360SP04-WKB-DIAG02

Cisco Expert-Level Training for CCIE Service Provider Exercise Workbook Diagnostic Lab 02 Answer Key

This diagnostic module is a technology-specific task that focuses on issue spotting and analysis of networking issues.

Cisco Expert-Level Training for CCIE Service Provider Exercise Workbook Diagnostic Lab 02 Answer Key

Table of Contents

Diagnostic Module

Duration

- Diagnostic section duration: 30 minutes total (6 minutes per task in average)

Diagnostic Guidelines

Caution Read the Following Guidelines Before Starting the Section.

- The diagnostic section is comprised of a set of support tasks for which you are playing the role of Senior Network engineer.
- You have a fixed time of 30 minutes to complete the section, 6 minutes per task in average.
- The real CCIE SP DIAG lab contain 10 tasks and the duration is the 60 minutes.
- The final score of this section is combined with the Troubleshooting and the Configuration sections to comprise your final Pass or Fail status on the CCIE lab exam.
- A student is required to achieve a minimum score in every individual section of the lab exam as well as achieve a minimum overall total score (sum of score in all three sections) in order to pass the Cisco Service Provider certification.
- Each task explicitly mentions what is expected of you.
- Ensure that you carefully read all information provided before selecting your answer(s).
- Select your answer(s) according to the requirement(s) of each task.
- All tasks are independent from each other, i.e. the resolution of a task does not depend on the resolution of any other task.
- The FIVE tasks of this DIAG module shares the same topology and device configuration. Only the e-mails threads and output resources are specific per task.

