Vendor: Juniper

Exam Code: JN0-694

Exam Name: Enterprise Routing and Switching Support, Professional (JNCSP-ENT)

Version: Demo

QUESTION 1
You are implementing Q-in-Q tunneling on an EX Series switch. You want the tunnel to support all C-VLANs; however, only some VLANs are able to send traffic across the tunnel. Switch-1 has the following configuration:

```
[edit vlans]
user@Switch-1# show
v100 {
vlan-id 100;
interface {
ge-0/0/0.10;
ge-0/0/1.20;
}
dot1q-tunneling {
customer-vlans [ ];
}
}
```

What would solve this problem?

A. Add family ethernet-switching to the tunnel-side interface on Switch-1.
B. Implement RSTP.
C. Q-in-Q tunneling will not work in this scenario; use a Layer 2 VPN instead.
D. Remove the customer-vlans statement.

Correct Answer: C

QUESTION 2
You are troubleshooting a problem where an OSPF adjacency between two neighboring routers will not form.
What are two reasons for this problem? (Choose two.)

A. One or both of the connected interfaces are missing the family inet statement.
B. One or both of the connected interfaces are missing the family iso statement.
C. The connected interfaces are not on the same subnet.
D. Another IGP is running on one or both of the routers, overriding OSPF.

Correct Answer: BD

QUESTION 3
Your Junos device is dropping certain traffic flows, while allowing other traffic flows to pass through the device unaffected.

Which CoS component is causing this problem?

A. BA classification
B. RED
C. MF classification
D. Rewrite rules

Correct Answer: D

QUESTION 4
Two neighboring routers are able to form an OSPF adjacency, but are not able to establish an IBGP neighborship.

What are two reasons for the IBGP neighborship problem? (Choose two.)

A. One of the devices has a misconfigured BGP peer address.
B. One or both of the connected interfaces are missing the family iso statement.
C. OSPF has a lower route preference than BGP.
D. A firewall filter on one of the interfaces is blocking TCP traffic.

Correct Answer: BC

QUESTION 5
Your switch is experiencing a problem where a port that should have only one host connected occasionally shows that multiple MAC addresses are being learned.

Which configuration setting would ensure that no extra hosts can join the network using this switch port?

A. mac-limit
B. no-mac-learning
C. persistent-learning
D. bpdu-block-on-edge

Correct Answer: D

QUESTION 6
You are having problems redistributing RIP routes into OSPF. Your Junos device has the following configuration:

```
[edit protocols ospf]
user@router# show
import my-policy;
area 0.0.0.0 {
interface ge-0/0/0.0;
interface ge-0/0/ ;
interface ge-0/0/ {
passive;
}
}
```

What would resolve the problem?

A. Apply my-policy as an export policy under the [edit protocols rip] hierarchy.
B. Apply my-policy as an import policy under the [edit protocols rip] hierarchy.
C. Apply my-policy as an export policy under the [edit protocols ospf] hierarchy.
D. Use the area-range parameter instead of a routing policy.

Correct Answer: D

QUESTION 7
You observe that a router is using an unusually high amount of CPU cycles. You determine that continuous SPF calculations in OSPF are being performed.

What are two reasons for this problem? (Choose two.)

A. The wrong authentication keys between the OSPF neighbors are used.
B. The interface MTU is mismatched between the OSPF neighbors.
C. There are duplicate router IDs within the OSPF area.
D. An OSPF adjacency is flapping.

Correct Answer: CD

QUESTION 8
Your Layer 2 network uses VLAN IDs 100 through 400 and you are required to load-balance these VLANs

between two different root bridges. You are currently using the default RSTP settings and notice that all VLANs are using the same root bridge.

How do you ensure the VLANs are load-balanced between two root bridges?

A. Configure MSTP with two MSTI regions and split the VLAN range between them.
B. Configure VSTP with two VLAN groups and split the VLAN range between them.
C. Configure two RSTP instances and split the VLAN range between them.
D. Configure STP and RSTP and split the VLAN range between them.

Correct Answer: A

QUESTION 9
An end user on interface ge-0/0/1.0 is trying to receive a multicast stream for 232.0.0.1 sourced from but is not receiving it. You use the show igmp group command and do not see this group in the list. You enable traceoptions for IGMP and find the following IGMPv3 report from the end user's host:
Jun 10 13:11:03.577641 RCV IGMP V3 Report len 16 from 192.168.1.13 intf ge-0/0/1.0 Jun 10 13:11:03.577984 Records 1
Jun 10 13:11:03.578027 Group 232.0.0.1, type IS_EX, aux_len 0, sources 0

Which configuration change is required to allow the group to be added in the router?

