Dumps with PDF and VCE（+Free VCE Software）

Juniper

Exam JN0-692

Service Provider Routing and Switching Support, Professional

Version: 4.0

[Total Questions: 171]

http://www.ensurepass.com/

Topic break down

Topic	No. of Questions
Topic 1: Volume A	70
Topic 2: Volume B	69
Topic 3: Volume C	32

Topic 1, Volume A

Question No : 1 - (Topic 1)

What is a limitation of LDP?

A. Traffic must follow explicitly configured paths.
B. It requires a full mesh of LSPs throughout the network.
C. It requires a traffic engineering database (TED).
D. It does not support traffic engineering.

Answer: D

Question No : 2 - (Topic 1)

Which two statements correctly describe BGP operation? (Choose two.)

A. IBGP does not advertise routes learned from other IBGP neighbors.
B. IBGP advertises routes learned from other IBGP neighbors.
C. EBGP advertises routes learned from other IBGP or EBGP neighbors.
D. EBGP does not advertise routes learned from other EBGP neighbors.

Answer: A,C

Question No : 3 - (Topic 1)

Click the Exhibit button.

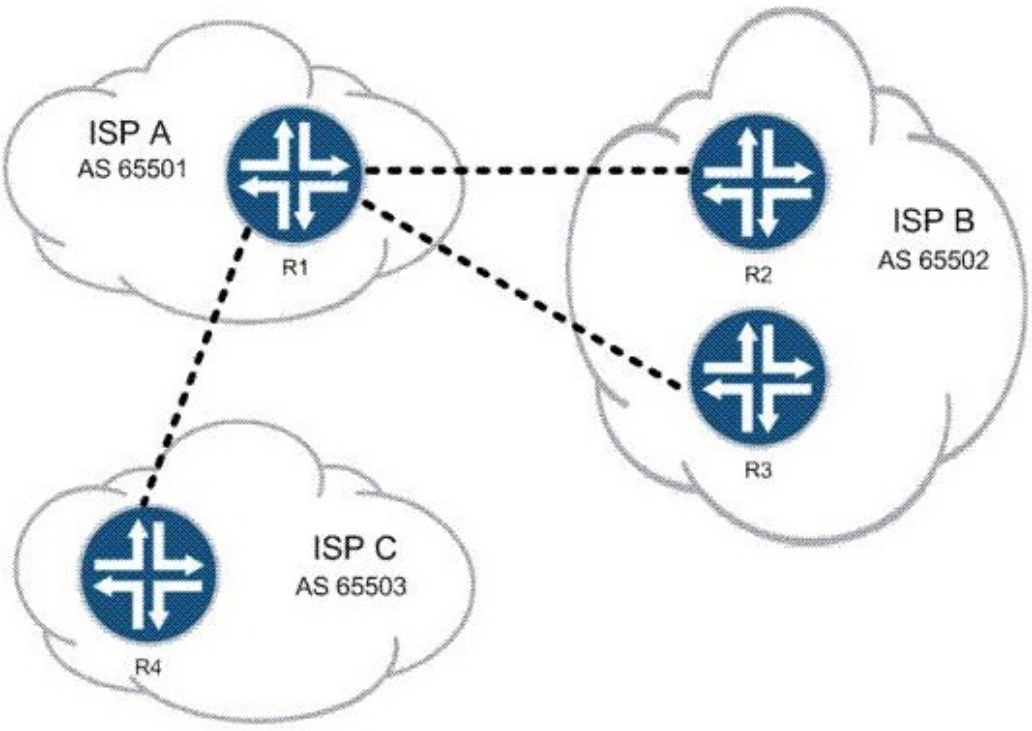

You work for ISP A. Customers of both ISP B and ISP C must be able to reach all of your customers, but your network must not allow transit traffic between ISP B and ISP C.

Referring to the exhibit, which two methods could you use? (Choose two.)

A. Use local preference to prefer the proper routes.
B. Use the well-known no-transit community.
C. Use policy to filter routes on AS number.
D. Use communities to identify and filter routes.

Answer: C,D

Question No : 4 - (Topic 1)

Click the Exhibit button.

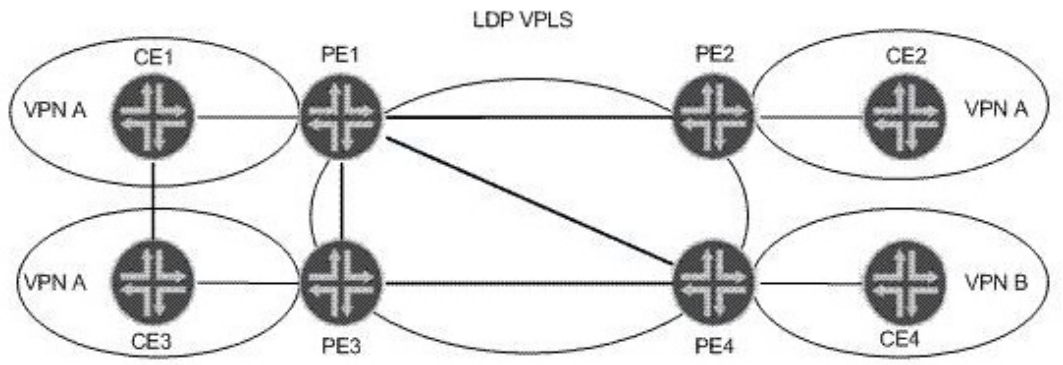

Your IT manager asks you to describe a benefit of migrating from LDP VPLS towards BGP VPLS considering the operational network shown in the exhibit. What can you tell your manager?

A. Using BGP signaling improves scaling, because a full mesh of transport LSPs is not needed.
B. MAC addresses are learned through BGP instead of LDP, which improves scaling.
C. Ingress PE replication can be reduced, because BGP VPLS supports P2MP LSPs.
D. Configuration overhead is reduced when adding new sites or new VPNs.

Answer: D

Question No : 5 - (Topic 1)

You have configured bidirectional LSPs between R1 and R3. You notice that the LSPs are not coming up.

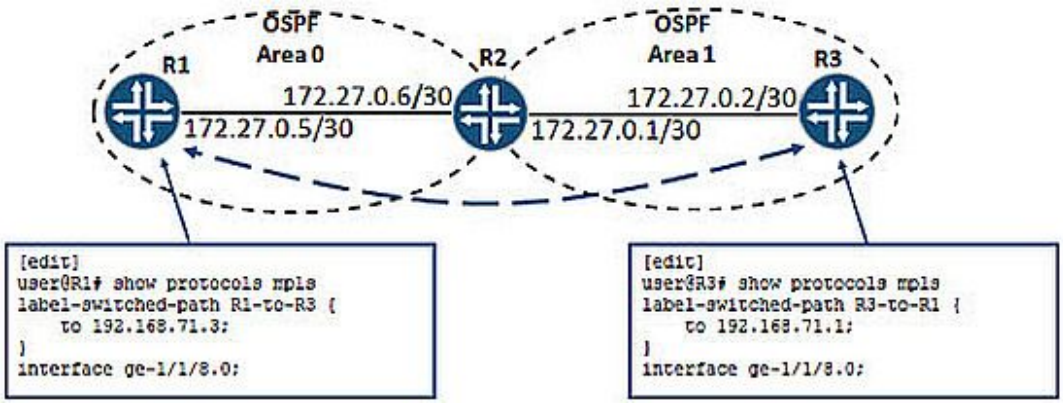

Referring to the exhibit, which statement is correct?

A. You must disable CSPF for the LSPs.
B. You must enable CSPF for the LSPs.
C. You must ensure type 5 LSAs are allowed in area 1.
D. You must ensure OSPF traffic engineering is configured.

Answer: A

Question No : 6 - (Topic 1)

```
user@PE-1# show routing-instances Cust
instance-type vrf;
interface ge-0/0/1.0;
route-distinguisher 5:5;
vrf-import vrf-imp;
vrf-target target:1:1;
protocols {
    bgp {
        group cust {
            peer-as 100;
            as-override;
            neighbor 204.56.78.1;
        }
    }
}

user@PE-1# show policy-options policy-statement vrf-imp
term 1 {
    from {
        protocol direct;
        community target;
    }
    then accept;
}
term 2 {
    then reject;
}

user@PE-1# show policy-options community target
members target:1:1;
```

PE-1 is not sending any routes to the CE router in the Cust routing instance.

What is the reason?

A. There is a mismatch in the target community and the route distinguisher.
B. The vrf-import policy is only accepting direct routes.

C. The as override parameter is causing an AS loop on the CE device.
D. The vrf-export policy must be explicitly defined.

Answer: B

Question No : 7 - (Topic 1)

What is the purpose of the no-cspf command?

A. to successfully signal the LSP across the network regardless of constraints
B. to delete the CSPF database
C. to ignore OSPF when calculating the ERO
D. to successfully signal the LSP only if the default IGP path (or named path) meets all constraints

Answer: D

Question No : 8 - (Topic 1)

You recently added your autonomous system to an existing BGP confederation. You notice that a route that had a local preference of 100 now has a local preference of 50. Which statement explains the change?

A. BGP path attributes such as next hop, local preference, and MED are normally restricted to a single AS but are allowed to propagate throughout the confederation's AS members.
B. The confederation has sub-ASs that require all IBGP routes to have a local preference of 50 or below.
C. When your Junos devices joined the confederation, they lost IBGP connectivity to the route in question; the local preference reverted to 50 once the BGP peering established.
D. The route is being shared through an EBGP peer, and the confederation is propagating the local preference from the peer.

Answer: A

Question No : 9 - (Topic 1)

Click the Exhibit button.

