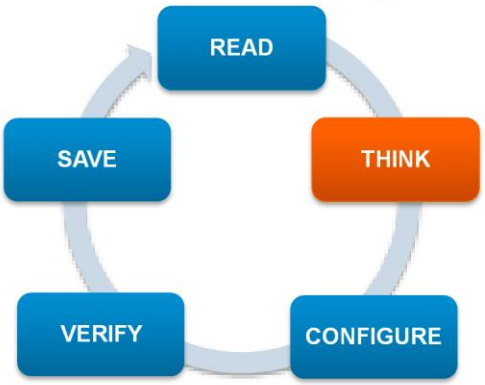

- **THINK** about the expected result
- **THINK** about the solution
- **THINK** about potential impact on previous questions

Lab Exam: Tips and Tricks

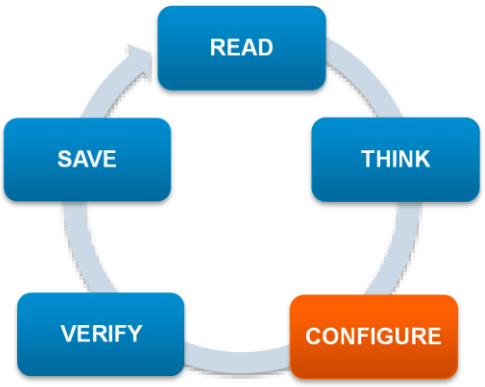

- **Speed up** configuration time
- Use NOTEPAD to build your configurations (copy/paste)
- Use CLI shortcuts and aliases
- Don’t lose time with useless additional configuration
- Organize your terminal windows!

Lab Exam: Tips and Tricks

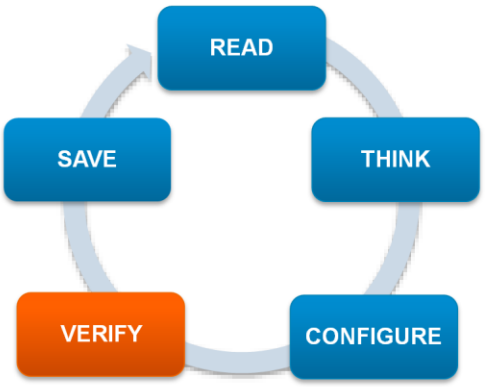

- **VERIFY** the effects of your configuration with ‘show’ commands! (not just ‘sh run’)!
- **VERIFY** effects on previous questions!
- Use **DEBUG** when appropriate
- Enable console logging!

- Especially “debug ip routing” or “debug ip mrouting”

Lab Exam: Tips and Tricks

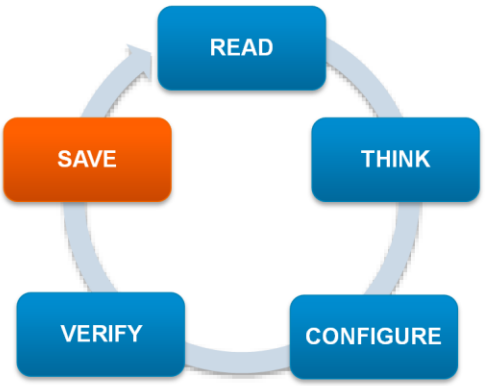

- **SAVE** your configurations at every changes!

Lab Exam: Tips and Tricks

- Once full reachability is established, keep a continuous ping spanning the longest path to easily monitor reachability issue during later questions

- Don’t lose too much time on 2 points questions
  - Move ahead with topics you are more comfortable with
  - Better use a prohibited solution and lose one question rather than losing ALL dependent questions (Ex. PPPoE…)

Lab Exam: Tips and Tricks

- Beware of rumors
- Visit the Cisco Learning Network for more on CCIE R&S certification and interaction with other candidates

  www.cisco.com/go/learningnetwork
- Contact support:

  www.cisco.com/go/certsupport
- Report cheating:

  ccie-nda-enforcement@cisco.com

Lab Exam: Tips and Tricks

- Proctor's first role is to keep exam **fair** !
- Proctors **do not trick** your rack behind your back!

Lab Exam: Tips and Tricks

- Ask the proctor if clarifications are required:
  - Don't ask or fish for answers!
  - Ask good questions:
    - Ex: Ask advises on options based on functional differences!
    - (Show understanding, not just CLI)

Lab Exam: Tips and Tricks

- Report any equipment or technical problems to proctor

  As soon as it occurs but expect he will ask for evidences

  Useless to report/complain after the exam

  (no extra time possible)

  Ex: Console excessively slow or inexistent, Device crash, …

  (very rare)

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

Troubleshooting Case Studies

- Big topology allows for **independent** incidents
  - Incident X+1 should not depend on Incident X's resolution!
- By nature, **timing** is part of the troubleshooting
- Reminders:
  - IOU + L2IOU
  - Guidelines!

Troubleshooting Case Studies

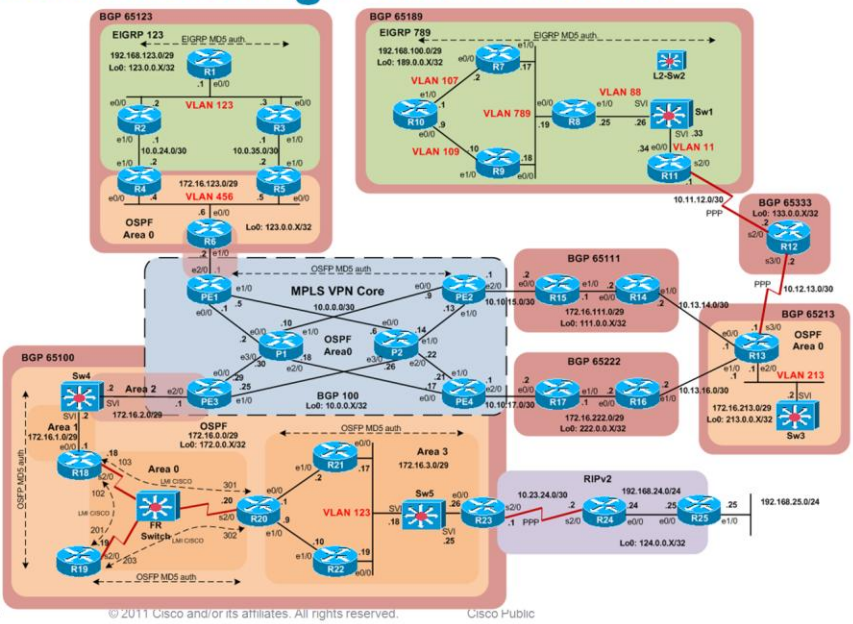

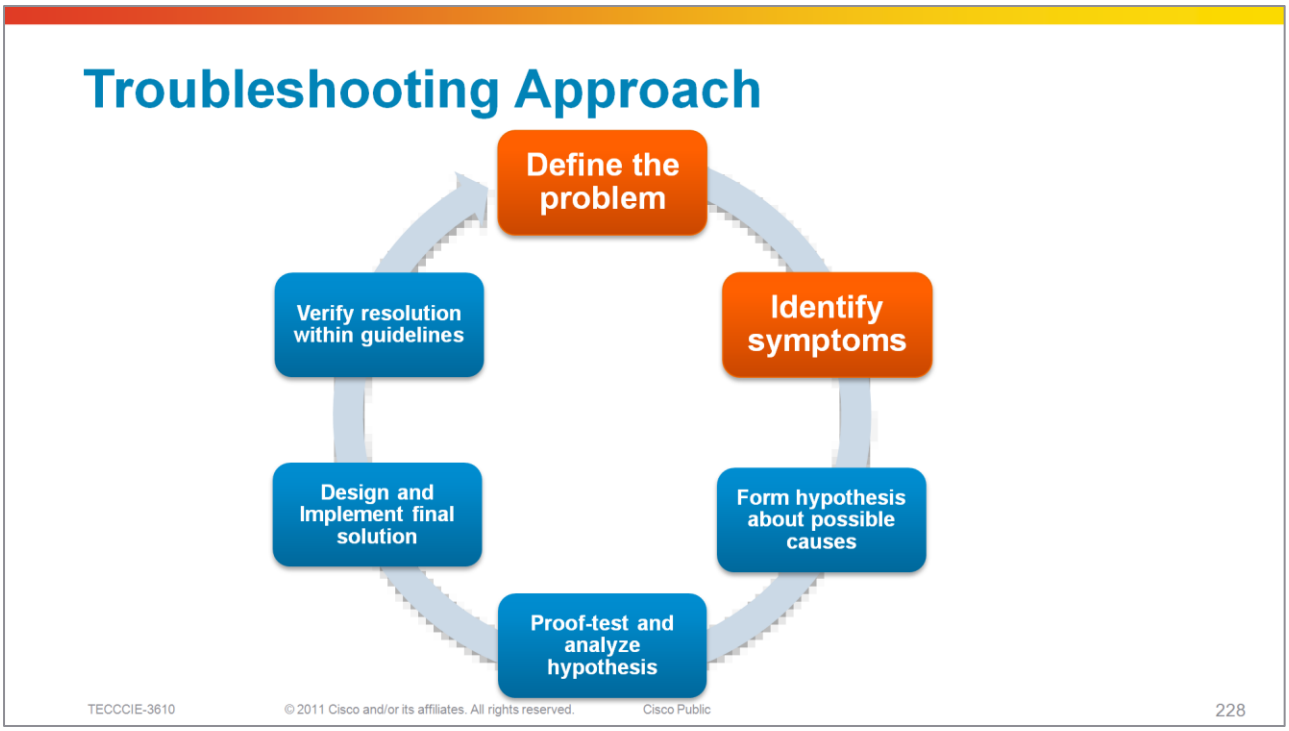

- Define the problem: “Question to the void” until the problem is stated in clear terms of networking components
- Identify the symptoms: This is where your knowledge and experience come into play as the more you understand and know, the easiest you will identify unexpected behavior or outputs
  - These are the most critical and difficult steps
  - Once the problem and symptoms are clarified, most of the time, the associated solution(s) will be obvious
- Most of the time, you will be able to determine multiple possible causes/hypothesis and will need to proof-test each until you find the most likely cause
- Once you found the root cause of the issue, you need to build the solution and verify it fits with the exam guidelines