A. set routing-options multicast ssm-groups 232.0.0.1/32
B. set routing-options multicast asm-override-ssm
C. set protocols igmp interface ge-0/0/1.0 promiscuous-mode
D. set protocols igmp interface ge-0/0/1.0 group-limit 2

Correct Answer: B

QUESTION 10
You recently deployed two Anycast RPs. Multicast clients in the network are reporting that they are receiving traffic from some, but not all, multicast sources.

What are two solutions for this problem? (Choose two.)

A. Configure MSDP between the Anycast RPs.
B. Configure rp-set for the Anycast RPs.
C. Configure multicast BGP between the Anycast RPs.
D. Configure the network to always use the RPT and not switch over to the SPT.

Correct Answer: AB

QUESTION 11
You use static routes for connectivity to the ISP. Your ISP recently switched to using different links for multicast and unicast traffic. Following the change, users in your company were unable to receive multicast traffic through the ISP.

What must you configure on your router to reestablish multicast connectivity to your ISP?

A. Add a static default route to the ISP in the inet.2 routing table.
B. Add the default-rpf-interface parameter under the [edit routing-options multicast] hierarchy.
C. Add the upstream-interface parameter under the [edit protocols pim] hierarchy.
D. Disable PIM on the interface used for unicast traffic.

Correct Answer: A

QUESTION 12
-- Exhibit --
user@router# run show log ospf-test

...
Jun 10 22:35:38.598494 OSPF sent Hello 10.100.0.1 -> 224.0.0.5 (ge-1/0/3.1000 IFL 77 area 0.0.0.0) Jun 10 22:35:38.598520 Version 2, length 44, ID 10.100.1.2, area 0.0.0.0 Jun 10 22:35:38.598543 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 128 Jun 10 22:35:38.598564 dead_ivl 32, DR 10.100.0.1, BDR 0.0.0.0 Jun 10 22:35:41.522956 OSPF periodic xmit from 10.200.26.1 to 224.0.0.5 (IFL 2684276196 area 0.0.0.1)
Jun 10 22:35:42.798220 OSPF rcvd Hello 10.100.0.2 -> 224.0.0.5 (ge-1/0/3.1000 IFL 77 area 0.0.0.0) Jun 10 22:35:42.798311 Version 2, length 48, ID 10.100.1.1, area 0.0.0.0 Jun 10 22:35:42.798334 checksum 0x0, authtype 0
Jun 10 22:35:42.798356 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 128 Jun 10 22:35:42.798377 dead_ivl 40, DR 10.100.0.2, BDR 10.100.0.1 Jun 10 22:35:45.189034 OSPF rcvd Hello 10.100.0.2 -> 224.0.0.5 (ge-1/0/3.1000 IFL 77 area 0.0.0.0) Jun 10 22:35:45.189097 Version 2, length 44, ID 10.100.1.1, area 0.0.0.0 Jun 10 22:35:45.189118 checksum 0x0, authtype 0
Jun 10 22:35:45.189140 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 128 Jun 10 22:35:45.189162 dead_ivl 40, DR 10.100.0.2, BDR 0.0.0.0 Jun 10 22:35:45.196969 OSPF DR is 10.100.1.2, BDR is 0.0.0.0 Jun 10 22:35:45.197050 OSPF sent Hello 10.200.26.1 -> 224.0.0.5 (ge-1/0/0.0 IFL 69 area 0.0.0.1) Jun 10 22:35:45.197076 Version 2, length 44, ID 10.100.1.2, area 0.0.0.1 Jun 10 22:35:45.197098 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 128 Jun 10 22:35:45.197119 dead_ivl 40, DR 10.200.26.1, BDR 0.0.0.0 Jun 10 22:35:46.746900 OSPF periodic xmit from 10.100.0.1 to 224.0.0.5 (IFL 2684276196 area 0.0.0.0)
-- Exhibit --

Click the Exhibit button.

Referring to the exhibit, what is preventing the OSPF neighborship with two directly connected routers using interface ge-1/0/3 from reaching the full state?