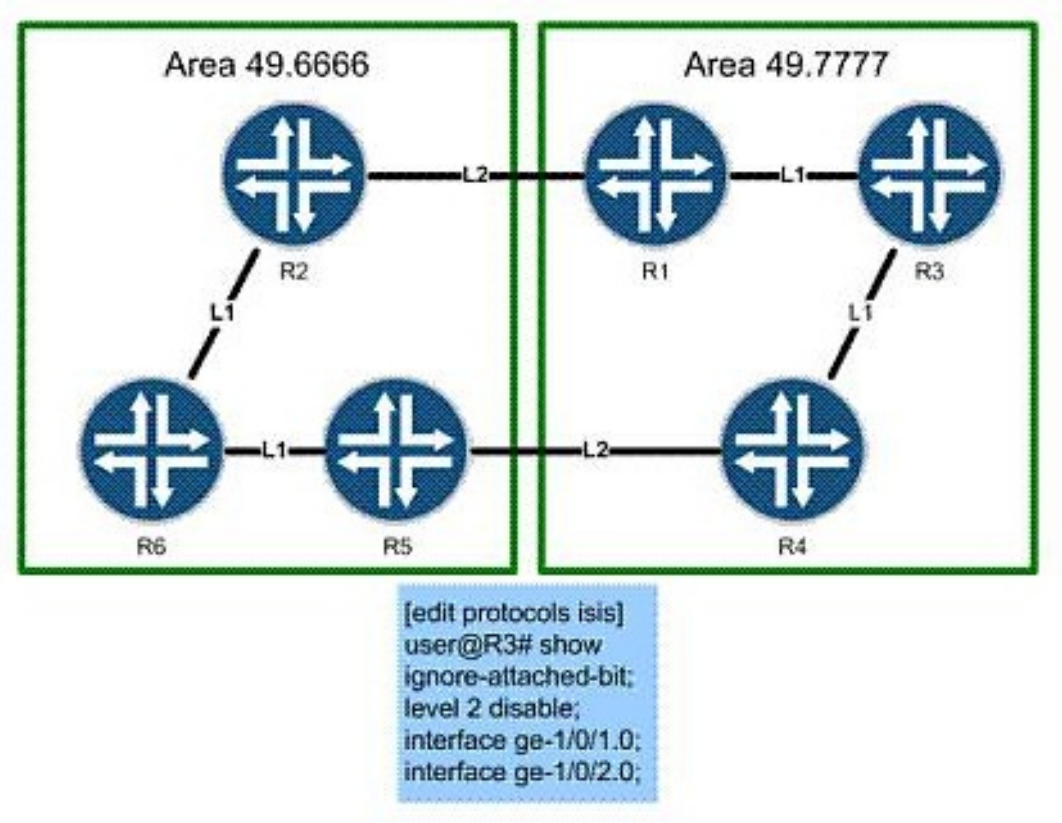

Based on the exhibit, what do you expect to find in the configuration on R1 and R4?

A. a policy leaking level 1 routes into level 2
B. a policy leaking level 2 routes into level 1
C. a policy setting the attached bit on level 2 routes
D. a policy setting the attached bit on level 1 routes

Answer: B

Question No : 10 - (Topic 1)

IS-IS is configured to support both IPv4 and IPv6 routing. Which statement is true?

A. Separate IPv4 and IPv6 hellos will be sent.
B. IPv6 will have a separate link-state database.
C. IS-IS v6 support must be enabled under protocols isis.
D. IS-IS sends IPv6 topology information as new TLVs in existing LSPs.

Answer: D

Question No : 11 - (Topic 1)

Click the Exhibit button.

```
user@router# run show class-of-service rewrite-rule name traffic-class
Rewrite rule: traffic-class, Code point type: exp, Index: 58855
  Forwarding class                Loss priority       Code point
  best-effort                     low                 000
  best-effort                     high                001
  expedited-forwarding            low                 111
  expedited-forwarding            high                011
  assured-forwarding              low                 100
  assured-forwarding              high                101
  network-control                 low                 110
  network-control                 high                111
```

Your router should be configured with a rewrite rule which alters the default behavior of expedited-forwarding as shown in the exhibit. Which configuration is correct?

A.
```
[edit]
user@router# show class-of-service
rewrite-rules {
    exp traffic-class {
        import default;
        forwarding-class expedited-forwarding {
            loss-priority low code-point 111;
        }
    }
}
```

B.
```
[edit]
user@router# show class-of-service
rewrite-rules {
    exp traffic-class {
        import rewrite-rule best-effort;
        import rewrite-rule expedited-forwarding;
        import rewrite-rule assured-forwarding;
        import rewrite-rule network-control;
        forwarding-class expedited-forwarding {
            loss-priority low code-point 111;
        }
    }
}
```

C.
```
[edit]
user@router# show class-of-service
rewrite-rules {
    exp traffic-class {
        import best-effort;
        import assured-forwarding;
        import network-control;
        forwarding-class expedited-forwarding {
            loss-priority low code-point 111;
        }
    }
}
```

D.
```
[edit]
user@router# show class-of-service
rewrite-rules {
    exp traffic-class {
        import best-effort;
        import assured-forwarding;
        import expedited-forwarding;
        import network-control;
        }
    }
}
```

A. Option A
B. Option B
C. Option C

D. Option D

Answer: A

Question No : 12 - (Topic 1)

You are facing BGP scaling issues and decide to add dedicated route reflectors to your network. You notice that VPN routes are not being advertised by your route reflectors. Which three actions can you take to solve this? (Choose three.)

A. Add a static default route to inet.3 and/or inet6.3 on the route reflectors.
B. Add a full mesh of MPLS LSPs between all of the route reflectors.
C. Add MPLS LSPs between the route reflectors and their client routers.
D. Add a static default route to inet.3 and/or inet6.3 on all of the client routers.
E. Use rib-groups to add IGP routes to inet.3 and/or inet6.3 on the route reflectors.

Answer: A,C,E

Question No : 13 - (Topic 1)

```
user@R3> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 2.2.2.2
  Flags: sparse,rptree,wildcard
  Upstream interface: xe-5/3/0.0
  Upstream neighbor: 17.2.1.2
  Upstream state: None
  Uptime: 00:23:49
  Downstream neighbors:
    Interface: xe-2/1/0.0
      17.4.1.3 State: Join Flags: SRW Timeout: Infinity
      Uptime: 00:23:49 Time since last Join: 00:01:17

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

user@R3> show multicast route
Instance: master Family: INET

Instance: master Family: INET6
```

Referring to the exhibit, which statement is true about the multicast state on R3?

A. R3 has no PIM neighbors.
B. R3 has no receivers.
C. R3 has no route to RP.
D. R3 is not receiving any traffic

Answer: D

Question No : 14 - (Topic 1)

```
user@router> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped : 2006-02-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate : 0 bps (0 pps)
  Output rate : 0 bps (0 pps)
  Active alarms : None
  Active defects : None
  Logical interface ge-0/0/0.0 (Index 66) (SNMP ifIndex 198)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
...

user@router> show ospf neighbor
Address       Interface           State       ID           Pri    Dead
1.1.1.2       ge-0/0/0.0          Exstart     2.2.2.2      128    36

user@router> show log trace_ospf
Apr 24 12:19:01 Version 2, length 48, ID 1.1.1.2, area 0.0.0.0
Apr 24 12:19:01 checksum 0xbd12, authtype 0
Apr 24 12:19:01 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 128
Apr 24 12:19:01 dead_ivl 40, DR 2.2.2.2, BDR 2.2.2.3
Apr 24 12:19:01 checksum 0x66a2, authtype 0
Apr 24 12:19:01 options 0x42, i 1, m 1, ms 1, seq 0xa0f3843, mtu 1500
Apr 24 12:19:01 OSPF now slave for nbr 10.0.8.1
Apr 24 12:19:01 options 0x42, i 1, m 1, ms 1, seq 0xa04c360, mtu 1500
Apr 24 12:19:01 options 0x42, i 0, m 0, ms 0, seq 0xa04c360, mtu 2986
Apr 24 12:19:01 OSPF packet ignored: MTU mismatch from 2.2.2.2 on intf ge-0/0/0.0 area 0.0.0.0
Apr 24 12:19:01 Version 2, length 48, ID 1.1.1.2, area 0.0.0.0
Apr 24 12:19:01 checksum 0xbd12, authtype 0
Apr 24 12:19:01 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 128
Apr 24 12:19:01 dead_ivl 40, DR 2.2.2.2, BDR 2.2.2.3
Apr 24 12:19:01 checksum 0x66a2, authtype 0
Apr 24 12:19:01 options 0x42, i 1, m 1,
Apr 24 12:19:01 options 0x42, i 1, m 1, ms 1, seq 0xa04c360, mtu 1500
Apr 24 12:19:01 options 0x42, i 0, m 0, ms 0, seq 0xa04c360, mtu 2986
```

You have been asked to troubleshoot an OSPF problem where the OSPF session will not establish.

According to the outputs shown in the exhibit, which statement is true?

A. The hold-time interval is set too low.
B. There is an MTU mismatch.
C. The hold-time interval is set too high.
D. The dead-interval is set too low.

Answer: B

Question No : 15 - (Topic 1)

Click the Exhibit button.

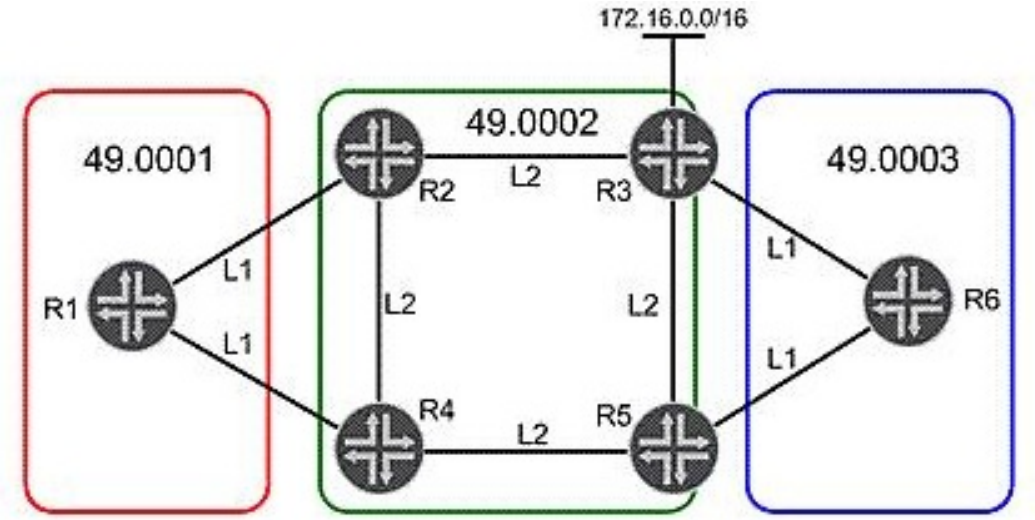

The IGP is IS-IS. Routes from R1 need to be present on R6. Referring to the exhibit, what will accomplish this task?

A. Create an L1 adjacency between R2 and R3 to allow the routes to pass through to R6.
B. Use policy on R3 to leak R1's routes from L2 to L1.
C. Change the area address from 49.0003 to 49.0001 on R6 to allow R6 to accept routes from R1.
D. Use policy on R2 to leak R1's routes from L1 to L2.