Network Topology

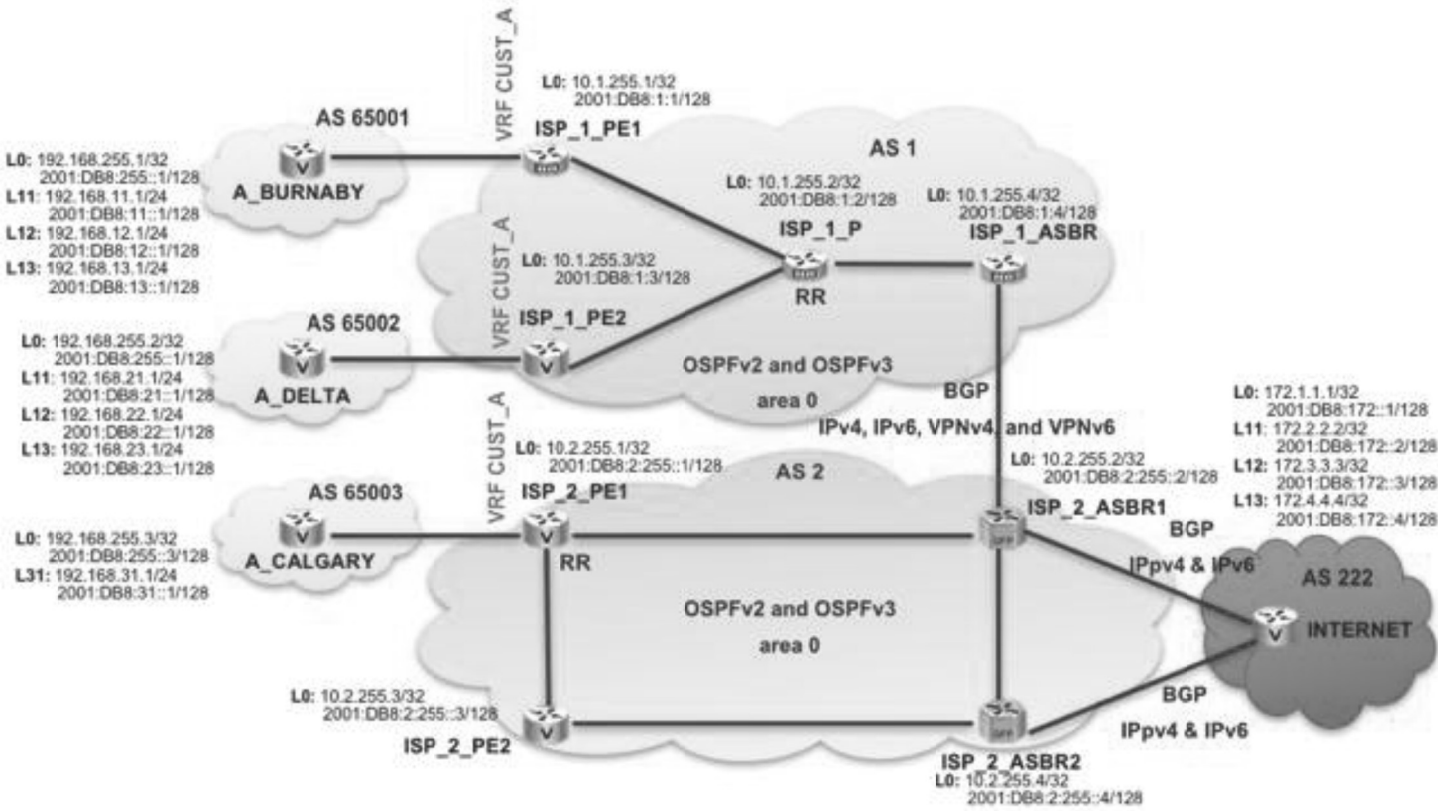

<Click here to return to Task 1 description>
<Click here to return to Task 2 description>
<Click here to return to Task 3 description>
<Click here to return to Task 4 description>
<Click here to return to Task 5 description>

Task 1: Interior Gateway Protocol

Answer

1. There is a duplicate router ID; both routers, ISP_2_PE1 and ISP_2_PE2, use the same router ID.

Clues

1. Focus only between ISP_2_PE1 and ISP_2_PE2:
   a. The e-mail says "deployed exactly the same architecture", it implies that the same area ID design is used for OSPFv2. For OSPFv2 everything is ok. Most likely it's a configuration mistake specific on OSPFv3. So, you can compare the configuration between OSPFv2 and OSPFv3 for these two routers.
   b. Or you can verify that:
      i. in the output of the **show ospfv3 neighbor** command, we do not see any neighborship with regards to this peer, which means that the parameters used for the adjacency criteria is not the issue (such as hello and dead-interval timers, wrong network type, mismatch MTU, etc…)
      ii. in the output of the **show ospfv3** command you can verify that Router ID is the same for both routers: 10.2.255.1

```
ISP_2_PE1# show ospfv3

OSPFv3 1 address-family ipv6
Router ID 10.2.255.1
Supports NSSA (compatible with RFC 3101)
Event-log enabled, Maximum number of events: 1000, Mode: cyclic
Router is not originating router-LSAs with maximum metric
[output omitted]
```

```
ISP_2_PE2# show ospfv3

OSPFv3 1 address-family ipv6
Router ID 10.2.255.1
Supports NSSA (compatible with RFC 3101)
Event-log enabled, Maximum number of events: 1000, Mode: cyclic
Router is not originating router-LSAs with maximum metric
[output omitted]
```

   c. Another approach is to verify the options against the configuration to identify what is the issue.
      i. **IPv6 link local**: by default, the host part of the IPv6 link local address comes from the MAC address. It can be manually configured as well. When manually configured, it can increase the chances of having duplicate addresses due to configuration errors. If you check the interface configuration between these two routers you will not find the command that redefined the IPv6 link local address. So, most likely this is not the issue.
      ii. **Area-ID mismatch**: If you verify the configuration for both routers, you will identify that all interfaces belong to the same area-ID, area 0. So, most likely this is not the issue.
      iii. **IPv6 not enabled.** If the IPv6 is configured under the interface, it implies IPv6 is enabled. So, this option is not the reason of the issue we are troubleshooting either.
      iv. **Network type mismatch**: If you pay attention, you will identify that the **network point-to-point** command is applied only for OSPFv2. For all interfaces enabled on OSPFv3, the network type is broadcast (which is the default for Ethernet interfaces). So, this is not the issue either.

v. **Duplicate Router-**ID: By default, the Router ID, is the loopback address, the highest ipv4 address configured at the time you enable OSPF, in case there is more than one loopback address. ISP_2_PE2 has a manually configured Router ID, and if you verify the ipv4 address used, it's the same as the ISP_2_PE1 loopback 0 address. This is the issue.