A. dead interval mismatch
B. authentication type mismatch
C. subnet mismatch
D. hello interval mismatch

Correct Answer: A

QUESTION 13
-- Exhibit --
user@router> show ospf database

Area 0.0.0.1
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 172.24.255.1 172.24.255.1 0x800000d4 182 0x22 0x59f3 36 Router 172.24.255.2 172.24.255.2 0x800000d4 177 0x22 0x57f2 36 Router *172.24.255.4 172.24.255.4 0x800000dc 176 0x22 0x75fa 72 Network 172.24.124.2 172.24.255.2 0x80000007 177 0x22 0x7957 36 Summary 172.24.13.0 172.24.255.1 0x80000004 2370 0x22 0x3f62 28 Summary 172.24.23.0 172.24.255.1 0x80000002 471 0x22 0xdeb9 28 Summary 172.24.255.1 172.24.255.1 0x800000cb 2037 0x22 0x2bbb 28 Summary 172.24.255.2 172.24.255.2 0x800000cc 487 0x22 0x19ca 28 Summary 172.24.255.3 172.24.255.1 0x80000003 140 0x22 0xb2f9 28 OSPF AS SCOPE link state database
Type ID Adv Rtr Seq Age Opt Cksum Len
Extern *1.47.82.0 172.24.255.4 0x80000002 1037 0x22 0x4225 36 Extern *100.0.0.0 172.24.255.4 0x80000001 2643 0x22 0xfc88 36

user@router> show ospf neighbor
Address Interface State ID Pri Dead
172.24.124.2 ge-0/0/1.0 Full 172.24.255.2 128 36
172.24.124.1 ge-0/0/1.0 Full 172.24.255.1 128 30

user@router> show ospf interface ge-0/0/1.0 extensive
Interface State Area DR ID BDR ID Nbrs
ge-0/0/1.0 PtToPt 0.0.0.1 0.0.0.0 0.0.0.0 2
Type: P2MP, Address: 172.24.124.4, Mask: 255.255.255.0, MTU: 1500, Cost: 1 Adj count: 2
Hello: 10, DeaD. 40, ReXmit: 5, Not Stub
Auth type: None
Protection type: None

Topology default (ID 0) -> Cost: 1
user@router> show route protocol ospf table inet.0

inet.0: 11133 destinations, 11135 routes (11133 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both

224.0.0.5/32 *[OSPF/10] 1w0d 00:01:14, metric 1
MultiRecv
-- Exhibit --

Click the Exhibit button.

Referring to the exhibit, why are the OSPF routes missing from the routing table for this router?

A. mismatching OSPF interface type with the neighbor
B. MTU mismatch with the neighbor
C. incorrect IP address configured on the interface
D. no Type 4 LSAs in the OSPF database

Correct Answer: A

QUESTION 14
-- Exhibit

```
user@R1> show route
inet.0: 5 destinations, 5 routes (5 active, 0
holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[Direct/0] 00:01:10
                     > via lo0.0
2.2.2.2/32          *[OSPF/10] 00:00:13, metric 1
                     > to 172.10.1.2 via ge-0/0/1.0
172.10.1.0/24       *[Direct/0] 00:01:10
                     > via ge-0/0/1.0
172.10.1.1/32       *[Local/0] 00:01:10
                       Local via ge-0/0/1.0
224.0.0.5/32        *[OSPF/10] 00:01:10, metric 1
                       MultiRecv
```

```
user@R1> show ospf interface ge-0/0/1.0 detail
Interface  State Area        DR ID         BDR ID    Nbrs
ge-0/0/1.0  DR  0.0.0.0  200.200.200.200 2.2.2.2     1
  Type: LAN, Address: 172.10.1.1, Mask: 255.255.255.0,
  MTU: 1500, Cost: 1
  DR addr: 172.10.1.1, BDR addr: 172.10.1.2, Priority:
128
  …
user@R1> show ospf neighbor detail
Address      Interface   State     ID           Pri  Dead
172.10.1.2   ge-0/0/1.0  Full      2.2.2.2      128    31
 …
```