Answer: B

Question No : 16 - (Topic 1)

You have assigned target:65432:100 as the route target for Customer A's BGP Layer 2 VPN. The PE1 router VRF is configured with vrf-target export target:65432:100.Which configuration on PE2 correctly assigned Customer A's routes to their VRF?

A. vrf-target target:65432:100
B. route-target target:65432:100
C. vrf-target export target:65432:100
D. route-target export target:65432:100

Answer: A

Question No : 17 - (Topic 1)

Click the Exhibit button.

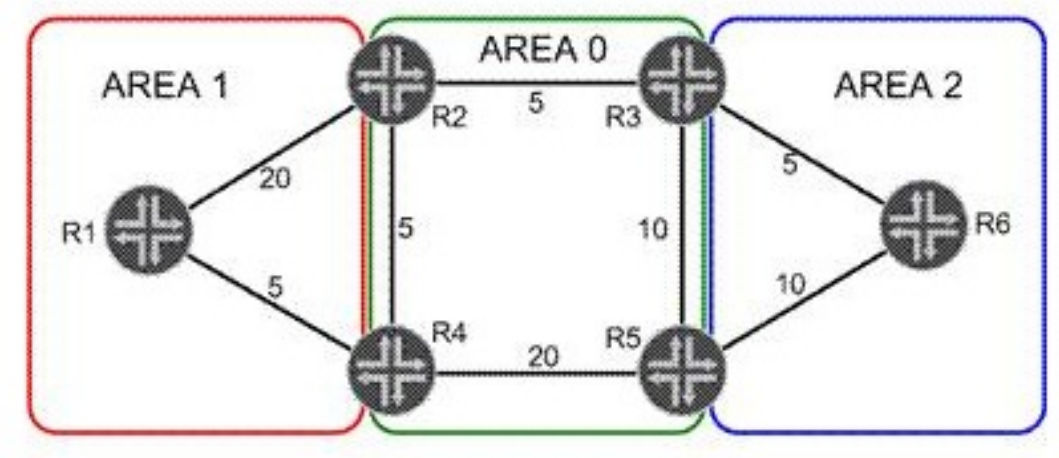

Referring to the OSPF link metrics in the exhibit, which path will traffic from R6 take to reach R1?

A. R6, R3, R2, R4, R1
B. R6, R3, R2, R1
C. R6, R5, R4, R1
D. R6, R5, R3, R2, R4, R1

Answer: B

Question No : 18 - (Topic 1)

Click the Exhibit button.

```
[edit]
user@router# show firewall
policer policerA {
    logical-interface-policer;
    if-exceeding {
        bandwidth-limit 10m;
        burst-size-limit 500k;
    }
    then discard;
}

[edit]
user@router# show interfaces
ge-0/0/2 {
    unit 0 {
        family inet {
            policer {
                input policerA;
            }
        }
        family inet6 {
            policer {
                input policerA;
            }
        }
    }
    unit 1 {
        family inet {
            policer {
                input policerA;
            }
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            policer {
                input policerA;
            }
        }
        family inet6 {
            policer {
                input policerA;
            }
        }
```

http://www.ensurepass.com/

Traffic is flowing through the interfaces in the exhibit as follows:

On ge-0/0/2.0, IPv4 traffic has a throughput rate of 4 Mbps, and the burst size counter is at 200 KB.

On ge-0/0/2.0, IPv6 traffic has a throughput rate of 7 Mbps, and the burst size counter is at 550 KB.

On ge-0/0/3.0, IPv4 traffic has a throughput rate of 5 Mbps, and the burst size counter is at 250 KB.

On ge-0/0/3.1, IPv6 traffic has a throughput rate of 12 Mbps, and the burst size counter is at 450 KB.

Which statement describes what is happening?

A. IPv6 traffic on ge-0/0.3.1 is being dropped; all other traffic is unaffected.
B. IPv4 traffic on ge-0/0/2.0 is unaffected; IPv6 traffic on ge-0/0/2.0 is being dropped; IPv4 traffic on ge-0/0/3.0 is unaffected; IPv6 traffic on ge-0/0/3.1 is being dropped.
C. IPv4 traffic on ge-0/0/2.0 is being dropped; IPv6 traffic on ge-0/0/2.0 is being dropped; IPv4 traffic on ge-0/0/3.0 is unaffected; IPv6 traffic on ge-0/0/3.1 is unaffected.
D. All IPv4 and IPv6 traffic on ge-0/0/2 and ge-0/0/3 is being dropped.

Answer: B

Question No : 19 - (Topic 1)

R1 and R2 are IS-IS Level 1 routers, but are not forming an adjacency.

Which two reasons would account for this happening? (Choose two.)

A. The remote router is configured for both IPv4 and IPv6.
B. The remote router does not have family iso enabled on its interface,
C. The remote router does not have family clns enabled on its interface.
D. The remote router is not configured to participate in the same Level 1 area

Answer: B,D

Question No : 20 - (Topic 1)

```
Feb 9 16:58:56.580693 task_reset_socket: task BGP_66666_65001.10.10.17.253+179 socket 35
Feb 9 16:58:56.580966 bgp_event: peer 10.10.17.253 (External AS 66666) old state Connect event OpenFail new state Idle
Feb 9 16:58:56.581592 bgp_event: peer 10.10.17.253 (External AS 66666) old state Idle event Start new state Active
Feb 9 17:00:07.868665 task_timer_dispatch: calling BGP_66666_65001.10.10.17.253_Connect, late by 0.004
Feb 9 17:00:07.869042 bgp_connect_timeout: BGP_66666_65001.10.10.17.253_Connect
Feb 9 17:00:07.869158 bgp_connect_start: peer 10.10.17.253 (External AS 66666)
Feb 9 17:00:07.869252 bgp_event: peer 10.10.17.253 (External AS 66666) old state Active event ConnectRetry new state
Connect
Feb 9 17:00:07.869951 task_get_socket: domain AF_INET type SOCK_STREAM protocol 0 socket 35
Feb 9 17:00:07.870180 task_set_socket: task BGP_66666_65001.10.10.17.253 socket 35
Feb 9 17:00:07.870398 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option NonBlocking(8) value 1
Feb 9 17:00:07.871810 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option ReUseAddress(3) value 1
Feb 9 17:00:07.872041 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option PathMTUDiscovery(26)
value 0
Feb 9 17:00:07.872206 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option RoutingTable(27) value 5
Feb 9 17:00:07.872347 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option TOS(16) value 192
Feb 9 17:00:07.872478 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option DontRoute(5) value 1
Feb 9 17:00:07.872605 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option IifRestrict(36) value 1
Feb 9 17:00:07.874664 task_set_option_internal: task BGP_66666_65001.10.10.17.253 socket 35 option TTL(15) value 1
Feb 9 17:00:07.874967 task_addr_local: task BGP_66666_65001.10.10.17.253 address 10.10.17.238
Feb 9 17:00:07.875442 task_connect: task BGP_66666_65001.10.10.17.253+179 addr 10.10.17.253+179task_timer_reset: reset BGP_
66666_65001.10.10.17.253+179_Connect
Feb 9 17:00:07.875634 task_timer_set_oneshot_latest: timer BGP_66666_65001.10.10.17.253+179_Connect interval set to 2:28
Feb 9 17:00:07.875757 task_timer_dispatch: returned from BGP_66666_65001.10.10.17.253+179_Connect, rescheduled in 2:28
Feb 9 17:01:24.602351 task_process_events: connect ready for BGP_66666_65001.10.10.17.253+179
Feb 9 17:01:24.602915 bgp_connect_complete: error connecting to 10.10.17.253 (External AS 66666): Socket is not connected
Feb 9 17:01:24.603030 bgp_close_socket: peer 10.10.17.253 (External AS 66666)
Feb 9 17:01:24.603158 task_close: close socket 35 task BGP_66666_65001.10.10.17.253+179
Feb 9 17:01:24.603255 task_reset_socket: task BGP_66666_65001.10.10.17.253+179 socket 35
Feb 9 17:01:24.603484 bgp_event: peer 10.10.17.253 (External AS 66666) old state Connect event OpenFail new state Idle
Feb 9 17:01:24.604095 bgp_event: peer 10.10.17.253 (External AS 66666) old state Idle event Start new state Active
Feb 9 17:02:35.923738 task_timer_dispatch: calling BGP_66666_65001.10.10.17.253_Connect, late by 0.055
Feb 9 17:02:35.923964 bgp_connect_timeout: BGP_66666_65001.10.10.17.253_Connect
```

You are asked to configure an EBGP peering, but the peering does not establish and transitions between the Connect and Active states.

Referring to the exhibit, what is causing the problem?

A. The number of messages from the BGP peer has exceeded the receive socket buffer.
B. An incorrect local AS has been configured in the BGP group configuration.
C. BGP is waiting to receive a keepalive or notification message from the peer.
D. The BGP peer has an incorrect address configured, causing a connectivity failure.

Answer: D

Question No : 21 - (Topic 1)

Which statement is true about ASM and/ or SSM multicast?

A. ASM requires an external mechanism to find the source.
B. SSM only builds RPT trees, since the RP is replaced by an external mechanism.
C. ASM and SSM for IPv6 multicast use embedded RP.
D. SSM does not require MSDP.

Answer: D

Question No : 22 - (Topic 1)

You have been asked to make a configuration which inherits the statements in a predefined configuration group. What will accomplish this?