- on ISP_2_PE1

```
interface Loopback0
 ip address 10.2.255.1 255.255.255.255
 ipv6 address 2001:DB8:2:255::1/128
 ospfv3 1 ipv6 area 0
```

- on ISP_2_PE2

```
router ospfv3 1
 router-id 10.2.255.1
 !
 address-family ipv6 unicast
 exit-address-family
```

Task 2: MPLS Traffic Engineering

Answer

2. There is not enough bandwidth resource available for the MPLS TE tunnel in the path between ISP_1_P and ISP_1_PE2.

Clues

1. For this troubleshooting analysis, we need to verify three routers, which are in the path for this MPLS TE tunnel: ISP_1_PE1, ISP_1_P, ISP_1_PE2.
   - First, let's check the output of the **show mpls traffic-eng tunnels** command on ISP_1_PE1 (head end). In this output, you can verify the reason why the tunnel status is down: **PCALC Error ... info: No path to destination, 10.1.255.3 (bw).**
   - This indicates that there isn't enough bandwidth available in the path that the MPLS TE tunnels is requesting. The tunnel is requesting 45000 Kbps, you can also verify this in the same output.

```
ISP_1_PE1# show mpls traffic-eng tunnels

Name: tunnel-te1  Destination: 10.1.255.3  Ifhandle:0x480
Signalled-Name: ISP_1_PE1_t1
Status:
Admin:    up Oper: down   Path: not valid   Signalling: Down

path option 1,  type dynamic
Last PCALC Error: Tue Jun 27 05:26:12 2017
Info: No path to destination, 10.1.255.3 (bw)
G-PID: 0x0800 (derived from egress interface properties)
Bandwidth Requested: 45000 kbps  CT0
Creation Time: Tue Jun 27 05:12:51 2017 (1d00h ago)
Config Parameters:
Bandwidth:    45000 kbps (CT0) Priority:  7  7 Affinity: 0x0/0xffff
Metric Type: TE (default)
Hop-limit: disabled
Cost-limit: disabled
AutoRoute:  enabled  LockDown: disabled   Policy class: not set
Forward class: 0 (default)
Forwarding-Adjacency: disabled
Loadshare:          0 equal loadshares
Auto-bw: disabled
Fast Reroute: Disabled, Protection Desired: None
Path Protection: Not Enabled
BFD Fast Detection: Disabled
Reoptimization after affinity failure: Enabled
Soft Preemption: Disabled
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 0 up, 1 down, 0 recovering, 0 recovered heads
```

   - You need to verify the bandwidth allocated in the RSVP command in the path of this tunnel. In the output of the **show rsvp interface** command of the ISP_1_P, you will identify that the interface Gi0/0/0/0 only has 10Mbps of bandwidth enabled on RSVP. This is the root cause of this issue.
2. Alternatively, you check the following options to identify what is the root-cause of the issue.
   - Missing a static route: This is not the reason why the MPLS TE tunnel is down.
   - RSVP not enabled. If you verify the configuration and the output for ISP_1_P you will see that RSVP is enabled on interfaces that are part in the path of the MPLS TE tunnel we are troubleshooting.

- MPLS traffic-eng not enabled globally. If you verify the configuration, you will see the MPLS TE enabled. So, this is not the root cause of the issue we are analyzing.
- Not enough bandwidth,
    i. on ISP_1_P you can see the following configuration:

```
rsvp
 interface GigabitEthernet0/0/0/0
  bandwidth 10000
 !
 interface GigabitEthernet0/0/0/1
  bandwidth 100000
```

    ii. on ISP_1_PE1 you can see the following configuration:

```
interface tunnel-te1
 ipv4 unnumbered Loopback0
 signalled-bandwidth 45000
 autoroute announce
 !
 destination 10.1.255.3
 path-option 1 dynamic
 logging events link-status
```

    iii. GigabitEthernet 0/0/0/0 does not have enough bandwidth to meet the MPLS TE tunnel requirements.