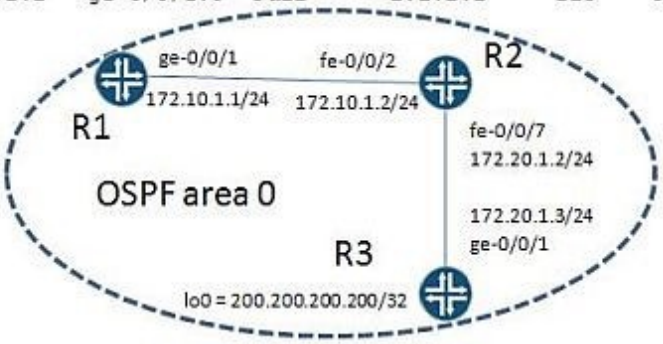


```
user@R1> show ospf database
Jun 12 03:33:34
    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   2.2.2.2          2.2.2.2         0x80000005    30  0x22 0xeb10  60
Router  *200.200.200.200  200.200.200.200 0x80000009     7  0x22 0x1d42  48
Network *172.10.1.1       200.200.200.200 0x80000005     2  0x22 0xcc62  32
Network *172.20.1.3       200.200.200.200 0x80000004  3600  0x22 0x42e1  32

user@R1> show ospf database
Jun 12 03:33:46
    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   2.2.2.2          2.2.2.2         0x80000005    42  0x22 0xeb10  60
Router  *200.200.200.200  200.200.200.200 0x8000000d     3  0x22 0x1546  48
Network *172.10.1.1       200.200.200.200 0x80000006     6  0x22 0xca63  32
Network *172.20.1.3       200.200.200.200 0x80000005  3600  0x22 0x40e2  32
```

-- Exhibit --

Click the Exhibit button.

Referring to the exhibit, you are configuring an OSPF network. All OSPF adjacencies come up and stay stable. But neither R1 nor R2 has the prefix 200.200.200.200/32 in its routing table.

What is causing this problem?

A. R2 does not have the export policy for prefix 200.200.200.200/32.
B. R1 does not have routes to network 172.10.1.0/24.
C. R2 is BDR on both network 172.10.1.0/24 and 172.20.1.0/24.
D. The router ID of R1 is the same as the router ID of R3.

Correct Answer: D

QUESTION 15
-- Exhibit


```
user@R1# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/2.0 {
        hello-interval 10;
        dead-interval 40;
    }
}

[edit]
user@R1# show interfaces ge-0/0/2
mtu 1500;
unit 0 {
    family inet {
        address 192.168.1.1/24;
    }
}
```

```
user@R2# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/2.0;
}

[edit]
user@R2# show interfaces ge-0/0/2
unit 0 {
    family inet {
        address 192.168.1.2/24;
    }
}
```

-- Exhibit --

Click the Exhibit button.

You are troubleshooting an OSPF adjacency problem between R1 and R2.

Referring to the exhibit, what is causing this OSPF adjacency problem?

A. There is a hello interval mismatch.
B. There is a dead interval mismatch.
C. There is an MTU mismatch.
D. There is an LSA refresh timer mismatch.

Correct Answer: C

QUESTION 16
-- Exhibit --
Jun 12 02:56:06 R1 rpd[60735]: RPD_OSPF_NBRDOWN: OSPF neighbor 10.50.10.25 (realm ospf-v2 fe-0/0/4.0 area 0.0.0.0) state changed from Full to Init due to 1WayRcvd (event reason: neighbor is in one-way mode)
Jun 12 02:59:36 R1 rpd[60735]: RPD_OSPF_NBRUP: OSPF neighbor 10.50.10.25 (realm ospf-v2 fe-0/0/4.0 area 0.0.0.0) state changed from Init to ExStart due to 2WayRcvd (event reason: neighbor detected this router)
Jun 12 02:59:36 R1 rpd[60735]: RPD_OSPF_NBRUP: OSPF neighbor 10.50.10.25 (realm ospf-v2 fe-0/0/4.0 area 0.0.0.0) state changed from Exchange to Full due to ExchangeDone (event reason: DBD exchange of slave completed)
-- Exhibit --

Click the Exhibit button.

You notice that there is a problem with the OSPF adjacency between two routers, R1 and R2. The relevant system logs from R1 are shown in the exhibit.

What would cause this behavior?

A. R2 was dropping R1's OSPF hello packets.
B. R1 was dropping R2's OSPF hello packets.
C. R1's interface went down and came back up.
D. There is an OSPF hello timer mismatch between the two routers.