A.
```
groups {
    group-name {
        configuration-data;
    }
}
```

B.
```
apply-groups <apply-group-name>;
```

C.
```
apply-macro <apply-macro-name>;
```

D.
```
event-options {
    event-script {
        file file-name;
    }
}
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Question No : 23 - (Topic 1)

Click the Exhibit button.

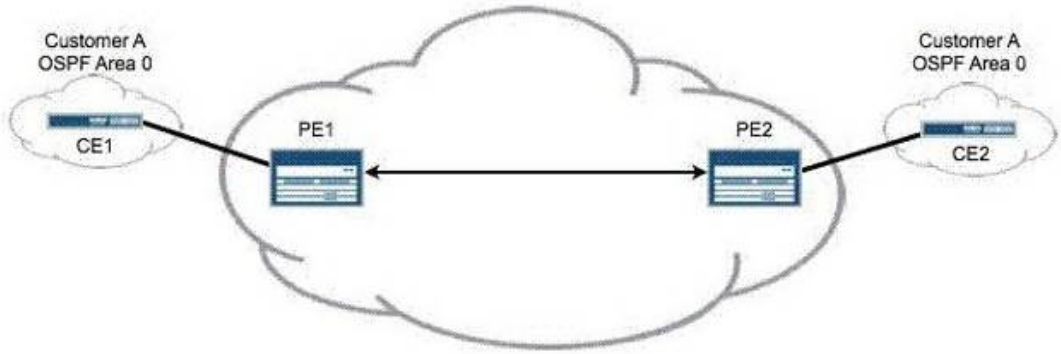

Referring to the exhibit, your network management systems have alerted you to a loss of connectivity to the CE2 router in your Layer 3 VPN. The loopback address of the CE router is 10.10.1.1/32.Which operational command on PE2 verifies connectivity across the PE-CE link?

A. ping 10.10.1.1
B. ping 10.10.1.1 table customer-a
C. ping 10.10.1.1 instance customer-a
D. ping 10.10.1.1 routing-instance customer-a

Answer: D

Question No : 24 - (Topic 1)

```
user@PE> show configuration routing-instances
CUST {
    instance-type vrf;
    interface xe-5/2/0.0;
    route-distinguisher 1:1;
    vrf-target target:1:1;
    vrf-table-label;
}

user@PE> show configuration interfaces xe-1/3/0
unit 0 {
    description CORE-INTERFACE;
    family inet {
        address 17.2.1.2/24;
    }
    family mpls;
}

user@PE> show configuration class-of-service classifiers
dscp DIFFSERVCLASS {
    forwarding-class expedited-forwarding {
        loss-priority low code-points 101110;
    }
}
exp DIFFSERVCLASS-EXP {
    forwarding-class expedited-forwarding {
        loss-priority low code-points 101;
    }
}
```

Referring to the exhibit, which configuration is required on the Layer 3 VPN PE router to classify EF packets from the core MPLS network into the expedited-forwarding queue?

A. Set class-of-service interfaces xe-1/3/0 unit 0 classifiers dscp default
B. Set class-of-service interfaces xe-1/3/0 unit 0 classifiers dscp DIFFSERVCLASS
C. Set class-of-service interfaces xe-1/3/0 unit 0 classifiers exp default
D. Set class-of-service routing-instances CUST classifiers exp DIFFSERVCLASS-EXP

Answer: D

Question No : 25 - (Topic 1)

You have recently implemented a new Layer 2 VPN and it is not establishing. The show l2vpn connections command displays an out-of-range state.

What is causing this problem?

A. The local site ID value is too high for the label block.
B. The VRF target community values do not match between sites.
C. The VRF route distinguisher values do not match between sites.
D. The local site ID is not defined.

Answer: A

Question No : 26 - (Topic 1)

```
user@router> show mpls lsp extensive
Ingress LSP: 1 sessions

4.4.4.4
  From: 1.1.1.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R4
  ActivePath: pathR4 (primary)
  Node/Link protection desired
  LSPtype: Static Configured
  LoadBalance: Least-fill
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary pathR4     State: Up, No-decrement-ttl
    Priorities: 7 0
    OptimizeTimer: 900
    SmartOptimizeTimer: 180
      Include Any: JUNOS
    Reoptimization in 51 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
  17.1.1.2 S 17.2.1.3 S 17.4.1.4 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
       2.2.2.2(flag=0x2b) 17.1.1.2(flag=0xb Label=299856) 17.3.1.4
  21 May 21 14:11:14.236 CSPF failed: no route toward 4.4.4.4
  20 May 21 14:11:14.236 17.1.1.2: Tunnel local repaired
  19 May 21 14:10:29.235 CSPF failed: no route toward 4.4.4.4
  18 May 21 14:10:29.234 17.1.1.2: Tunnel local repaired
  17 May 21 14:09:44.234 CSPF failed: no route toward 4.4.4.4
  16 May 21 14:09:44.234 17.1.1.2: Tunnel local repaired
  15 May 21 14:08:59.230 CSPF failed: no route toward 4.4.4.4
  14 May 21 14:08:59.229 17.1.1.2: Tunnel local repaired
  13 May 21 14:08:55.123 CSPF failed: no route toward 4.4.4.4
  12 May 21 14:08:55.123 17.1.1.2: Tunnel local repaired[2 times]
  11 May 21 14:08:53.227 Record Route: 2.2.2.2(flag=0x2b) 17.1.1.2(flag=0xb Label=299856) 17.3.1.4
  10 May 21 14:08:50.433 CSPF failed: no route toward 4.4.4.4
   9 May 21 14:08:50.433 CSPF: link down/deleted: 0.0.0.0(17.2.1.3:0)(17.2.1.3)->0.0.0.0(3.3.3.3:0)(3.3.3.3)
   8 May 21 14:08:50.371 CSPF failed: no route toward 4.4.4.4
   7 May 21 14:08:50.371 CSPF: link down/deleted: 17.2.1.2(2.2.2.2:0)(2.2.2.2)->0.0.0.0(17.2.1.3:0)(17.2.1.3)
   6 May 21 14:01:43.347 Selected as active path
   5 May 21 14:01:43.346 Record Route: 2.2.2.2(flag=0x29) 17.1.1.2(flag=9 Label=299856) 3.3.3.3(flag=0x21) 17.2.1.3(flag=1 Label=299840)
4.4.4.4(flag=0x20) 17.4.1.4(Label=3)
   4 May 21 14:01:43.346 Up
   3 May 21 14:01:43.318 Originate Call
   2 May 21 14:01:43.318 CSPF: computation result accepted 17.1.1.2 17.2.1.3 17.4.1.4
  Created: Tue May 21 14:01:14 2013
Total 1 displayed, Up 1, Down 0
```

Referring to the exhibit, which two statements are correct about the LSP? (Choose two.)

A. The primary LSP flapped following a link flap.
B. The LSP will go down after 51 seconds.
C. The head-end is signaled that the LSP is protected following a link flap.
D. The LSP constraint checks are failing following a link flap.

Answer: C,D

Question No : 27 - (Topic 1)

Click the Exhibit button.

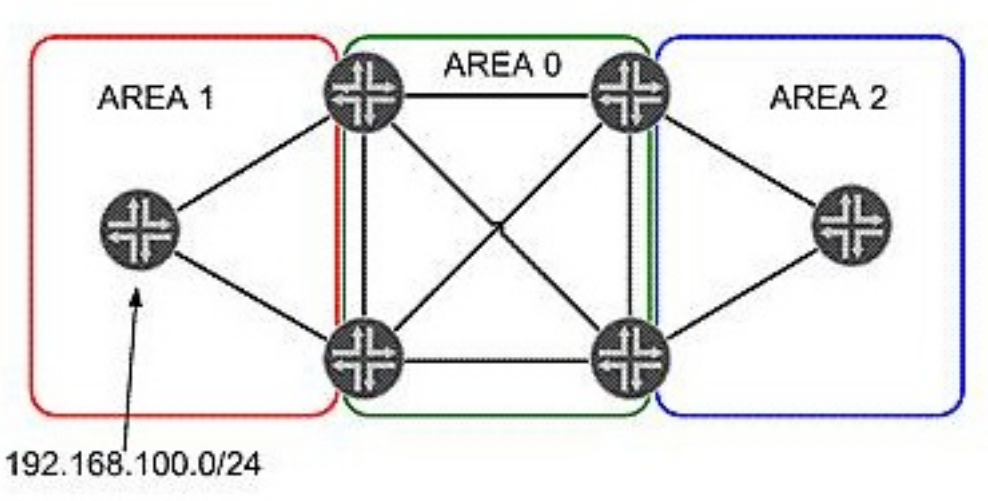

In the exhibit, Area 1 and Area 2 are configured as not-so-stubby areas. RIP network 192.168.100.0/24 is redistributed into OSPF in Area 1.Which three statements are true? (Choose three.)

A. Network 192.168.100.0/24 is advertised in a Type 7 LSA in Area 1.
B. Network 192.168.100.0/24 is advertised in a Type 7 LSA in Area 0.
C. Network 192.168.100.0/24 is advertised in a Type 5 LSA in Area 0.
D. The area border router between Area 0 and Area 2 converts network 192.168.100.0/24 to a Type 7 LSA.
E. Area 2 does not see the network 192.168.100.0/24 in its link-state database.

Answer: A,C,E

Question No : 28 - (Topic 1)

Click the Exhibit button.

```
customer-vpn {
    instance-type vrf;
    interface ge-0/0/0.0;
    route-distinguisher 172.16.1.1:1;
    vrf-target target:65000:100;
}
```

You are configuring a new PE router in your Layer 3 VPN. A remote PE router is using the configuration shown in the exhibit. Which configuration is needed to receive customer-vpn routes from the remote PE?

A.
```
customer-vpn {
    instance-type vrf;
    interface ge-0/0/1.0;
    route-distinguisher 172.16.1.2:1;
    vrf-target {
        export target:65000:100;
        import target:65000:200;
    }
}
```

B.
```
customer-vpn {
    instance-type vrf;
    interface ge-0/0/1.0;
    route-distinguisher 172.16.1.2:1;
    vrf-target {
        export target:65000:200;
        import target:65000:200;
    }
}
```

C.
```
customer-vpn {
    instance-type vrf;
    interface ge-0/0/1.0;
    route-distinguisher 172.16.1.2:1;
    vrf-target target:65000:100;
}
```

D.
```
customer-vpn {
    instance-type vrf;
    interface ge-0/0/1.0;
    route-distinguisher 172.16.1.2:1;
    vrf-target target:65000:200;
}
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

Question No : 29 - (Topic 1)

Click the Exhibit button.

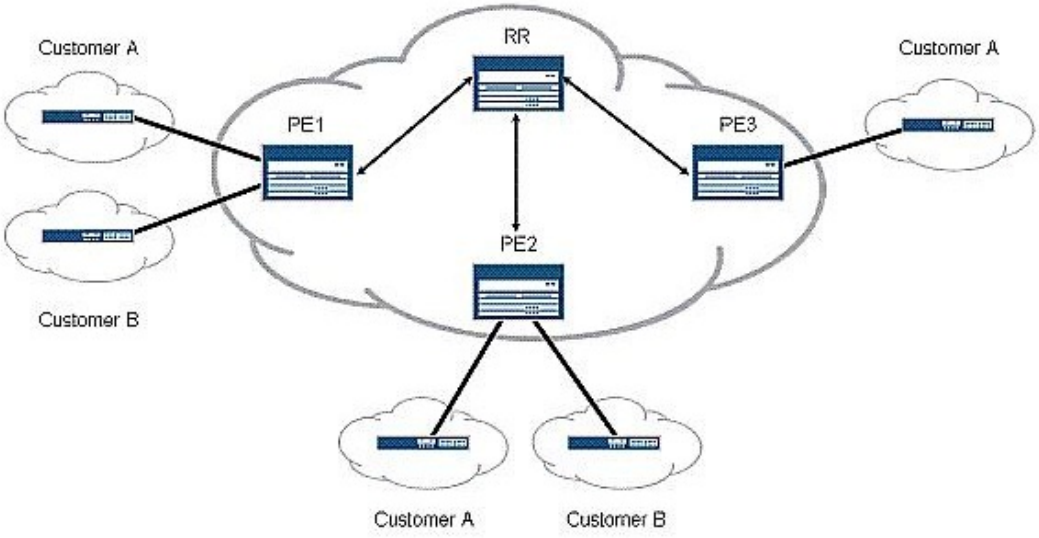

Referring to the exhibit, you want to save CPU processing load on the PE3 router by preventing the reception of routes belonging to Customer B. Which Layer 3 VPN scaling mechanism provides this functionality?