Task 3: L3VPN

Answer

ISP_1_ASBR

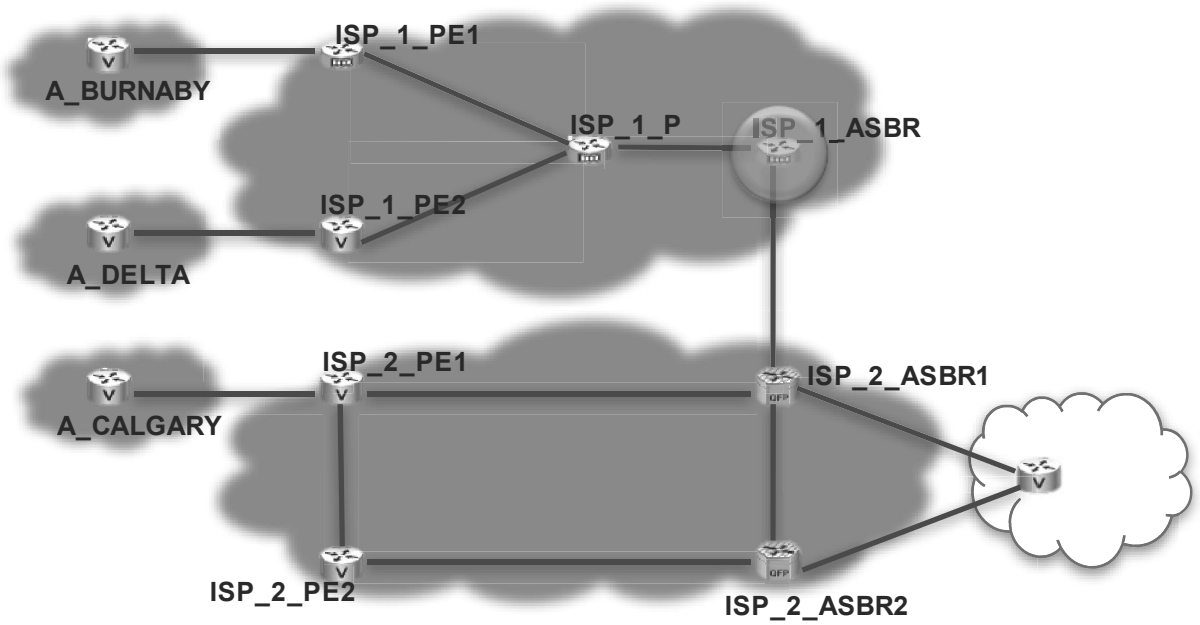

Clues

1. For this type of task, you do not need to identify the root cause of the issue, just ensure you identify where the problem is originated.
    - As A_BURNABY and A_DELTA can communicate to each other, this implies that the label path is fine between ISP_1_PE1, ISP_1_P, and ISP_1_PE2. So, we only have the option to check between ISP_1_ASBR and ISP_2_ASBR1.
    - In the **show bgp vpnv4 unicast** and **show bgpnv4 unicast label** commands for both ASBRs, you can verify that they are exchanging VPN prefixes. With this you can assume that we are talking about the MPLS Inter-AS option B solution. As the VPN prefixes are being exchanged, the next step of the verification is to check the BGP next-hop labels between these two ASBRs. For that, you need to focus on the output of the **show mpls forward** command with regards to the BGP next-hop, which is the IPv4 address used for the BGP peering between the ASBRs (192.168.101.2 and 192.168.101.1). You must have a /32 entry for these addresses and a local and outgoing label associated to it.

```
ISP_2_ASBR1# show mpls forwarding
Load for five secs: 1%/0%; one minute: 1%; five minutes: 1%

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
16         Pop Label  10.2.13.0/24     0             Gi2        10.2.12.1
17         Pop Label  10.2.14.0/24     0             Gi5        10.2.34.4
18         Pop Label  10.2.255.1/32    0             Gi2        10.2.12.1
19         16         10.2.255.3/32    0             Gi2        10.2.12.1
    20         10.2.255.3/32    0             Gi5        10.2.34.4
20         Pop Label  10.2.255.4/32    214137        Gi5        10.2.34.4
21         Pop Label  192.168.101.1/32 0             Gi3        192.168.101.1
<output omitted>
```

    - On ISP_1_ASBR you can't find the /32 entry. So, this is the device cause the issue.


```
RP/0/0/CPU0:ISP_1_ASBR# show mpls forwarding
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ----------
--
24000  24000       10.1.255.1/32      Gi0/0/0/0    10.1.24.2       1620
24001  Pop         10.1.255.2/32      Gi0/0/0/0    10.1.24.2       286613
24002  Pop         10.1.12.0/24       Gi0/0/0/0    10.1.24.2       0
24003  Pop         10.1.23.0/24       Gi0/0/0/0    10.1.24.2       0
24005  24004       1:1011:192.168.11.0/24   \
10.1.255.1      0
24006  24005       1:1011:192.168.12.0/24   \
10.1.255.1      0
24007  24006       1:1011:192.168.13.0/24   \
10.1.255.1      0
24008  24007       1:1011:192.168.91.0/24   \
10.1.255.1      520
24009  24008       1:1011:192.168.255.1/32   \
10.1.255.1      1040
24010  22          2:1011:192.168.31.0/24   \
192.168.101.2   0
24011  23          2:1011:192.168.255.3/32   \
192.168.101.2   1560
24012  24          2:1011:2001:db8:255::3/128   \
192.168.101.2   0
24013  24002       10.1.255.3/32      Gi0/0/0/0    10.1.24.2       540
24014  31          2:1011:10.18.0.0/16   \
       10.1.255.3      0
24015  32          2:1011:192.168.21.0/24   \
       10.1.255.3      0
24016  33          2:1011:192.168.22.0/24   \
       10.1.255.3      0
24017  34          2:1011:192.168.23.0/24   \
       10.1.255.3      0
24018  35          2:1011:192.168.92.0/24   \
       10.1.255.3      0
24019  36          2:1011:192.168.255.2/32   \
       10.1.255.3      520
24020  29          2:1011:2001:db8:192:92::/64   \
       10.1.255.3      0
24021  30          2:1011:2001:db8:255::2/128   \
       10.1.255.3      0
24022  38          2:1011:192.168.93.0/24   \
       192.168.101.2   0
24023  39          2:1011:2001:db8:31::1/128   \
       192.168.101.2   0
24024  40          2:1011:2001:db8:192:93::/64   \
       192.168.101.2   0
24025  Aggregate   192.168.101.0/24   default
```