Incident#1

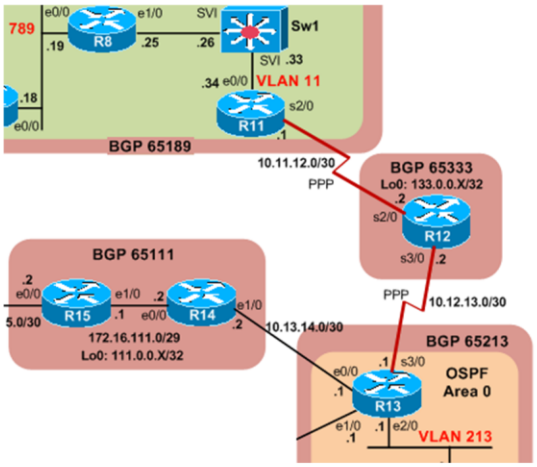

- R15 is not able to use Telnet to connect to R11's loopback0.
- Fix the problem so that the following Telnet establishes:

```
R15#telnet 189.0.0.11 /so lo0
Trying 189.0.0.11 ... Open


User Access Verification

Password:
R11>
```

- While resolving this issue, you are not allowed to create any new interfaces anywhere.

Lets start with an easy one ☺

Incident#1: Define the problem

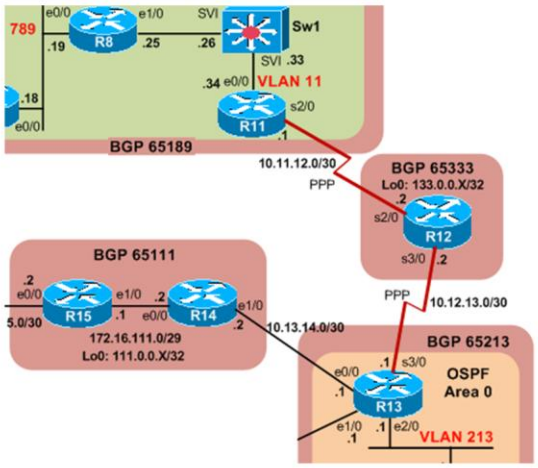

- Can R15 ping R11's lo0?
- Can R12 ping R11's lo0?
- Does R12 have a route to R11's lo0?

- R15:

R15#ping 189.0.0.11

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 189.0.0.11, timeout is 2 seconds:

.....

Success rate is 0 percent (0/5)

R15#

R15#sh ip ro 189.0.0.11

% Network not in table

R15#

- R12

R12#sh ip ro 189.0.0.11

% Network not in table

R12#

Incident#1: Identify symptoms

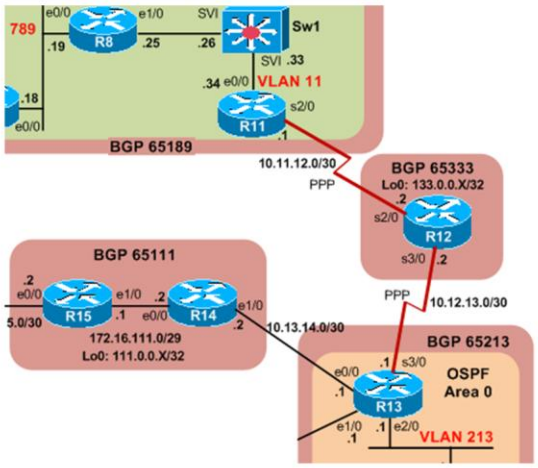

- Is the prefix seen in R12 BGP table?
- Does R11 advertise the prefix in BGP ?

R12#sh ip bgp 189.0.0.11

% Network not in table

R12#

R11#sh ip ro 189.0.0.11

Routing entry for 189.0.0.11/32

Known via "connected", distance 0, metric 0 (connected, via interface)

Redistributing via eigrp 789 **< not into BGP !!!**

Routing Descriptor Blocks:

* directly connected, via Loopback0

Route metric is 0, traffic share count is 1

R11#

R11#sh run | s bgp

router bgp 65189

no synchronization

bgp log-neighbor-changes

network 189.0.0.0 mask 255.255.255.0

aggregate-address 189.0.0.0 255.255.255.0 summary-only

neighbor 10.11.12.2 remote-as 65333

Incident#1: Form hypothesis

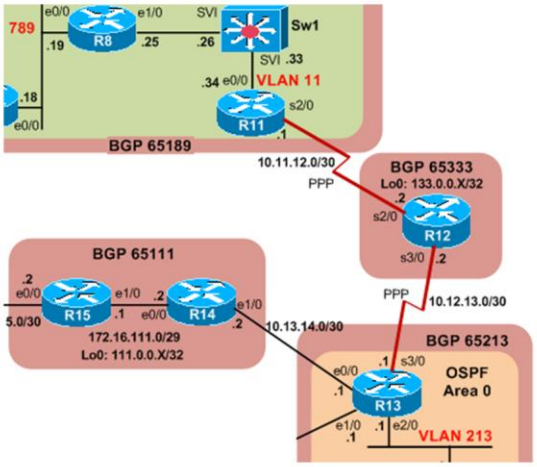

- What is needed for a prefix to be advertised with the network command ?

R11#**sh ip bgp 189.0.0.11**

% Network not in table

R11#

R11#sh run | s bgp

router bgp 65189

no synchronization

bgp log-neighbor-changes

network 189.0.0.0 mask 255.255.255.0

aggregate-address 189.0.0.0 255.255.255.0 summary-only

neighbor 10.11.12.2 remote-as 65333

neighbor 189.0.0.10 remote-as 65189

neighbor 189.0.0.10 update-source Loopback0

neighbor 189.0.0.10 next-hop-self

no auto-summary

R11#

If the mask keyword is configured, then an exact match must exist in the routing table…

Incident#1: Proof-test and Analyze hypothesis

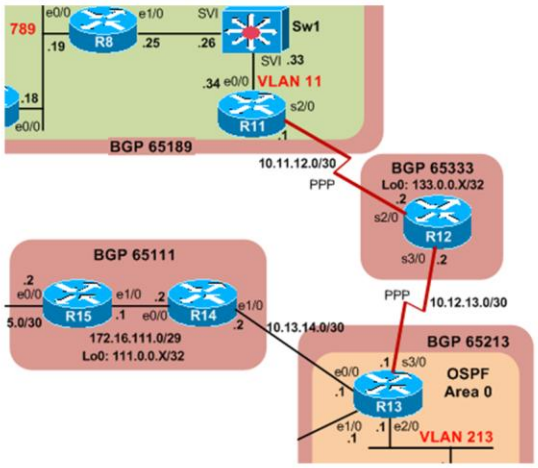

- Correct the mask statement to match the interface's mask…


```
R11#
conf t
router bgp 65189
 no network 189.0.0.0 mask 255.255.255.0
 network 189.0.0.11 mask 255.255.255.255
end

clear ip bgp *

sh ip ro 189.0.0.11
Routing entry for 189.0.0.11/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via eigrp 789
  Advertised by bgp 65189
  Routing Descriptor Blocks:
  * directly connected, via Loopback0
      Route metric is 0, traffic share count is 1

R11#
```

Incident#1: Verify resolution within guidelines

- Check the required Telnet:

```
R15#telnet 189.0.0.11
Trying 189.0.0.11 ... Open


User Access Verification

Password:
R11>
```

- Check if initial BGP configuration is still there:

```
R15#sh ip bgp 189.0.0.11
BGP routing table entry for 189.0.0.0/24, version 202
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        2
  65213 65333 65189, (aggregated by 65189 189.0.0.11)
    10.13.14.1 (metric 20) from 111.0.0.14 (111.0.0.14)
      Origin IGP, metric 0, localpref 100, valid, internal, atomic-aggregate, best
R15#
```

Incident#2

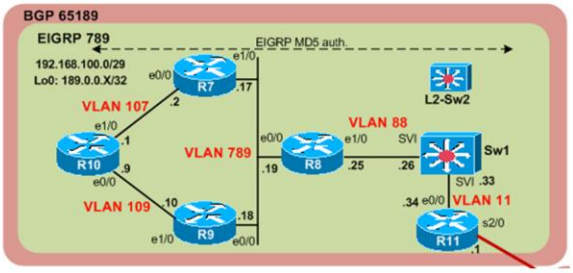

- R11 is not able to use Telnet to connect to R8's loopback0.
- Fix the problem so that the following Telnet establishes:

```
R11#telnet 189.0.0.8
Trying 189.0.0.8 ... Open


User Access Verification

Password:
R8>
```

- While resolving this issue, you are not allowed to create any new interfaces anywhere.