Correct Answer: A

QUESTION 17
-- Exhibit --
user@R1> show ospf neighbor
Address Interface State ID Pri Dead
10.222.0.2 ge-0/0/1.0 Init 10.222.1.2 128 32

user@R1> show ospf interface detail
Interface State Area DR ID BDR ID Nbrs
ge-0/0/1.0 DR 0.0.0.0 10.222.1.1 0.0.0.0 1
Type: LAN, Address: 10.222.0.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1 DR addr: 10.222.0.1, Priority: 128
Adj count: 0
Hello: 10, DeaD. 40, ReXmit: 5, Not Stub
Auth type: MD5, Active key ID. 10, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None
Topology default (ID 0) -> Cost: 1
lo0.0 DR 0.0.0.0 10.222.1.1 0.0.0.0 0
Type: LAN, Address: 10.222.1.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0 DR addr: 10.222.1.1, Priority: 128
Adj count: 0
Hello: 10, DeaD. 40, ReXmit: 5, Not Stub
Auth type: None
Protection type: None
Topology default (ID 0) -> Cost: 0

user@R2> show ospf neighbor

user@R2> show ospf interface detail
Interface State Area DR ID BDR ID Nbrs
ge-0/0/1.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 0
Type: P2P, Address: 10.222.0.2, Mask: 255.255.255.252, MTU: 1500, Cost: 1 Adj count: 0
Hello: 10, DeaD. 40, ReXmit: 5, Not Stub
Auth type: MD5, Active key ID. 10, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None
Topology default (ID 0) -> Cost: 1

lo0.0 DR 0.0.0.0 10.222.1.2 0.0.0.0 0
Type: LAN, Address: 10.222.1.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0 DR addr: 10.222.1.2, Priority: 128
Adj count: 0
Hello: 10, DeaD. 40, ReXmit: 5, Not Stub
Auth type: None
Protection type: None
Topology default (ID 0) -> Cost: 0
-- Exhibit --

Click the Exhibit button.

You are trying to establish an OSPF adjacency between R1 and R2, but the adjacency does not establish.

Referring to the exhibit, what is causing the adjacency to fail?

A. The MD5 key ID values are mismatched between R1 and R2.
B. R1 has both family inet and family iso configured on the link toward R2.
C. The IP subnet mask is mismatched between R1 and R2.
D. The interface type is mismatched between R1 and R2.

Correct Answer: D

QUESTION 18
-- Exhibit

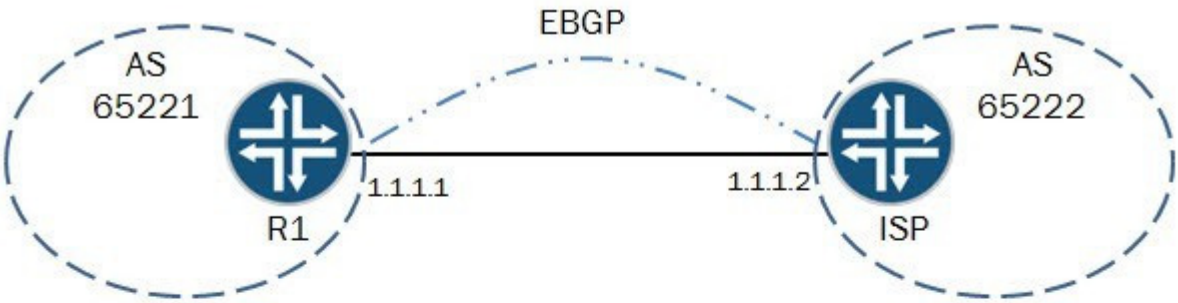


```
[edit]
user@R1# show protocols bgp
group ebgp {
    type external;
    multihop;
    local-address 1.1.1.1;
    peer-as 65222;
    neighbor 172.16.1.2;
}

[edit]
user@R1# show routing-options
static {
    route 172.16.1.2/32 next-hop 1.1.1.1;
}
router-id 10.222.1.1;
autonomous-system 65221;
```

Loopback IP:
172.16.1.2

-- Exhibit --

Click the Exhibit button.
You are asked to configure a multihop EBGP peering to a loopback address of your ISP. The peering does not establish, and the ISP has verified that the settings are correct on their side.

Referring to the exhibit, what is causing the problem?

A. The peer-as parameter is misconfigured.
B. The local-address parameter is misconfigured.
C. The static route is misconfigured.
D. The neighbor parameter is misconfigured.