A. route origin
B. route refresh
C. route reflection
D. route target filtering

Answer: D

Question No : 30 - (Topic 1)

Click the Exhibit button.

```
[edit]
user@R4# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence  Checksum  Lifetime  Attributes
R4.00-00                    0x4    0xe888      1154  L1 L2
R3.00-00                    0x3    0x2ce1      1150  L1 L2
R3.02-00                    0x2    0x46c7      1150  L1 L2
  3 LSPs

IS-IS level 2 link-state database:
LSP ID                 Sequence  Checksum  Lifetime  Attributes
R4.00-00                    0x5    0xee7d      1154  L1 L2
R3.00-00                    0x4    0xed1f      1150  L1 L2
R3.02-00                    0x3    0x44c8      1151  L1 L2
  3 LSPs

[edit]
user@R4#
```

Based on the output shown in the exhibit, which statement is correct?

A. R3 is the designated intermediate system.
B. R3 is the backup designated intermediate system.
C. R3 has been configured with an export policy and is announcing external routes to IS-IS neighbors.
D. R3 is using both IPv4 and IPv6 resulting in two pseudonodes.

Answer: A

Question No : 31 - (Topic 1)

Which two statements describe advantages of using BGP for VPLS signaling instead of LDP signaling? (Choose two.)

A. There is no need for MPLS signaling protocol.
B. There is a well-defined scaling hierarchy.
C. There is a separation of signaling from other services.
D. There is auto discovery.

Answer: B,D

Question No : 32 - (Topic 1)

You have recently added LDP-signaled VPLS to your network. The VPLS connections are established and you have been asked to verify that the forwarding plane is working properly. Which three commands would you use? (Choose three.)

A. ping mpls ldp <fec>
B. show vpls connections
C. traceroute mpls ldp <fec>
D. show vpls mac-table
E. show ldp statistics

Answer: A,C,D

Question No : 33 - (Topic 1)

Click the Exhibit button.

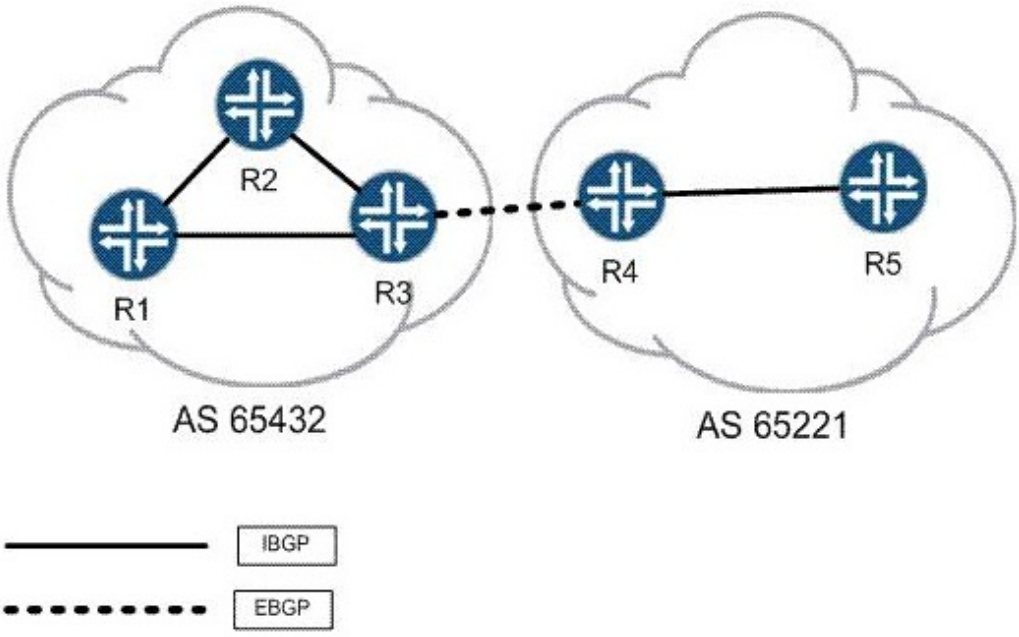

R3 and R4 want to establish an EBGP session between each other's loopback addresses. They have each configured static routes to the other's loopback address and can ping from loopback to loopback. Their EBGP session is configured with correct neighbor and local addresses. The correct AS numbers have been specified at the [routing-options] hierarchy as well. Considering the topology in the exhibit, which statement is true?

A. BGP's protocol preference must be adjusted to be lower than protocol static for the session to establish.

B. Each side must configure multipath for the session to establish.
C. Each peer must specify a local-as within their EBGP configuration for the session to establish.
D. Each peer must configure multihop for the session to establish.

Answer: D

Question No : 34 - (Topic 1)

Click the Exhibit button.

```
user@router> monitor traffic detail interface so-0/1/0 size 1514
Listening on so-0/1/0
11:55:48.470418 In ISIS(186), 30:30:30:30:30:30 > 30:30:30:30:30:30, hlen: 27, v: 1,
    sys-id-len: 6 (0), max-area: 3 (0), L2 LSP
    lsp-id: 1921.6804.8001.00-00, seq: 0x00000008, lifetime: 1189s
    chksum: 0x86c9 (correct), PDU length: 186, L1L2 IS
    Area address(es) TLV #1, length: 4
        Area address (3): 49.0001
    Protocols supported TLV #129, length: 2
        NLPID(s): IPv4, IPv6
    Traffic Engineering Router ID TLV #134, length: 4
        Traffic Engineering Router ID: 192.168.48.1
    IPv4 Interface address(es) TLV #132, length: 4
        IPv4 interface address: 192.168.48.1
    Hostname TLV #137, length: 8
        Hostname: SaoPaulo
    IPv4 Internal reachability TLV #128, length: 24
        IPv4 prefix: 192.168.48.1/32
            Default Metric: 00, Internal, Distribution: up
        IPv4 prefix: 10.222.60.0/24
            Default Metric: 10, Internal, Distribution: up
    Extended IPv4 reachability TLV #135, length: 17
        IPv4 prefix: 192.168.48.1/32
            Metric: 0, Distribution: up, no sub-TLVs present
        IPv4 prefix: 10.222.60.0/24
            Metric: 10, Distribution: up, no sub-TLVs present
    IPv4 External reachability TLV #130, length: 12
        IPv4 prefix: 192.168.49.0/24
            Default Metric: 00, Internal, Distribution: up
    Extended IPv4 reachability TLV #135, length: 8
        IPv4 prefix: 192.168.49.0/24
        Metric: 0, Distribution: up, no sub-TLVs present
    IS Reachability TLV #2, length: 12
        IsNotVirtual
        IPv4 prefix: 192.168.49.0/24
            Default Metric: 00, Internal, Distribution: up
    Extended IPv4 reachability TLV #135, length: 8
        IPv4 prefix: 192.168.49.0/24
        Metric: 0, Distribution: up, no sub-TLVs present
    IS Reachability TLV #2, length: 12
        IsNotVirtual
        IS Neighbor: 1921.6805.2001.00, Default Metric: 10, Internal
    Extended IS Reachability TLV #22, length: 23
        IS Neighbor: 1921.6805.2001.00, Metric: 10, sub-TLVs present (12)
            IPv4 interface address: 10.222.60.2
            IPv4 neighbor address: 10.222.60.1
    Authentication TLV #10, length: 17
        HMAC-MD5 password: 00bb32fd7712bcea6003e516e2333077
```

The output in the exhibit was captured on an interface. Which three statements are true about the configuration on the router with hostname SaoPaulo? (Choose three.)

A. Wide metrics is not in use.
B. The router has the overload bit set to "on".
C. Authentication is enabled.
D. System ID is 1921.6805.2001.
E. Level 2 routing is enabled.

Answer: A,C,E

Question No : 35 - (Topic 1)

Click the Exhibit button.

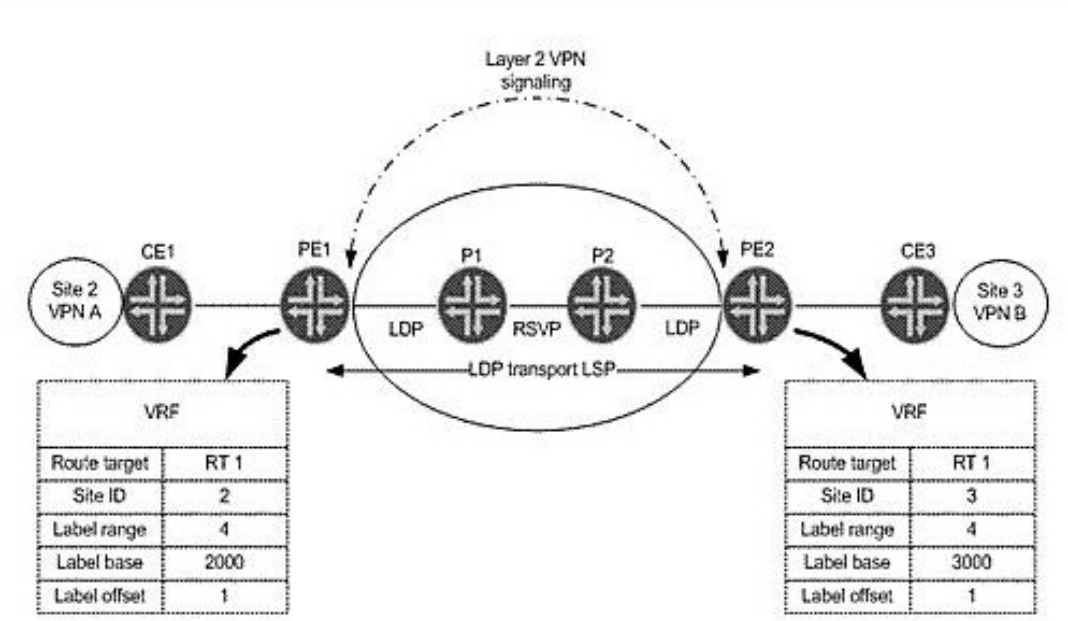

In the exhibit, on which label value does PE1 expect to receive traffic from CE3 for VPN A?