Incident#2: Define the problem

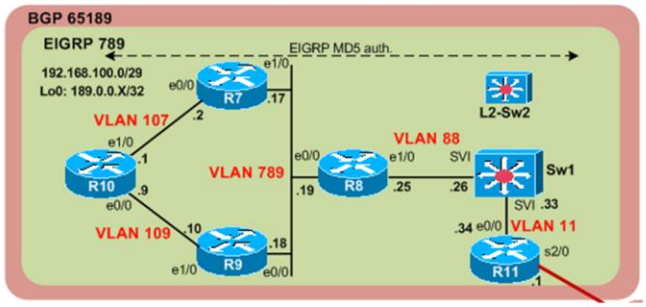

- Can R11 ping R8's e1/0? SW1's SVI 11? SW1's SVI 88?
- Can SW1 ping R8's lo0? e1/0?
- Can SW1 or R9 ping R8's lo0?

- **R11 ping R8.e1/0:**

R11#ping 192.168.100.25

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.100.25, timeout is 2 seconds:

.....

Success rate is 0 percent (0/5)

R11#

- **R11 ping SW1 SVI 11:**

R11#ping 192.168.100.33

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.100.33, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/20 ms

R11#

- **R11 ping SW1 SVI 88**

R11#ping 192.168.100.26

Type escape sequence to abort.

Incident#2: Identify symptoms

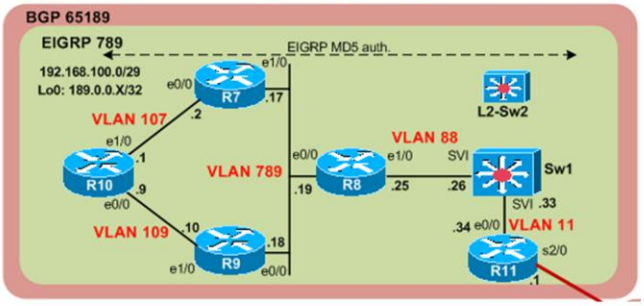

- Are routes present in R11 for VLAN 88 and vice versa?
- Are EIGRP neighborships established?
- Does ARP resolve in VLAN 11 and 88?
- Are interfaces operational?

Top Down approach…

SW1#sh ip ei nei vl 88

EIGRP-IPv4 neighbors for process 789

SW1#

SW1#sh int vl 88 | i line

Vlan88 is down, line protocol is down

SW1#

Incident#2: Form hypothesis

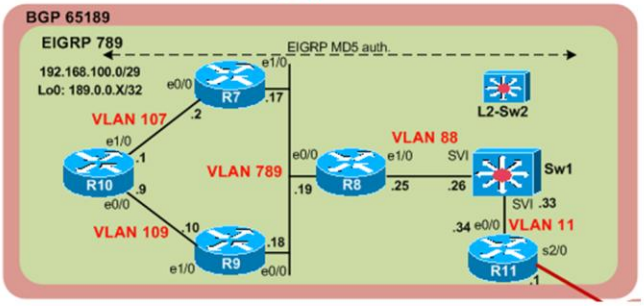

- What can cause interface down/down ?

Incident#2: Proof-test and Analyze hypothesis

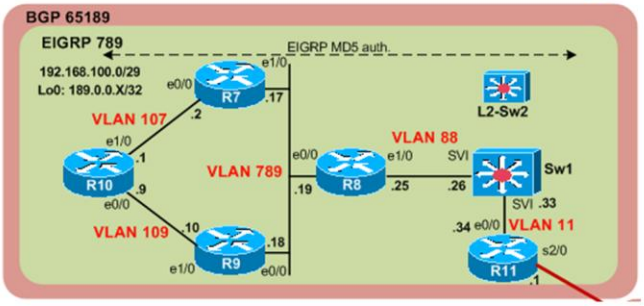

- Hardware or cabling issue?
- Mis-configuration?
- Underlying protocol? (Dot1q, VTP, DTP, STP)

Incident#2: Proof-test and Analyze hypothesis

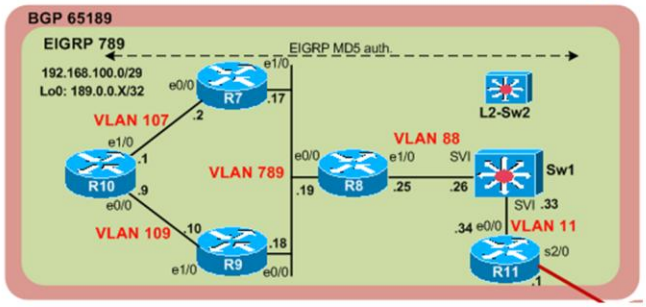

- Underlying protocols:
  - Check VLAN configuration
  - Check Trunk sanity
  - Check STP state

- VLAN config is ok
- Trunk has manual pruning but doesn't look to cause an issue at first
- STP for VLAN 88 is running and blocking on e2/1 only!!!

SW1#**sh span mst 2 | i Et2/**

port Et2/0 cost 2000000 rem hops 19

Et2/0 **Root FWD** 2000000 128.65 Shr

Et2/1 **Altn BLK** 2000000 128.66 Shr

SW1# ^^^^^^^^^^^^^^^

On Sw2:

L2-SW2#sh span mst 2 | i Et2/

Et2/0 **Desg FWD** 2000000 128.65 Shr

Et2/1 **Desg FWD** 2000000 128.66 Shr

L2-SW2#

BUT VLAN 88 is not allowed on SW1.e2/0! And SW1.e2/1 is still blocking due to MST not converging PER-VLAN but PER-INSTANCE!!!

SW1#sh span vl 88 | i Et2

Et2/1 Altn BLK 2000000 128.66 Shr

SW1#

Incident#2: Design and Implement solution

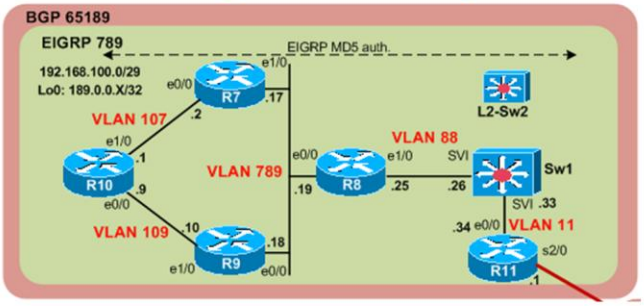

- Force SW1.e2/1 to become root port for MST#2

  or

- Allow VLAN on trunk

Both valid solutions… pick one ☺

- Add VLAN 88 in the allowed list on SW1.e2/0

Or

- Select SW1.e2/1 as root port for MST instance 2 (ex: SW1.e2/1 port cost 1 or Sw2.e2/1 port-priority 0)

Incident#2: Verify resolution within guidelines

- Check the required Telnet:

```
R11#telnet 189.0.0.8
Trying 189.0.0.8 ... Open


User Access Verification

Password:
R8>
```

- Check if initial MST configuration is still there:

```
SW1#sh span mst 1 | i MST|Root
##### MST1    vlans mapped:   1,11,107
Root          this switch for MST1
SW1#
```

```
L2-SW2#sh span mst 2 | i MST|Root
##### MST2    vlans mapped:   88,109,789
Root          this switch for MST2
L2-SW2#
```

Incident#3

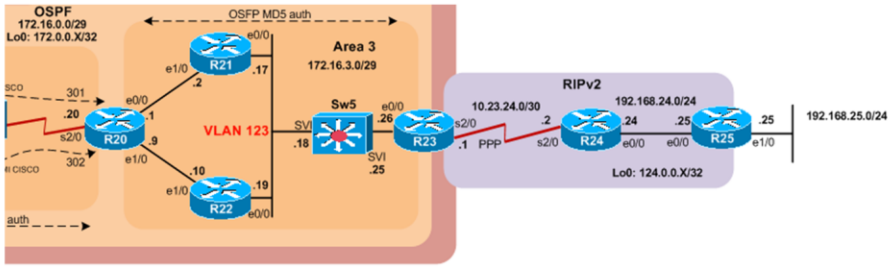

- R20 is not able to use Telnet to connect to R25's loopback0.
- Fix the problem so that the following Telnet and Traceroute establish:

```
R20#telnet 124.0.0.25 /source lo0
Trying 124.0.0.25 ... Open


User Access Verification

Password:
R25>
```

```
R20#traceroute 124.0.0.25

Type escape sequence to abort.
Tracing the route to 124.0.0.25

  1 172.16.3.10 0 msec
    172.16.3.2 4 msec
    172.16.3.10 0 msec
  2 172.16.3.18 0 msec 0 msec 4 msec
  3 172.16.3.26 0 msec 4 msec 0 msec
  4 10.23.24.2 16 msec 20 msec 20 msec
  5 192.168.24.25 20 msec *  20 msec
R20#
```

- While resolving this issue, you are not allowed to change anything in R23

R20#telnet 124.0.0.25 /so lo0

Trying 124.0.0.25 ...