Correct Answer: C

QUESTION 19
-- Exhibit --
user@router# run show log bgp-test
...
Jun 10 23:50:43.056697 BGP SEND 192.168.133.1+179 -> 192.168.133.0+64925 Jun 10 23:50:43.056739 BGP SEND message type 3 (Notification) length 23 Jun 10 23:50:43.056760 BGP SEND Notification code 2 (Open Message Error) subcode 7 (unsupported capability)
Jun 10 23:50:43.056781 BGP SEND Data (2 bytes): 00 04
Jun 10 23:50:52.215104 advertising receiving-speaker only capabilty to neighbor ::192.168.133.0 (External AS 300)
Jun 10 23:50:52.215173 bgp_senD. sending 59 bytes to ::192.168.133.0 (External AS 300) Jun 10 23:50:52.215200
Jun 10 23:50:52.215200 BGP SEND ::192.168.133.1+179 -> ::192.168.133.0+57107 Jun 10 23:50:52.215233 BGP SEND message type 1 (Open) length 59 Jun 10 23:50:52.215256 BGP SEND version 4 as 23456 holdtime 90 id 10.200.1.1 parmlen 30 Jun 10 23:50:52.215276 BGP SEND MP capability AFI=2, SAFI=1 Jun 10 23:50:52.215294 BGP SEND Refresh capability, code=128 Jun 10 23:50:52.215312 BGP SEND Refresh capability, code=2 Jun 10 23:50:52.215332 BGP SEND Restart capability, code=64, time=120, flags= Jun 10 23:50:52.215353 BGP SEND 4 Byte AS-Path capability (65), as_num 2123456789 Jun 10 23:50:52.216018
Jun 10 23:50:52.216018 BGP RECV ::192.168.133.0+57107 -> ::192.168.133.1+179 Jun 10 23:50:52.216058 BGP RECV message type 3 (Notification) length 21 Jun 10 23:50:52.216079 BGP RECV Notification code 2 (Open Message Error) subcode 2 (bad peer AS number)
Jun 10 23:51:15.058112 advertising receiving-speaker only capabilty to neighbor 192.168.133.0 (External AS 300)
Jun 10 23:51:15.058192 bgp_senD. sending 59 bytes to 192.168.133.0 (External AS 300) Jun 10 23:51:15.058217
Jun 10 23:51:15.058217 BGP SEND 192.168.133.1+50083 -> 192.168.133.0+179 Jun 10 23:51:15.058250 BGP SEND message type 1 (Open) length 59 Jun 10 23:51:15.058273 BGP SEND version 4 as 65001 holdtime 90 id 10.200.1.1 parmlen 30 Jun 10 23:51:15.058294 BGP SEND MP capability AFI=1, SAFI=128 Jun 10 23:51:15.058312 BGP SEND Refresh capability, code=128 Jun 10 23:51:15.058331 BGP SEND Refresh capability, code=2 Jun 10 23:51:15.058386 BGP SEND Restart capability, code=64, time=120, flags= Jun 10 23:51:15.058416 BGP SEND 4 Byte AS-Path capability (65), as_num 65001 Jun 10 23:51:15.058651 bgp_pp_recv:3140: NOTIFICATION sent to 192.168.133.0 (External AS 300): code 6 (Cease) subcode 7 (Connection collision resolution), Reason: dropping 192.168.133.0 (External AS 300), connection collision prefers 192.168.133.0+53170 (proto) Jun 10 23:51:15.058680 bgp_senD. sending 21 bytes to 192.168.133.0 (External AS 300) Jun 10 23:51:15.058702
Jun 10 23:51:15.058702 BGP SEND 192.168.133.1+50083 -> 192.168.133.0+179 Jun 10 23:51:15.058735 BGP SEND message type 3 (Notification) length 21 Jun 10 23:51:15.058755 BGP SEND Notification code 6 (Cease) subcode 7 (Connection collision resolution)
Jun 10 23:51:15.059557 advertising receiving-speaker only capabilty to neighbor 192.168.133.0 (External AS 300)
Jun 10 23:51:15.059594 bgp_senD. sending 59 bytes to 192.168.133.0 (External AS 300) Jun 10 23:51:15.059617
Jun 10 23:51:15.059617 BGP SEND 192.168.133.1+179 -> 192.168.133.0+53170 Jun 10 23:51:15.059649 BGP SEND message type 1 (Open) length 59 Jun 10 23:51:15.059671 BGP SEND version 4 as 65001 holdtime 90 id 10.200.1.1 parmlen 30 Jun 10 23:51:15.059691 BGP SEND MP capability AFI=1, SAFI=128 Jun 10 23:51:15.059709 BGP SEND Refresh capability, code=128 Jun 10 23:51:15.059727 BGP SEND Refresh capability, code=2 Jun 10 23:51:15.059747 BGP SEND Restart capability, code=64, time=120, flags= Jun 10 23:51:15.059768 BGP SEND 4 Byte AS-Path capability (65), as_num 65001 Jun 10 23:51:15.060383 bgp_process_caps: mismatch NLRI with 192.168.133.0 (External AS 300):
peer: (1) us: (4)