A. 2002
B. 3001
C. 3002
D. 2001

Answer: A

Question No : 36 - (Topic 1)

Which two configuration parameters are required to configure a BGP-signaled VPLS service? (Choose two.)

A. vpls-id
B. site-identifier
C. route-distinguisher
D. site-address

Answer: B,C

Question No : 37 - (Topic 1)

Click the Exhibit button.

```
user@R5# run show bgp neighbor
Peer: 192.168.56.1+179 AS 65000 Local: 192.168.56.5+56710 AS 65000
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Open Message Error
  Options: <Preference LocalAddress Refresh>
  Local Address: 192.168.56.5 Holdtime: 90 Preference: 170
  Number of flaps: 1
  Last flap event: RecvNotify
  Error: 'Open Message Error' Sent: 2 Recv: 0
  Error: 'Cease' Sent: 0 Recv: 1
  Peer ID: 192.168.56.1     Local ID: 192.168.56.5     Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet6-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 65000)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              0
    Received prefixes:            0
    Accepted prefixes:            0
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Last traffic (seconds): Received 4    Sent 4    Checked 4
  Input messages:  Total 3      Updates 1       Refreshes 0     Octets 101
  Output messages: Total 7      Updates 0       Refreshes 0     Octets 284
  Output Queue[0]: 0
```

The exhibit shows the output of a Junos show bgp neighbor command. Which two

statements are true? (Choose two.)

A. IPv4 routes will be exchanged over this session.
B. IPv6 routes will be exchanged over this session.
C. The local router initiated the BGP session.
D. BFD keepalive is configured to 30 seconds.

Answer: A,C

Question No : 38 - (Topic 1)

Click the Exhibit button.

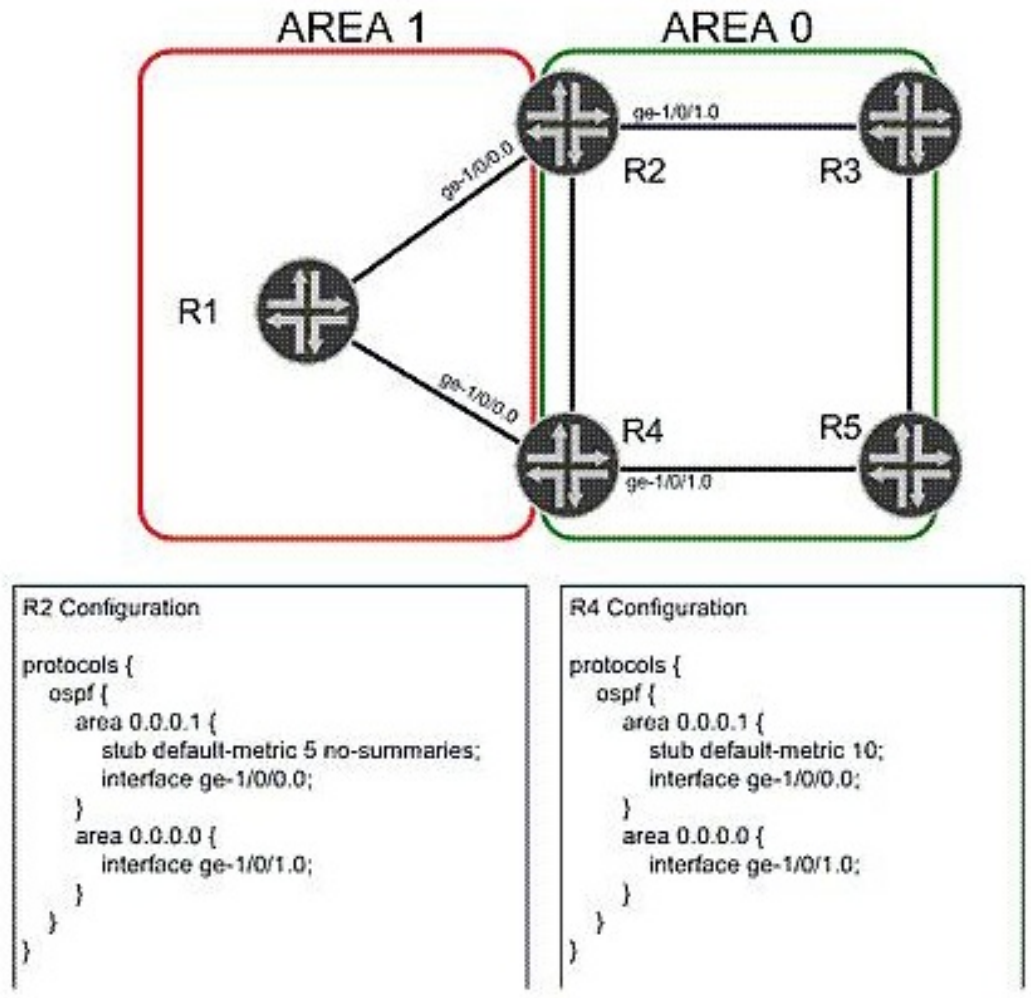

R2 and R3 advertise a default route into Area 1.Based on the configurations in the exhibit, which statement is true? (Choose two.)

A. Traffic from R1 to internal OSPF destinations in Area 0 will always transit R4.
B. Traffic from R1 to internal OSPF destinations in Area 0 will always transit R2.
C. Traffic from R1 to external OSPF destinations in Area 0 will always transit R2.
D. Traffic from R1 to external OSPF destinations in Area 0 will always transit R4.

Answer: A,C

Question No : 39 - (Topic 1)

Click the Exhibit button.

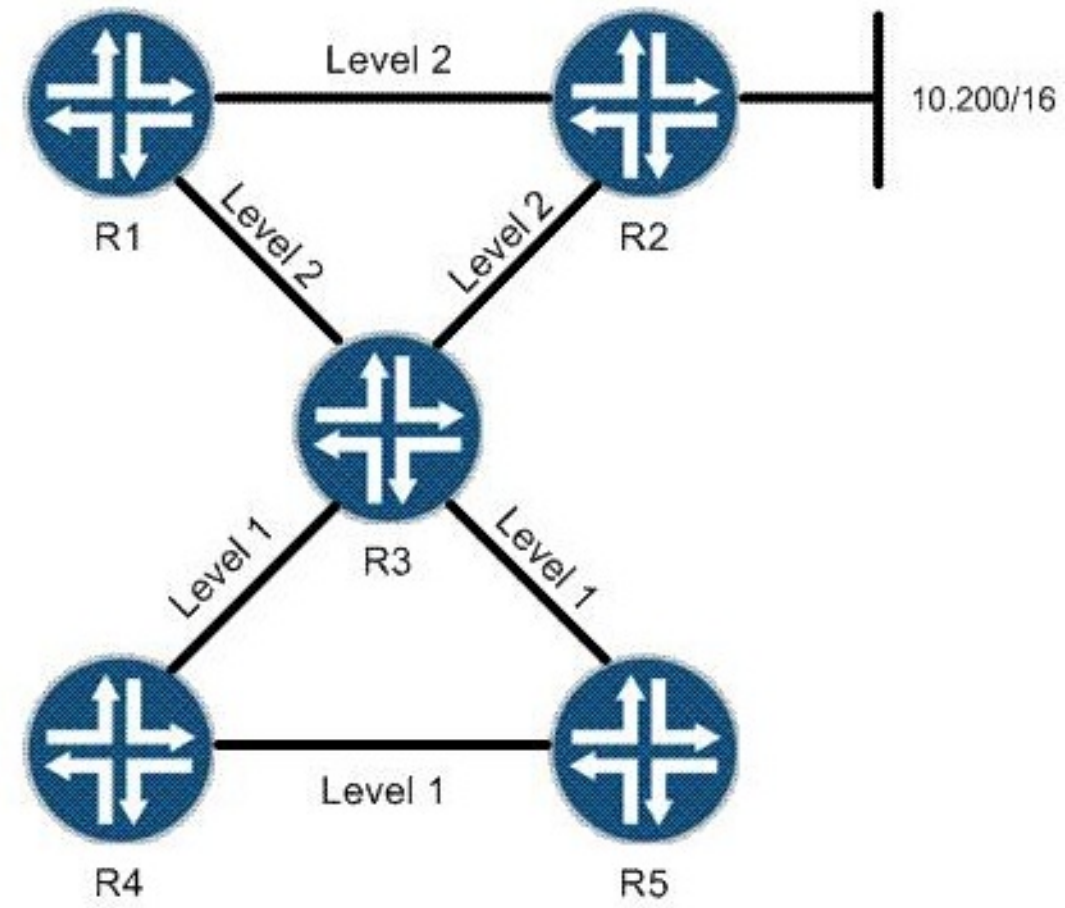

R2 is announcing the 10.200/16 network to its IS-IS neighbors. No routing policies have been applied to R3. Referring to the exhibit, will R5 have 10.200/16 as an IS-IS route?

A. Yes; IS-IS level 2 externals are passed from level 2 to level 1 by default.
B. No; IS-IS level 2 externals are only passed to level 1 if wide-metrics-only is configured on all routers.
C. Yes; all level 2 routing information is shared throughout an IS-IS domain by default.
D. No; IS-IS does not announce routes from level 2 to level 1 unless a routing policy is

applied.

Answer: D

Question No : 40 - (Topic 1)

Click the Exhibit button.

```
user@router# show
traffic-control-profiles {
    L3-unit-profile {
        scheduler-map "sched-map-example;";
        shaping-rate 30m;
        guaranteed-rate 20m;
    }
}
interfaces {
    ge-0/1/1 {
        output-traffic-control-profile "l1-port-profile;";
        unit 100 {
            output-traffic-control-profile L3-unit-profile;
        }
    }
}
```

What would happen if the guaranteed-rate command is removed from the configuration shown in the exhibit?

A. The logical interface gets a minimal bandwidth reservation.
B. The minimum-rate command should be configured instead.
C. The logical interface receives no bandwidth constraints.
D. The transmit-rate command should be configured instead.