% Destination unreachable; gateway or host down

R20#

R20#traceroute 124.0.0.25

Type escape sequence to abort.

Tracing the route to 124.0.0.25

1 * * *

2

R20#

Incident#3: Define the problem

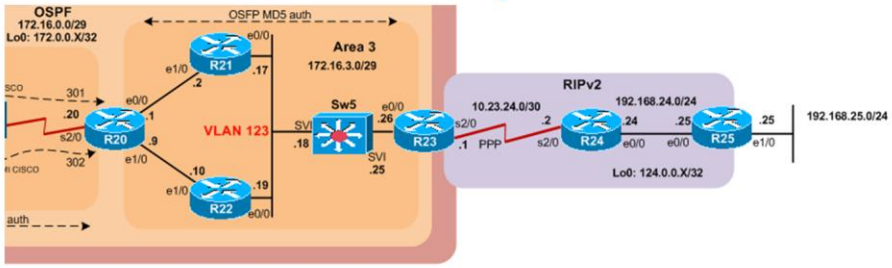

- Can R20 ping R25's e0/0? R23's s2/0?
- Can R23 ping R25's lo0? R25's e0/0?
- Can R21 ping R25's lo0? R25's e0/0?

- R20

R20#ping 124.0.0.25

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 124.0.0.25, timeout is 2 seconds:

.....

Success rate is 0 percent (0/5)

R20#ping 192.168.24.25

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.24.25, timeout is 2 seconds:

.....

Success rate is 0 percent (0/5)

R20#ping 10.23.24.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.23.24.1, timeout is 2 seconds:

!!!!!

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

R20#

Incident#3: Identify symptoms

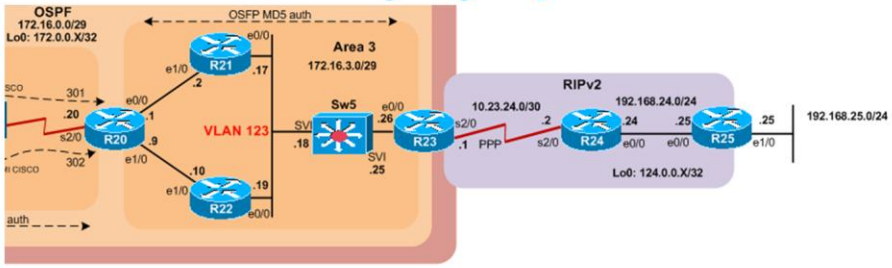

- Are RIP routes present in R20?
- Are OSPF routes present in R25?
- Are RIP routes seen anywhere in OSPF domain?
- ?

R20: RIP routes seen in OSPF domain:

R20#sh ip ro 124.0.0.25

% Network not in table

R20#sh ip ro 192.168.24.25

% Network not in table

R20#sh ip ro 10.23.24.1

Routing entry for 10.23.24.0/30

Known via "ospf 100", distance 110, metric 85, type intra area

Last update from 172.16.3.10 on Ethernet1/0, 00:51:43 ago

Routing Descriptor Blocks:

*** 172.16.3.10, from 172.0.0.23, 00:51:43 ago, via Ethernet1/0**

Route metric is 85, traffic share count is 1

172.16.3.2, from 172.0.0.23, 00:51:43 ago, via Ethernet0/0

Route metric is 85, traffic share count is 1

R20#

R25: OSPF routes seen in RIP domain:

R25#sh ip ro 172.0.0.155

Routing entry for 172.0.0.155/32

Known via "rip", distance 120, metric 15

Incident#3: Form hypothesis

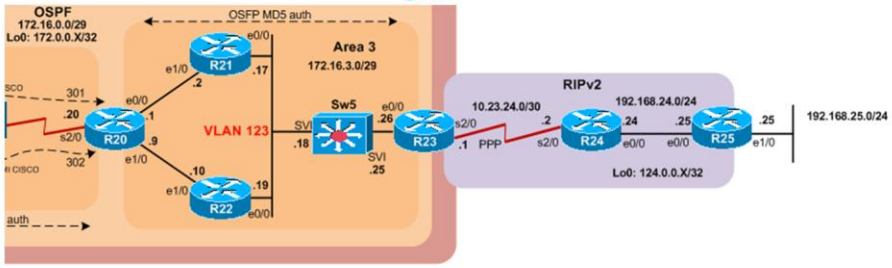

- What can cause a prefix to be stuck in OSPF DB?

- Why Are Some OSPF Routes in the Database but Not in the Routing Table?

  http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009481a.shtml

  - Reason 1: Network Type Mismatch
  - Reason 2: Wrong Address Assignment in Dual Serial Link Setup
  - Reason 3: One Side of Point-to-Point Link Included in Wrong Majornet or Subnet
  - Reason 4: One Side Is Unnumbered and the Other Side Is Numbered
  - Reason 5: Broken PVC in Fully Meshed Frame Relay Environment
  - Reason 6: Forwarding Address Known via an External Route
  - Reason 7: Distribute List Is Blocking the Routes

- Common Routing Problem with OSPF Forwarding Address

  http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009405a.shtml

The value of the forwarding address specified by the autonomous system boundary router (ASBR) (Router 1 in the previous figure) can be either 0.0.0.0 or non-zero. The 0.0.0.0 address indicates that the originating router (the ASBR) is the next hop. The forwarding address is determined by these conditions:

The forwarding address is set to 0.0.0.0 if the ASBR redistributes routes and OSPF is not enabled on the next hop interface for those routes. This is true in the figure if Router 1 does not have OSPF enabled on the Ethernet interface.

These conditions set the forwarding address field to a non-zero address:

OSPF is enabled on the ASBR's next hop interface AND

ASBR's next hop interface is non-passive under OSPF AND

ASBR's next hop interface is not point-to-point AND

ASBR's next hop interface is not point-to-multipoint AND

Incident#3: Proof-test and Analyze hypothesis

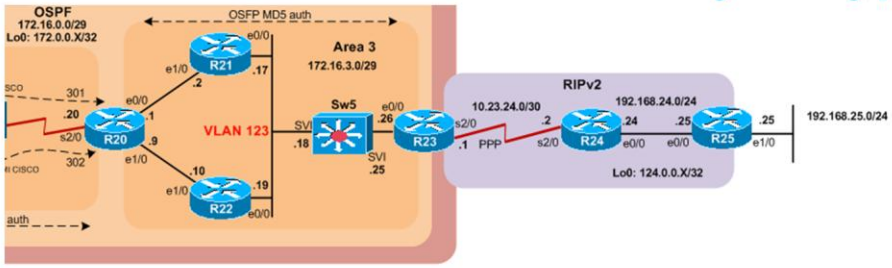

- Is FA set for the Type5 LSA ?
- How is the FA seen by OSPF ?
- Why ?

- R20:

R20#sh ip ospf data ext 124.0.0.25

OSPF Router with ID (172.0.0.20) (Process ID 100)

Type-5 AS External Link States

<<<<<< Routing bit not set !!

LS age: 205

Options: (No TOS-capability, DC)

LS Type: AS External Link

Link State ID: 124.0.0.25 (External Network Number)

Advertising Router: 172.0.0.23

LS Seq Number: 80000004

Checksum: 0xD038

Length: 36

Network Mask: /32

Metric Type: 2 (Larger than any link state path)

TOS: 0

Metric: 20

Forward Address: 10.23.24.2

External Route Tag: 0

Incident#3: Design and Implement solution

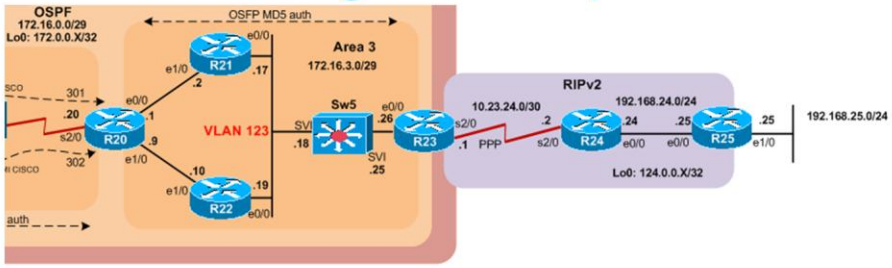

- Resolve the Type5 LSA known via another Type5 How?
- Are the Type5 LSA now seen in R20's Routing table?