Jun 10 23:51:15.060445 bgp_process_caps:2578: NOTIFICATION sent to 192.168.133.0 (External AS 300): code 2 (Open Message Error) subcode 7 (unsupported capability) value 4 Jun 10 23:51:15.060470 bgp_senD. sending 23 bytes to 192.168.133.0 (External AS 300) Jun 10 23:51:15.060492
Jun 10 23:51:15.060492 BGP SEND 192.168.133.1+179 -> 192.168.133.0+53170 Jun 10 23:51:15.060556 BGP SEND message type 3 (Notification) length 23 Jun 10 23:51:15.060578 BGP SEND Notification code 2 (Open Message Error) subcode 7 (unsupported capability)
Jun 10 23:51:15.060600 BGP SEND Data (2 bytes): 00 04
-- Exhibit --

Click the Exhibit button.

Referring to the exhibit, what is causing the IPv4 BGP peering to stay in an active state?

A. The peer AS is incorrect.
B. The peer does not support 4-byte AS values.
C. The peer has an NLRI mismatch.
D. The peer has an incorrect IP address.

Correct Answer: C

QUESTION 20
-- Exhibit --
user@router> show route protocol bgp detail

inet.0: 20 destinations, 20 routes (19 active, 0 holddown, 1 hidden) 10.222.1.3/32 (1 entry, 1 announced)
*BGP Preference: 170/-101
Next hop type: Indirect
Address: 0x15ec944
Next-hop reference count: 3
Source: 1.1.1.1
Next hop type: Router, Next hop index: 536
Next hop: 1.1.1.1 via ge-0/0/1.0, selected
Protocol next hop: 1.1.1.1
Indirect next hop: 14081d0 262142
State:
Local AS: 65222 Peer AS: 65221
Age: 2:12 MetriC. 1 Metric2: 0
Task: BGP_65221.1.1.1.1+56417
Announcement bits (2): 0-KRT 4-Resolve tree 1
AS path: 65221 I
Communities: no-advertise
Accepted
Localpref: 100
Router ID. 10.222.1.1
-- Exhibit --

Click the Exhibit button.

You are troubleshooting a problem where an EBGP route is not being advertised to your local IBGP peers. You have received a 10.222.1.3/32 route from an EBGP peer as shown in the exhibit, but the route is not being advertised.

What is causing the problem?

A. The route shows as a hidden route and cannot be advertised.
B. The next hop for the route is indirect and prevents the route from being advertised.
C. The community prevents the route from being advertised.
D. The local preference value is too high for the route to be advertised.

Correct Answer: C

Trying our product !

★ **100%** Guaranteed Success

★ **100%** Money Back Guarantee

★ **365 Days** Free Update

★ **Instant Download** After Purchase

★ **24x7** Customer Support

★ Average **99.9%** Success Rate

★ More than **69,000** Satisfied Customers Worldwide

★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:

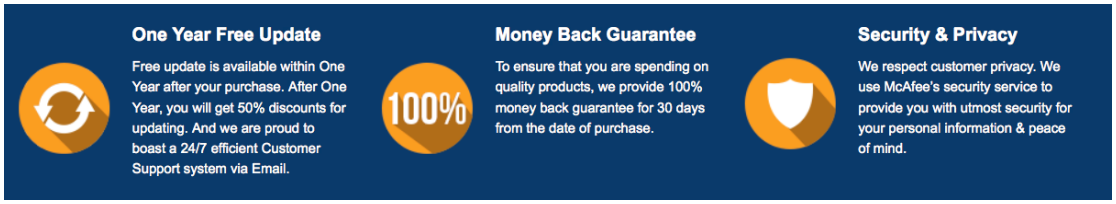

Guarantee & Policy | Privacy & Policy | Terms & Conditions