- For Cisco IOS XR, you must have a static route using a /32 prefix for the IPv4 address used for the BGP peering, in order to allocate label for the /32 prefix. If you check the configuration, the static route is missing. For Cisco IOS, this is done automatically, you do not need to add the static route.

Task 4: PE-CE Connectivity

Answer

device	A_DELTA
issue	The **route-map** applied uses an **access-list** that does not exist

Clues

1. You can exclude ISP_1_ASBR, let's check if the issue is on the CE or PE side.
   a. From the output of the **show bgp ipv4 unicast neighbor 192.168.92.3 advertised-routes** command, you can verify that A_DELTA is advertising everything. So, this concludes that A_DELTA is the device with the problem.
   b. Next step is to identify what is the mistake.

```
A_DELTA# show route-map

route-map ALLOW_SUBNETS, permit, sequence 10
Match clauses:
ip address (access-lists): access-list ALLOW_SUBNETS
Set clauses:
Policy routing matches: 0 packets, 0 bytes
route-map ALLOW_SUBNETS, deny, sequence 100
Match clauses:
Set clauses:
Policy routing matches: 0 packets, 0 bytes

A_DELTA# show access-list

Standard IP access list ALLOW_SUBNETS
10 permit 192.168.22.0, wildcard bits 0.0.0.255
```

   Based on these two outputs you can verify that the access-list named as **access-list** does not exist.

2. If you check all other reasons why the PE-CE filtering is not working are all incorrect.

Task 5: Routing/Fast Convergence

Answer

drop down list
Apply advertise external path feature on the 2nd ASBR device.

Clues

1. This task is easy if you know all the BGP advanced features well, which means you do not need to check any configuration to identify the correct option. The task is assigned as Intermediate because those BGP features are not widely known in the community yet. Each one of the new BGP advanced feature has a different impact in the BGP database, the BGP advertisement, and the RIB/FIB population.
    a. The bandwidth community does not enable the ICMP requirement. It's only an alternative in case of a tie-break of the BGP best path selection process.
    b. An enabled BGP multipath on a route reflector does not change how many prefixes the route reflector will learn.
    c. BPG PIC core is enabled by default.
    d. IP FRR and MPLS TE/FRR do not change the BGP advertisement behavior.
    e. LDP session protection does not change the BGP advertisement behavior.
    f. Apply advertise external path, makes the ASBR to advise the prefix to the route-reflector, even though for this ASBR the best path is an internal path. This command changes the BGP advertisement behavior.

Conclusion

1. In the Diagnostic module, you do not have too much time. You must look for clues to where to start the verification in order to identify the correct answer as fast as possible.
2. Double check if your answers are correct by checking other output or device configuration.
3. You should not spend time checking the entire configuration for all devices. You will waste a lot of time and it is not always easy to check for a mistake if you do not know where to start. If you do, most likely you will run out of the time before you start answering the questions.
4. Outputs are your best friend. Check the configuration only to verify if your answer is the correct one. Focus on in a section of the device configuration, and not the entire device configuration.
5. Do not get stuck in one task. Move to the next one. You can always return back.
6. Check the title of the task, it tells you which blueprint topic the task is related to. Start with a topic that is more familiar to you. Leave the most difficult topic for later. The last task can be the easier for you. You do not need to follow the order of the task numbers as they are display in your exam.
    a. Remember, you have 60 minutes to answer 10 tasks. This gives you 6 minutes per task on average.
    b. Remember all tasks are worth 1 point.