Answer: A

Question No : 41 - (Topic 1)

```
user@R1> show interfaces terse
Interface         Admin Link  Proto Local      Remote
so-0/0/0.0        up    up    inet  10.34.0.1/30
            iso
so-1/0/0.0        up    up    inet  10.34.0.5/30
            iso
```

```
user@R1# show interfaces
so-0/0/0 {
    no-keepalives;
    encapsulation cisco-hdlc;
    unit 0 {
        family inet {
            address 10.34.0.1/30;
        }
    }
}
so-1/0/0 {
    no-keepalives;
    encapsulation cisco-hdlc;
    unit 0 {
        family inet {
            address 10.34.0.5/30;
        }
    }
}
```

```
user@R2> show interfaces terse
Interface               Admin Link  Proto Local           Remote
so-0/0/0.0              up    up    inet  10.34.0.2/30
                iso
so-1/0/0.0              up    up    inet  10.34.0.6/30
                iso
```

```
user@R2# show interfaces
so-0/0/0 {
    no-keepalives;
    encapsulation cisco-hdlc;
    unit 0 {
        family inet {
            address 10.34.0.2/30;
        }
    }
}
so-1/0/0 {
    no-keepalives;
    encapsulation cisco-hdlc;
    unit 0 {
        family inet {
            address 10.34.0.6/30;
        }
    }
}
```

You have created a new IS-IS adjacency between identical routers over two STM-4 circuits. After enabling no-keepalives on the interfaces, they come up and you are able to ping between the routers. However, the IS-IS adjacency still will not establish.

Which step will determine the cause of the problem?

A. Disable no-keepalives from the interfaces and issue a show isis adjacency command.
B. Enable family iso on all the SONET interfaces, commit, and test the circuits.
C. Disable cisco-hdlc encapsulation, enable ppp encapsulation, commit, and test the circuits.
D. Disable no-keepalives from the interfaces and issue a show interfaces extensive command for each of the SONET interfaces.

Answer: D

Question No : 42 - (Topic 1)

Click the Exhibit button.

```
[edit]
root@R4# run show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence  Checksum  Lifetime  Attributes
R4.00-00                     0x2    0xcfbc      1072  L1 L2
R3.00-00                     0x3    0xf316      1192  L1 L2  Overload
R3.02-00                     0x2    0xc17e      1192  L1 L2
  3 LSPs

IS-IS level 2 link-state database:
LSP ID                  Sequence  Checksum  Lifetime  Attributes
R4.00-00                     0x2    0x4baa      1073  L1 L2
  1 LSPs
```

Based on the output in the exhibit, which statement is correct?

A. R4 has been configured with an IS-IS export policy and is announcing external routing information.
B. R3 and R4 have an adjacency at both level 1 and level 2.
C. R3 has been configured so that it is not used for transit traffic.
D. R3 and R4 are both attached to other IS-IS areas.

Answer: C

Question No : 43 - (Topic 1)

You have applied a customized EXP rewrite rule on router R1 on the egress of the interface connecting to router R2. You want to verify if it is working properly. Which two methods would you use to accomplish this task? (Choose two.)

A. Apply a rewrite rule on the ingress of R2 for packets coming in from R1 and count those packets for the new EXP values with a firewall filter as they leave R2.
B. Use the traceroute utility.
C. Use an output filter on R1 that matches and counts on various EXP values on packets going to R2.
D. Use an input firewall filter on R2 that matches and counts on various EXP values on packets coming from R1.

Answer: C,D

Question No : 44 - (Topic 1)

Click the Exhibit button.

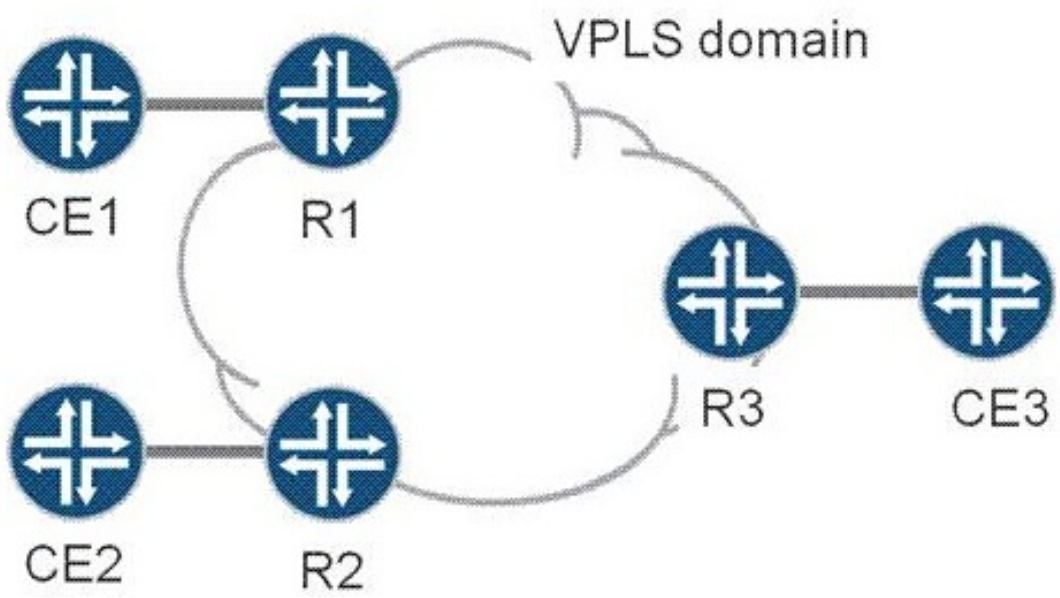

CE1, CE2, and CE3 are part of a single VPLS VPN. R1, R2, and R3 are PEs in the provider network, and have just been powered on. The VPLS domain has converged, and frames have passed between all CEs in the last minute. An Ethernet frame has just arrived at R3 from CE3. It has a source MAC address of CE3 and a destination MAC address of CE1.What does R3 do with the Ethernet frame?

A. Drops the packet as the destination MAC address is not for R3.
B. Drops the packet as the destination MAC address is not in R3's MAC table.
C. Forwards the packet to R1 only.
D. Forwards the packet to R1 and R2.

Answer: C

Question No : 45 - (Topic 1)

Click the Exhibit button.

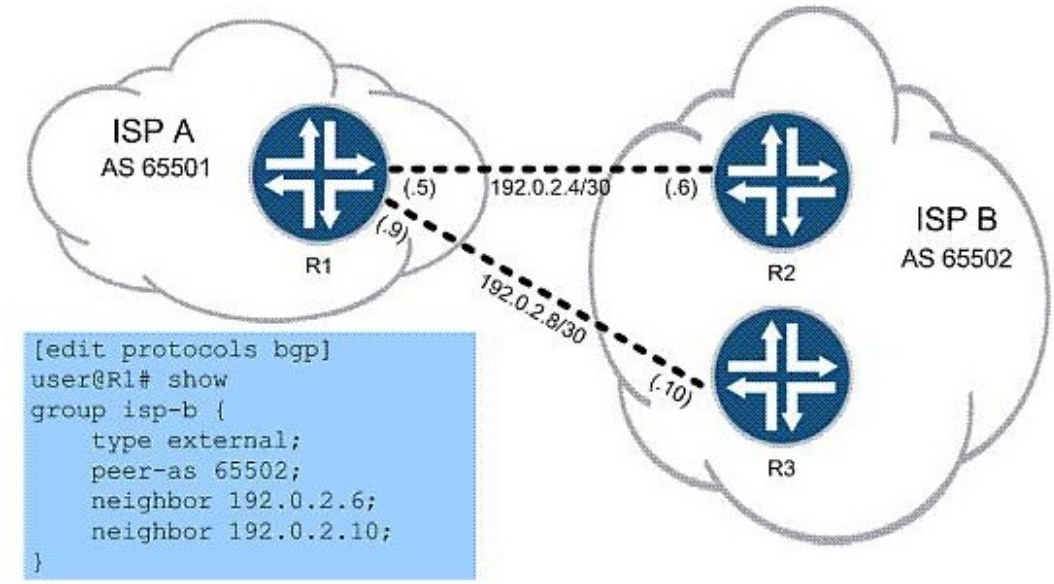

You work for ISP A, as shown in the exhibit, and must configure R1 to use load balancing across both available links to ISP B's network. Which command do you use to finish the configuration?

A. set protocols bgp group isp-b multipath
B. set routing-options forwarding-table export per-packet
C. set protocols bgp group isp-b multihop
D. set routing-options forwarding-table load-balance

Answer: A

Question No : 46 - (Topic 1)

Click the Exhibit button.

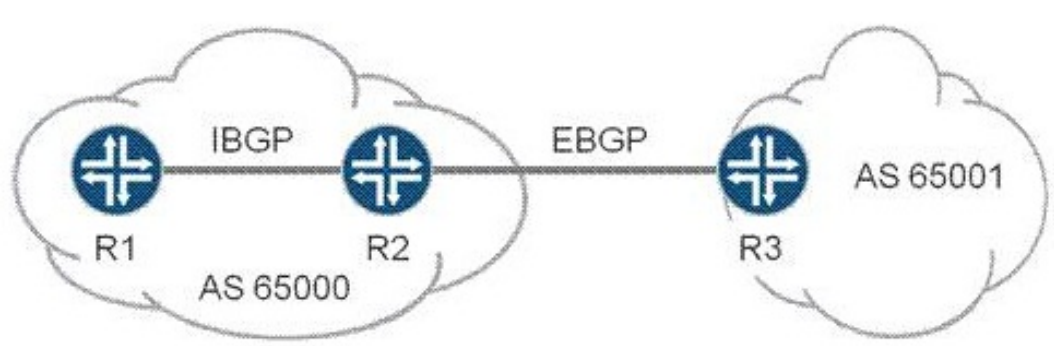

The exhibit contains a BGP topology. R1 and R2 are peering using IBGP. R2 and R3 are peering with EBGP. R1 is not installing any routes from R3 due to next-hop resolution issues. Which two configurations will resolve this issue? (Choose two.)

A. Use a policy to advertise the loopback on R2 into the IGP.
B. Advertise the R2-R3 subnet into the IGP.
C. Configure advertise-inactive on the IBGP peering session on R2.
D. Configure next-hop self on the IBGP peering session on R2.

Answer: B,D

Question No : 47 - (Topic 1)

An OSPF neighbor between routers R1 and R2 is stuck in loading state on R2. What are two causes? (Choose two.)

A. OSPF is not enabled on the interfaces.
B. A firewall filter is blocking OSPF hellos on both sides.
C. The R1 router has received a corrupted link-state request packet.
D. The interface MTU is mismatched between the routers.

Answer: C,D

Question No : 48 - (Topic 1)

What are three Junos automation scripts? (Choose three.)