- **Remove 'redistribute static subnet' from R22 THEN CHECK R20.....**

R22#sh run | s router

router ospf 100

router-id 172.0.0.22

log-adjacency-changes

redistribute static subnets

network 0.0.0.0 255.255.255.255 area 3

R22#conf t

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

R22(config)#router ospf 100

R22(config-router)#no redist stat

R22(config-router)#do sh run | s router

router ospf 100

router-id 172.0.0.22

log-adjacency-changes

network 0.0.0.0 255.255.255.255 area 3

R22(config-router)#^Z

R22#

CHECK R20 DIRECTLY:

Incident#3: Verify resolution within guidelines

- Check the required Telnet and Traceroute:

```
R20#telnet 124.0.0.25 /source lo0
Trying 124.0.0.25 ...
% Destination unreachable; gateway or
host down
R20#
R20#telnet 124.0.0.25 /source e0/0
Trying 124.0.0.25 ... Open


User Access Verification

Password:
R25>
```

```
R20#traceroute 124.0.0.25 source lo0

Type escape sequence to abort.
Tracing the route to 124.0.0.25

  1 172.16.3.10 4 msec
    172.16.3.2 0 msec
    172.16.3.10 0 msec
  2 172.16.3.18 4 msec
    172.16.3.10 !H
    172.16.3.18 0 msec
R20#
```

- Partial success is not acceptable!

- Telnet eventually completes thanks to the load-balancing between R21 and R22
- Traceroute is stuck at the 2nd hop due to !H (Host unreachable) from R22:

Incident#3: Form hypothesis

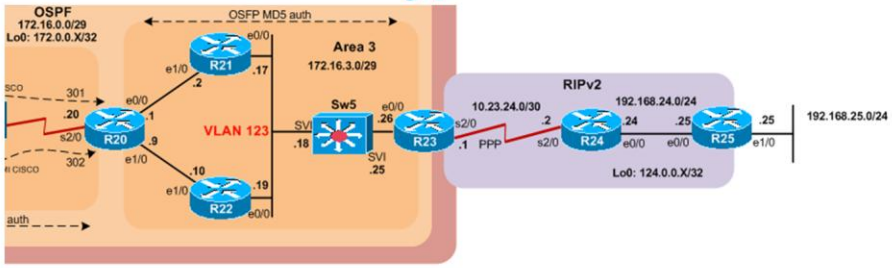

- What can cause intermittent telnet and failed traceroute?

- Intermittent ping/telnet might be caused by QoS or similar policies (rate-limit, MQC, CoPP, ACL?...)
- What else?
- Load-balancing…

Incident#3: Proof-test and Analyze hypothesis

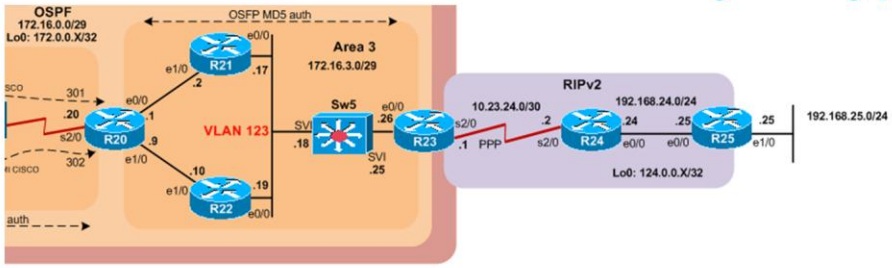

- Are there any Qos or CoPP?
- How is load-balancing achieved?

- **Check Load-balancing method:**

R20#sh cef int e0/0 | i load

Per packet load-sharing is disabled

R20#sh cef int e1/0 | i load

Per packet load-sharing is disabled

R20#sh ip cef 124.0.0.25 | i sharing

124.0.0.25/32, version 42, epoch 0, **per-destination sharing**

R20#

R20#sh ip cef exact-route 172.0.0.20 124.0.0.25

172.0.0.20 -> 124.0.0.25 : **Ethernet1/0 (next hop 172.16.3.10)**

R20#

- **Check route in both R21 and R22:**

R21#**sh ip ro 124.0.0.25**

Routing entry for 124.0.0.25/32

Known via "ospf 100", distance 110, metric 20, type extern 2, forward metric 75

Last update from 172.16.3.18 on Ethernet0/0, 00:21:05 ago

Routing Descriptor Blocks:

* 172.16.3.18, from 172.0.0.23, 00:21:05 ago, via Ethernet0/0

Route metric is 20, traffic share count is 1

Incident#3: Design and Implement solution

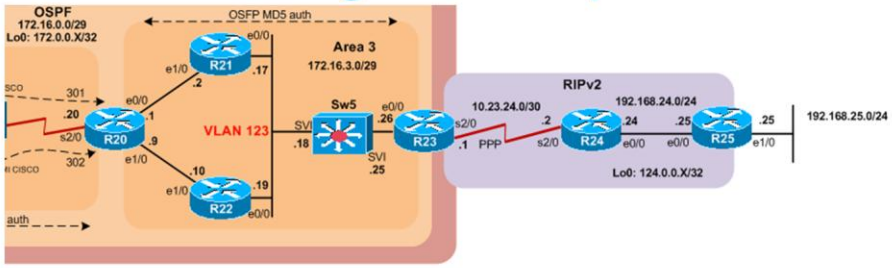

- Resolve the Type5 LSA known via another Type5 in R22 How?
- Are the Type5 LSA seen in R22's Routing table?

- **Remove the static route from R22**

R22#conf t

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

R22(config)#

R22(config)#no ip route 10.23.24.2 255.255.255.255 172.16.3.18

R22#clear ip ro *

R22#

R22#sh ip ospf data ext 124.0.0.25

OSPF Router with ID (172.0.0.22) (Process ID 100)

Type-5 AS External Link States

Routing Bit Set on this LSA

LS age: 1213

Options: (No TOS-capability, DC)

LS Type: AS External Link

Link State ID: 124.0.0.25 (External Network Number)

Advertising Router: 172.0.0.23

LS Seq Number: 80000007

Incident#3: Verify resolution within guidelines

- Check the required Telnet and Traceroute:

```
R20#telnet 124.0.0.25 /source lo0
Trying 124.0.0.25 ... Open


User Access Verification

Password:
R25>
```

```
R20#traceroute 124.0.0.25 source lo0

Type escape sequence to abort.
Tracing the route to 124.0.0.25

  1 172.16.3.10 4 msec
    172.16.3.2 0 msec
    172.16.3.10 0 msec
  2 172.16.3.18 0 msec 4 msec 0 msec
  3 172.16.3.26 4 msec 0 msec 0 msec
  4 10.23.24.2 24 msec 20 msec 20 msec
  5 192.168.24.25 20 msec *  20 msec
R20#
```

- Incident is resolved! ☺

Incident#4

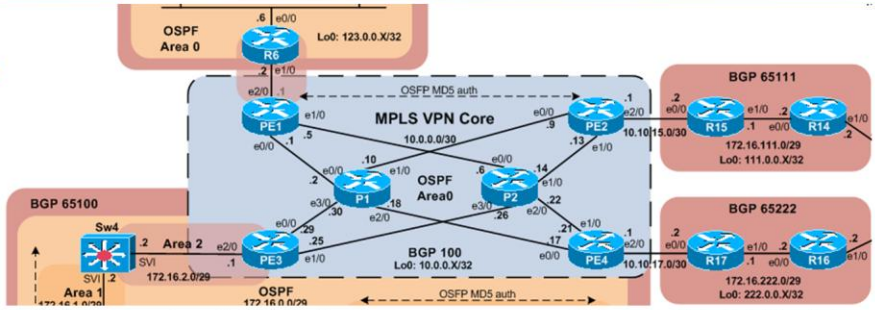

- In VPN-B, host 6.6.6.6 attached to R6 is not able to reach host 17.17.17.17 attached to R17.
- Fix the problem so that the following Ping results in 100 percents:

```
R6#ping 17.17.17.17 so lo1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 17.17.17.17, timeout is 2 seconds:
Packet sent with a source address of 6.6.6.6
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
R6#
```

TECCCIE-3610 © 2011 Cisco and/or its affiliates. All rights reserved. Cisco Public 254

R6#ping 17.17.17.17 so lo1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 17.17.17.17,
timeout is 2 seconds:

Packet sent with a source address of 6.6.6.6

.....

Success rate is 0 percent (0/5)

R6#

Incident#4: Define the problem

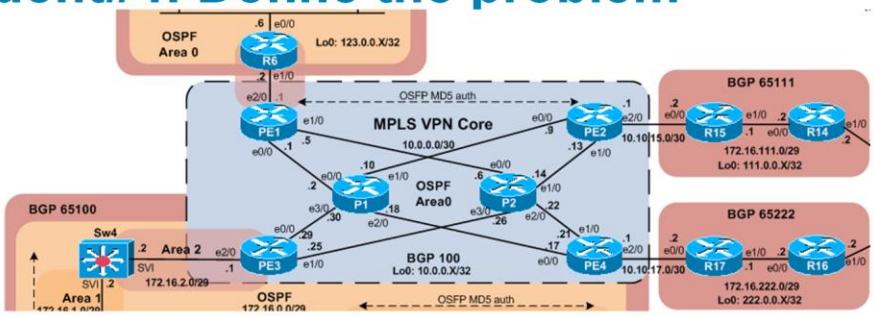

- Is there a route in R6 to R17?
- Is there a route in R17 to R6?
- Do the pings arrive at R17?

R6 has a route to R17:

R6#sh ip ro 17.17.17.17

Routing entry for 17.17.17.17/32

Known via "bgp 65123", distance 20, metric 0

Tag 100, type external

Last update from 172.16.6.1 00:03:33 ago

Routing Descriptor Blocks:

*** 172.16.6.1,** from 172.16.6.1, 00:03:33 ago

Route metric is 0, traffic share count is 1

AS Hops 2

Route tag 100

R6#

R17#sh ip ro 6.6.6.6

Routing entry for 6.6.6.6/32

Known via "bgp 65222", distance 20, metric 0

Tag 100, type external

Last update from 10.10.17.1 00:00:05 ago

Routing Descriptor Blocks:

Incident#4: Identify symptoms

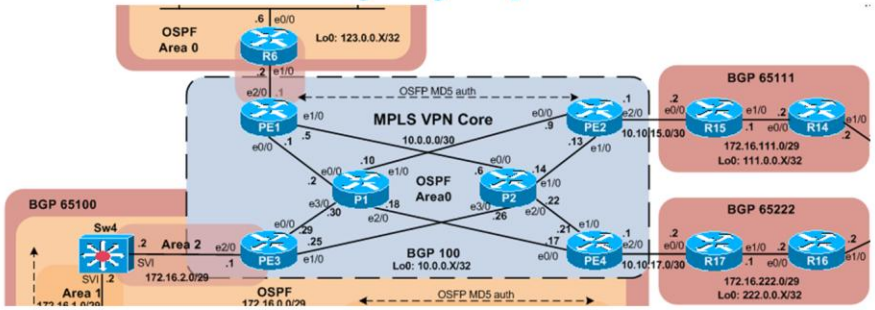

- How traffic is being forwarded inside the MPLS cloud?
- Is the LSP complete from PE1 to PE4?

PE1 looks good:

PE1#sh ip ro vrf VPN-B | i 17.17.17.17

B 17.17.17.17 [200/0] via **10.0.0.114**, 00:20:41

PE1#

PE1#**sh ip ro 10.0.0.114**

Routing entry for 10.0.0.114/32

Known via "ospf 1", distance 110, metric 21, type intra area

Last update from 10.0.0.2 on Ethernet0/0, 11:33:13 ago

Routing Descriptor Blocks:

*** 10.0.0.6, from 10.0.0.114, 11:33:23 ago, via Ethernet1/0**

Route metric is 21, traffic share count is 1

10.0.0.2, from 10.0.0.114, 11:33:13 ago, via Ethernet0/0

Route metric is 21, traffic share count is 1

PE1#sh mpls fo 10.0.0.114

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
23	**18**	10.0.0.114/32	0	Et1/0	10.0.0.6
	17	10.0.0.114/32	0	Et0/0	10.0.0.2

PE1#

Incident#4: Form hypothesis

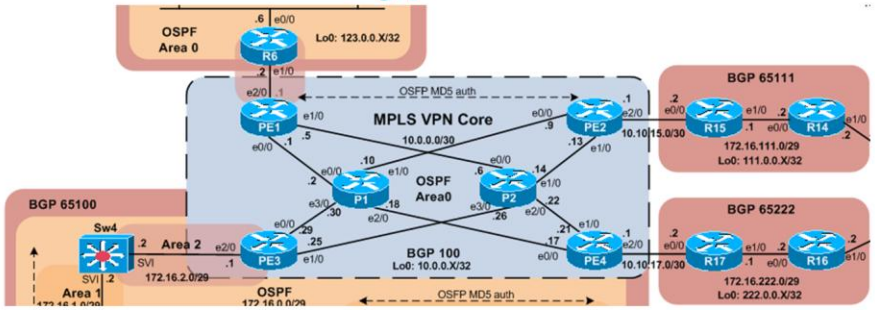

- What can prevent labels to be assigned to prefixes?

PE4#sh mpls ldp neighbor

PE4#

PE4#sh mpls int

Interface	IP	Tunnel	Operational
Ethernet0/0	Yes (tdp)	No	**Yes**
Ethernet1/0	Yes (tdp)	No	**Yes**

PE4#

P1#sh mpls int

Interface	IP	Tunnel	Operational
Ethernet0/0	Yes (ldp)	No	Yes
Ethernet1/0	Yes (ldp)	No	Yes
Ethernet3/0	Yes (ldp)	No	Yes

P1#

P2#sh mpls int

Interface	IP	Tunnel	Operational
Ethernet0/0	Yes (ldp)	No	Yes
Ethernet1/0	Yes (ldp)	No	Yes
Ethernet3/0	Yes (ldp)	No	Yes

P2#

Incident#4: Proof-test and Analyze hypothesis

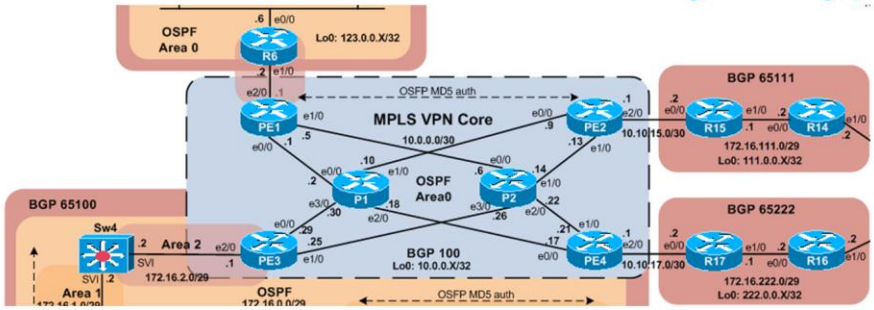

- How are LDP neighborships established ?

LDP is TCP#646, established between loopback interfaces

mpls label protocol ldp

mpls ldp router-id Loopback0

Interface X/Y

Mpls ip

P1#sh run | s ldp|mpls

mpls label protocol ldp

mpls ip

mpls ip

mpls ip

mpls ldp router-id Loopback0

P1#

>> P1 must have 4 “mpls ip” statement, et2/0 is missing it

>> same for P2

Incident#4: Design and Implement solution

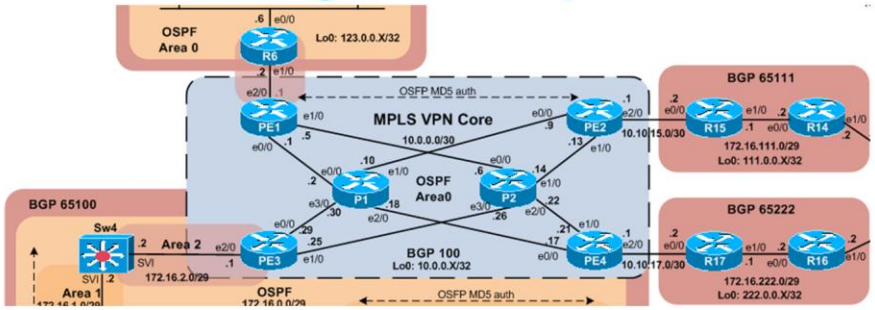

- Bring LDP up between PE4 and P1/P2

P1/P2:

```
int e2/0

mpls ip
```

PE4: has TDP enabled

```
mpls label protocol ldp
```

Incident#4: Verify resolution within guidelines

- Check the required Ping:

```
R6#ping 17.17.17.17 so lo1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 17.17.17.17, timeout is 2
seconds:
Packet sent with a source address of 6.6.6.6
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4
ms
R6#
```

- Incident is resolved! ☺


```
PE4#sh mpls ldp nei

    Peer LDP Ident: 10.0.0.100:0; Local LDP Ident 10.0.0.114:0
            TCP connection: 10.0.0.100.646 - 10.0.0.114.61707
            State: Oper; Msgs sent/rcvd: 17/17; Downstream
            Up time: 00:00:23
            LDP discovery sources:
              Ethernet0/0, Src IP addr: 10.0.0.18
          Addresses bound to peer LDP Ident:
            10.0.0.2        10.0.0.100      10.0.0.10       10.0.0.18
            10.0.0.30
    Peer LDP Ident: 10.0.0.200:0; Local LDP Ident 10.0.0.114:0
            TCP connection: 10.0.0.200.17556 - 10.0.0.114.646
            State: Oper; Msgs sent/rcvd: 17/17; Downstream
            Up time: 00:00:12
            LDP discovery sources:
              Ethernet1/0, Src IP addr: 10.0.0.22
          Addresses bound to peer LDP Ident:
            10.0.0.6        10.0.0.200      10.0.0.14       10.0.0.22
            10.0.0.26
PE4#
PE4#sh mpls fo 10.0.0.111
Local  Outgoing    Prefix          Bytes tag  Outgoing    Next Hop
```

Incident#5

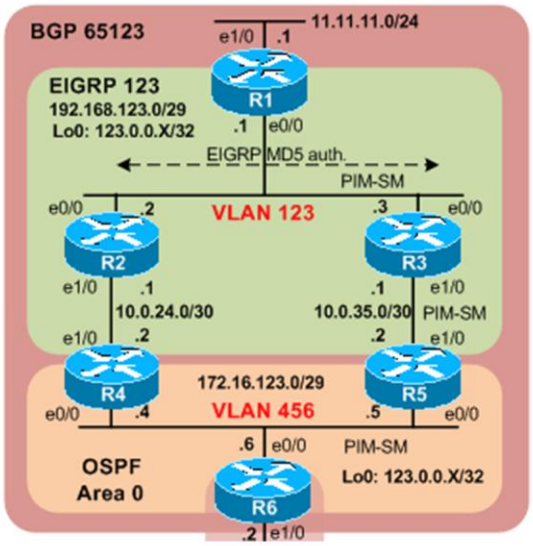

- Clients located behind R1 e1/0 must be able to receive Multicast traffic for the group 231.1.1.1 from a server located in VLAN 456.
- Fix the problem so that the following ping receives replies:

```
R6#ping 231.1.1.1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 231.1.1.1, timeout is 2
seconds:

Reply to request 0 from 192.168.123.1, 20 ms
Reply to request 0 from 192.168.123.1, 20 ms
R6#
```

- This incident contains two separate faults

- Remember: as per the guidelines, a single “fault” may involve one or multiple command lines on the same or multiple devices.

R6#ping 231.1.1.1 rep 5

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 231.1.1.1, timeout is 2 seconds:

.....

R6#

Incident#5: Define the problem

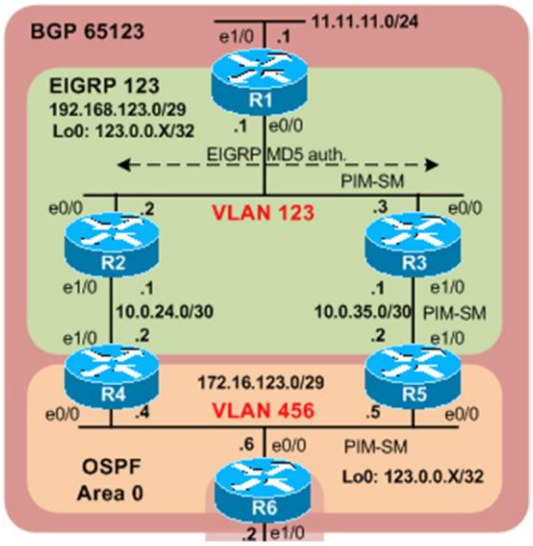

- Does ping fail consistently?
- Did R1 join the required group?

Incident#5: Identify symptoms

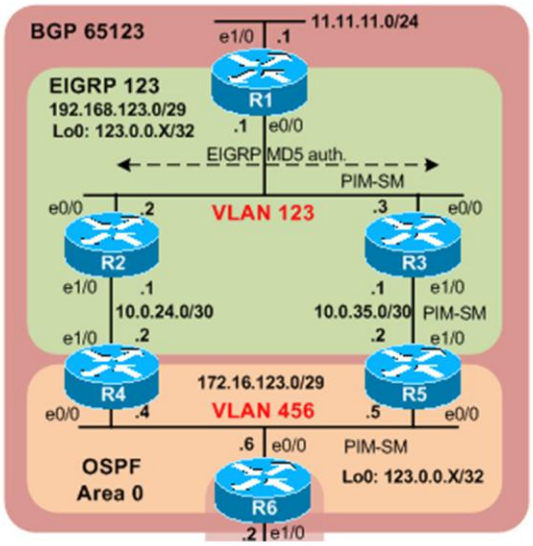

- Are all PIM neighborships established?
- Are the RP announcements propagated in OSPF?
- Are the RP announcements propagated in EIGRP?

PIM neighborships are all fine

RP is known in OSPF but not in EIGRP

R2#sh ip pim rp map

PIM Group-to-RP Mappings

R2#

R3#sh ip pim rp map

PIM Group-to-RP Mappings

R3#

R4#sh ip pim rp map

PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4

RP 123.0.0.6 (?), v2v1

Incident#5: Form hypothesis

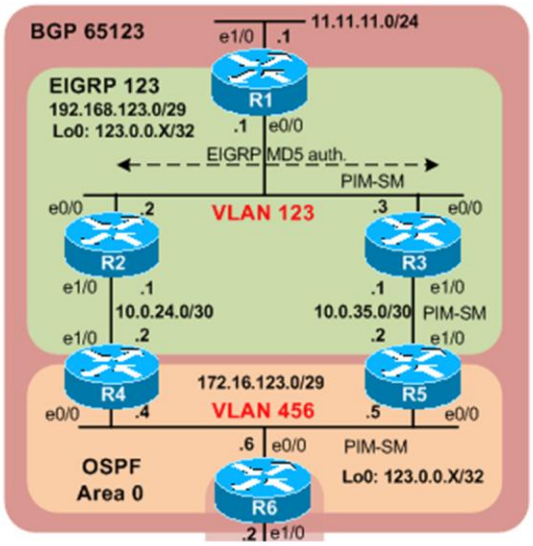

- What can prevent the RP announcements to be propagated?

- Some hypothesis might be:

- ACL blocking the group 224.0.1.40
- Dense mode not enabled or Auto-RP listener not enabled
- RP announcements' scope too short

Auto-RP listener is enabled in R1, R2, R3 but not in R4 and R5 while R4 and R5 are the ones seeing the RP announcements:

R1#sh ip pim auto

AutoRP Information:

AutoRP is enabled.

AutoRP groups over sparse mode interface is enabled

PIM AutoRP Statistics: Sent/Received

RP Announce: 0/0, RP Discovery: 0/0

R1#

R1#sh ip pim rp map

PIM Group-to-RP Mappings

R1#

R4#sh ip pim auto

Incident#5: Proof-test and Analyze hypothesis

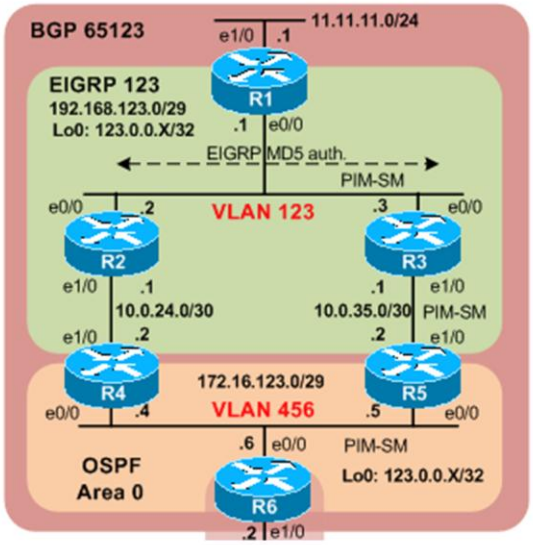

- How to ensure RP announces are flooded in Dense-mode as expected?

Add ‘ip pim autorp listener’ in R4 and R5 does not resolves RP announce up to R1.

Only R1 doesn’t receive any announcement.

Check scope in R6: value of 3 seems fine at first but is not enough: R6 decrement TTL to 2, R4/R5 decrement to 1 when sending to EIGRP; R2/R3 see the announce but don’t forward it !

conf t

ip pim send-rp-announce Loopback0 scope 4

ip pim send-rp-discovery Loopback0 scope 4

end

Incident#5: Implement solution

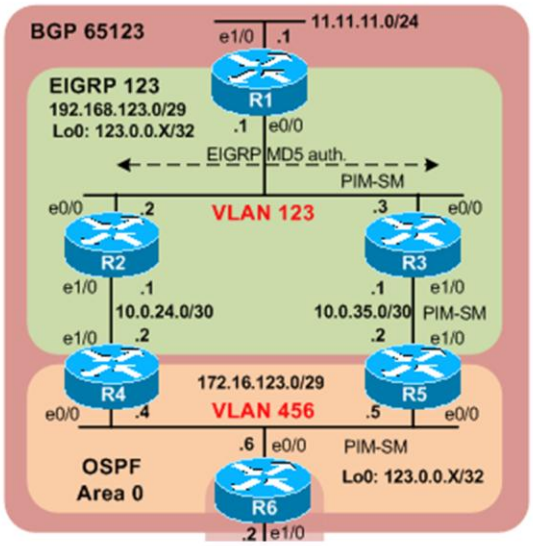

- Add the ‘ip pim autorp listener’ in R4 and R5;
- Increase scope to min 4 in R6’s announcements

Incident#5: Verify resolution within guidelines

- Check the Auto-RP is seen in R1:

```
R1#sh ip pim rp map
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 123.0.0.6 (?), v2v1
    Info source: 123.0.0.6 (?), elected via Auto-RP
         Uptime: 00:04:41, expires: 00:02:14
R1#
```

- Check if required Ping succeeds:

```
R6#ping 231.1.1.1 rep 5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 231.1.1.1, timeout is 2 seconds:
.....
R6#
```

Incident#5: Form hypothesis

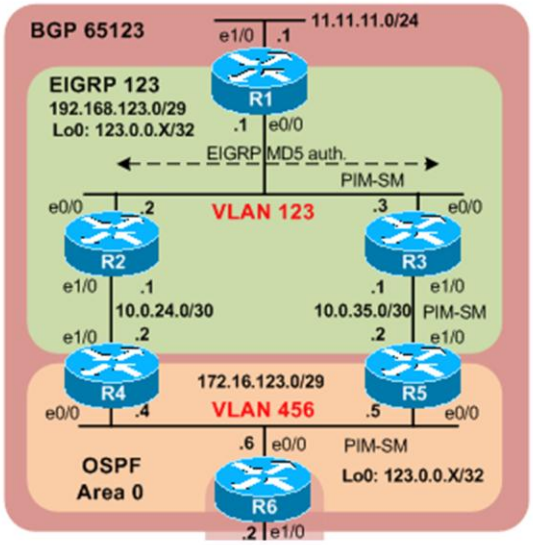

- What could prevent multicast traffic to flow while the control plane is converged?

ACL or similar is a likely reason, lets hunt for it:

Incident#5: Implement solution

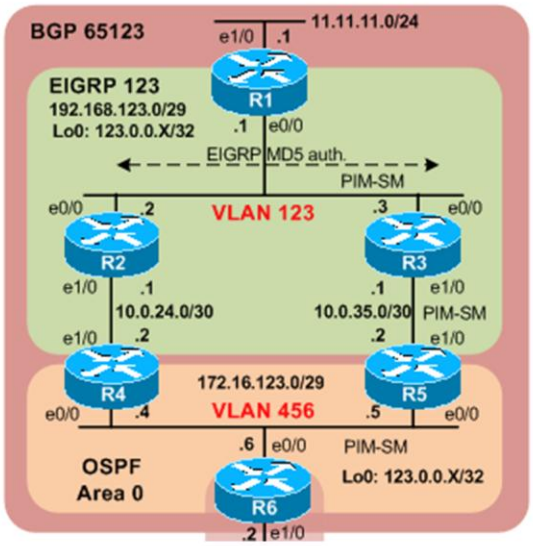

- Resolve ACL in R4 and R5…

- R4 & R5 have an outbound ACL on e1/0 which denies all mcast traffic but the two default groups for RP announce (224.0.0.0 15.255.255.255!!!)

R5#sh access-list

Extended IP access list 199

10 permit ip any host 224.0.1.39 (6 matches)

20 permit ip any host 224.0.1.40 (7 matches)

30 deny ip any 224.0.0.0 15.255.255.255 (3 matches)

40 permit ip any any (24 matches)

R5#

conf t

no access-list 199

access-list 199 permit ip any host 224.0.1.39

access-list 199 permit ip any host 224.0.1.40

access-list 199 **permit** ip any 224.0.0.0 15.255.255.255

access-list 199 permit ip any any

end

Incident#5: Verify resolution within guidelines

- Check the required Ping:

```
R6#ping 231.1.1.1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 231.1.1.1, timeout is 2 seconds:

Reply to request 0 from 192.168.123.1, 4 ms
Reply to request 0 from 192.168.123.1, 4 ms
R6#
```

- Check if initial ACL is still there but resolved:

```
R4#sh ip int e1/0 | i Outgoing
  Outgoing access list is 199
R4#
```

```
R5#sh ip int e1/0 | i Outgoing
  Outgoing access list is 199
R5#
```


- The Guidelines are very clear about ACL, one can not remove it completely but needs to resolve it.

Incident#6

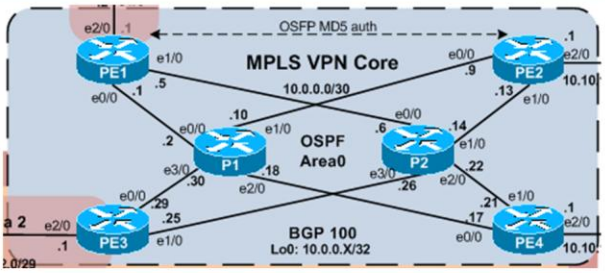

- PE1 is the NTP master for the whole MPLS VPN Core.
- Make sure that all devices are synchronized and authenticating the server.

```
<any MPLS router>#sh ntp assoc det | i auth
10.0.0.111 configured, authenticated, our_master, sane, valid, stratum 1
<any MPLS router>#
```

- Use “CCIE” as the MD5 authentication key (without quotes)
- Refer to the Troubleshooting guidelines to determine if your solution is appropriate.

Incident#6: Define the problem

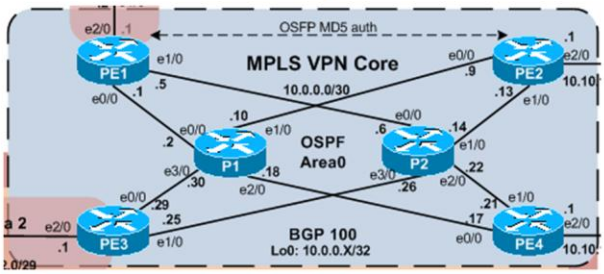

- Which device(s) are and are not synchronized and authenticated?

Only PE3 is not synchronized, all other devices are synchronized and authenticated

Incident#6: Identify symptoms

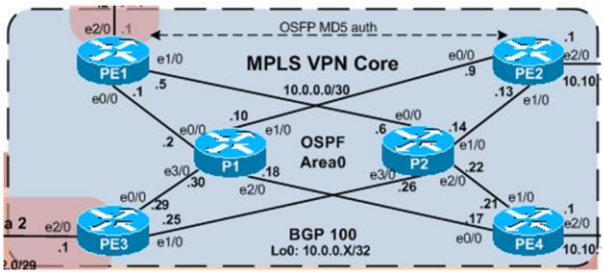

- How is NTP configured?

NTP server is fine:

PE1#sh run | s ntp

ntp authentication-key 1 md5 096F6D203C 7

ntp source Loopback0

ntp master 1

PE1#

NTP config is identical in all NTP client!

PE2#sh run | s ntp

ntp authentication-key 1 md5 0478282F2A 7

ntp authenticate

ntp trusted-key 1

ntp clock-period 17179866

ntp source Loopback0

ntp server 10.0.0.111 key 1

PE2#

PE3#sh run | s ntp

ntp authentication-key 1 md5 0478282F2A 7

Incident#6: Form hypothesis

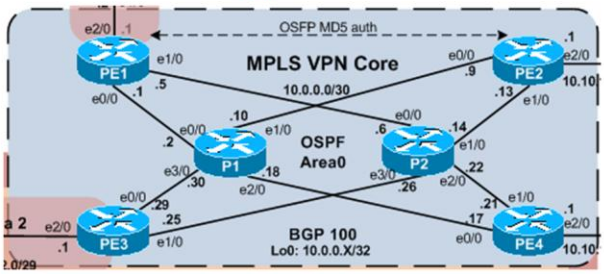

- What can prevent NTP to be synchronized and authenticated in one device only?

NTP is UDP#123

Incident#6: Proof-test and Analyze hypothesis

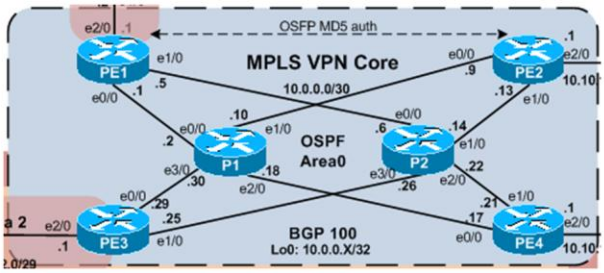

- How is NTP transported?
- What could prevent it to happen?

NTP is UDP#123

PE3 is the only device with CoPP configured

PE3#sh policy-map control-plane

Control Plane

Service-policy input: CoPP

Class-map: CoPP (match-all)

3335 packets, 280001 bytes

5 minute offered rate 0 bps, drop rate 0 bps

Match: access-group 123

police:

rate 1000000 bps, burst 31250 bytes

conformed 3335 packets, 280001 bytes; actions:

transmit

exceeded 0 packets, 0 bytes; actions:

Incident#6: Design and Implement solution

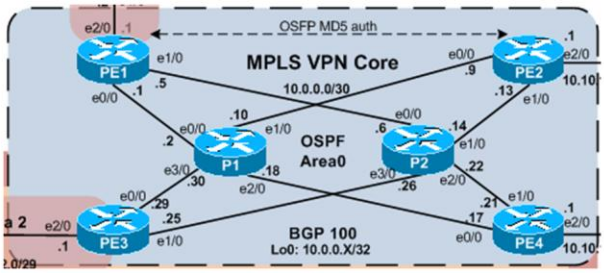

- How to permit NTP in the current config?

Add an entry in the ACL to permit UDP#123

PE3(config)#access-list 123 permit udp any any eq 123

PE3(config)#

NTP: synced to new peer 10.0.0.111

PE3(config)#do sh ntp assoc det | i sane

10.0.0.111 configured, authenticated, our_master, sane, valid, stratum 1

PE3(config)#

Incident#6: Verify resolution within guidelines

- Check NTP association on PE3:

```
PE3#sh ntp assoc det | i auth
10.0.0.111 configured, authenticated, our_master, sane, valid,
stratum 1
PE3#
```

- Incident is resolved! ☺

Agenda

Section 1	CCIE Program Overview
Section 2	CCIE Routing and Switching Version 4
Section 3	Study plan: Content topics review and Sample questions
Section 4	Study plan: Preparation materials
Section 5	Lab Exam: Tips and tricks
Section 6	Troubleshooting Case Studies
Section 7	Conclusion

- Here is again the main message that I would like you to remember when leaving the session.

Complete Your Online Session Evaluation

Complete your session evaluation:

- Directly from your mobile device by visiting www.ciscoliveaustralia.com/mobile and login by entering your badge ID (located on the front of your badge)
- Visit one of the Cisco Live internet stations located throughout the venue
- Open a browser on your own computer to access the Cisco Live onsite portal