A. op scripts
B. pulse scripts
C. commit scripts
D. event scripts
E. action scripts

Answer: A,C,D

Question No : 49 - (Topic 1)

Click the Exhibit button.

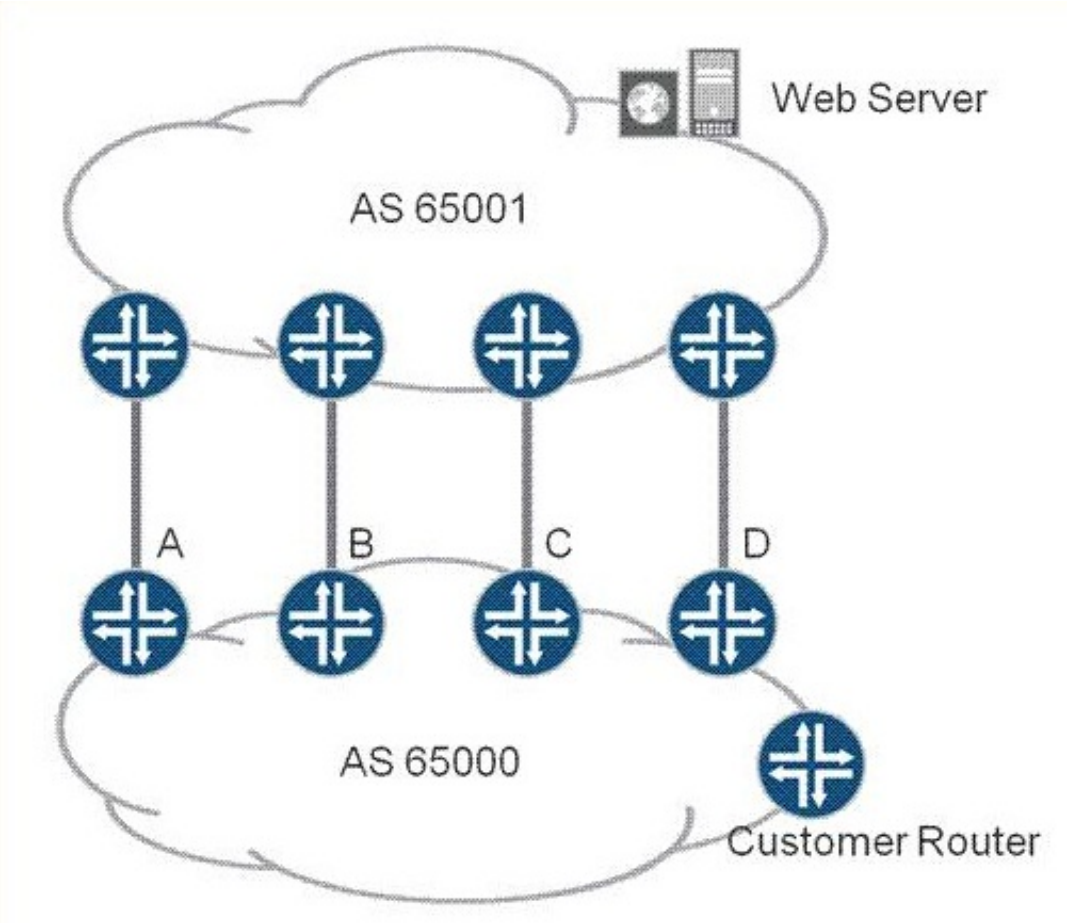

You are the administrator of AS 65000. In the exhibit, there are four links between your network (AS 65000) and your upstream provider (AS 65001).

You have an export policy on all of your routers to advertise your routes such that:

Router A: MED 100, AS Path (65000), Origin 1

Router B: MED 100, AS Path (65000 65000), Origin 0

Router C: MED 50, AS Path (65000 65000), Origin 1

Router D: MED 50, AS Path (65000), Origin 0

Through which link will traffic from the Web server enter your network (AS 65000)?

A. Router A
B. Router B
C. Router C
D. Router D

Answer: D

Question No : 50 - (Topic 1)

In an interdomain multicast deployment scenario, an RP1 is in AS1 and an RP2 is in AS2. MSDP is configured between RP1 and RP2.In which routing table on RP1 are source-active messages (SAs) received from RP2 by default?

A. inet.0
B. inet.2
C. inet.1
D. inet.4

Answer: D

Question No : 51 - (Topic 1)

Click the Exhibit button.

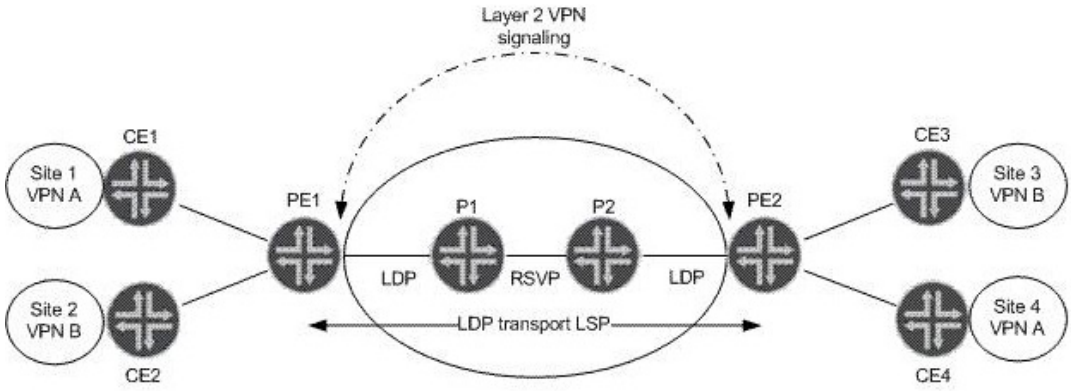

A LDP Layer 2 circuit is shown for VPN A and VPN B. LDP tunneling over RSVP is activated on P1 and P2. Referring to the exhibit, which statement is true about the LDP Layer 2 circuit?

A. MAC learning is needed and using the inner VPN label between PE1 and PE2 for VPN A or VPN B.
B. Targeted LDP sessions are established between PE1, P1 and P2, PE2.
C. Label stitching must be configured on P1 and P2 for end to end transport LSPs.

D. LDP must be enabled on the loopback interfaces of PE1 and PE2.

Answer: D

Question No : 52 - (Topic 1)

Click the Exhibit button.

```
user@PE1> show bgp neighbor | match nlri
  NLRI for restart configured on peer: inet-unicast inet-vpn-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast

user@PE2> show bgp neighbor | match nlri
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet-vpn-unicast
  NLRI for this session: inet-unicast
  NLRI that peer supports restart for: inet-unicast inet-vpn-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
```

Two PE routers in your Layer 3 VPN are not advertising customer VPN routes to each other. Referring to the output in the exhibit, which configuration parameter is missing?

A. family inet on PE1
B. family inet on PE2
C. family inet-vpn on PE1
D. family inet-vpn on PE2

Answer: D

Question No : 53 - (Topic 1)

Click the Exhibit button.

```
user@PE2> show route advertising-protocol bgp 192.168.3.1

customer-vpn.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  Prefix  Nexthop       MED     Lclpref    AS path
* 172.16.2.0/24          Self                          100        I
* 172.16.20.0/30         Self                          100        65001 I
* 172.16.20.4/30         Self                          100        65001 I
* 172.16.20.8/30         Self                          100        65001 I


user@PE1> show route advertising-protocol bgp 172.16.1.2


user@PE1> show route receive-protocol bgp 192.168.4.1

inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

customer-vpn.inet.0: 6 destinations, 6 routes (2 active, 0 holddown, 4 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 4 destinations, 4 routes (0 active, 0 holddown, 4 hidden)
```

Customer A is complaining that routes advertised from the CE2 router are not being received on the CE1 router. The physical topology of the network is CE1-PE1-PE2-CE2. The CE1-PE1 subnet is 172.16.1.0/24. The CE2-PE2 subnet is 172.16.2.0/24. PE1's loopback is 192.168.3.1 and PE2's loopback is 192.168.4.1.Referring to the output in the exhibit, what is the problem?

A. No LSP exists between PE1 and PE2.
B. Route targets are not properly configured.
C. as-override is not configured in the VRFs.
D. family inet-vpn is not configured on the PEs.

Answer: A

Question No : 54 - (Topic 1)

Your OSPF network includes an NSSA. Which LSA type is injected into the NSSA by the ASBR?

A. Type 3
B. Type 5
C. Type 7
D. Type 9

Answer: C

Question No : 55 - (Topic 1)

You are asked to retain several routes from an external BGP neighbor in the routing table on your local router, but you are not allowed to forward traffic to these destinations. You have configured a forwarding table firewall filter to block these routes, and applied it under the (edit forwarding-options] hierarchy, but the routes are still showing up in the forwarding table. What is required to achieve this task?

A. Configure an EBGP import policy on your local router to block the routes.
B. Have the EBGP neighbor configure an export policy to block the routes.
C. Configure an export policy for the forwarding table to block the routes.
D. Use the no-install configuration statement within the EBGP neighbor group on your local router.

Answer: C

Question No : 56 - (Topic 1)

Click the Exhibit button.

```
192.168.56.1
  From: 192.168.56.5, LSPstate: Up, ActiveRoute: 0
  LSPname: to-r6, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Tue Feb 22 21:38:36 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 18916 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.10.56.1 (ge-1/0/1.0) 7 pkts
  RESV rcvfrom: 10.10.56.1 (ge-1/0/1.0) 5 pkts
  Explct route: 10.10.56.1
  Record route: <self> 10.10.56.1
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 10.10.10.9 (ge-1/0/2.0) 4 pkts
    Detour RESV rcvfrom: 10.10.10.9 (ge-1/0/2.0) 3 pkts
    Detour Explct route: 10.10.10.9 10.10.10.6
    Detour Record route: <self> 10.10.10.9 10.10.10.6
    Detour Label out: 299856
```

Referring to the exhibit, which type of traffic protection mechanism is used for the LSP?

A. link-protection
B. fast-reroute
C. node-link-protection
D. bypass

Answer: B

Question No : 57 - (Topic 1)

Router R5 has the overload parameter configured. Which statement is true?

A. R5 will purge its LSAs from the network until the overload condition is cleared.
B. R5 will increase its link metrics to 65535 and will stop forwarding transit traffic to OSPF destinations.
C. R5 will increase its link metrics to 65535 and will continue to forward transit traffic to OSPF destinations.
D. R5 will send an overload LSA to its neighbors to indicate it is in the overload state.

Answer: C

Question No : 58 - (Topic 1)

Click the Exhibit button.

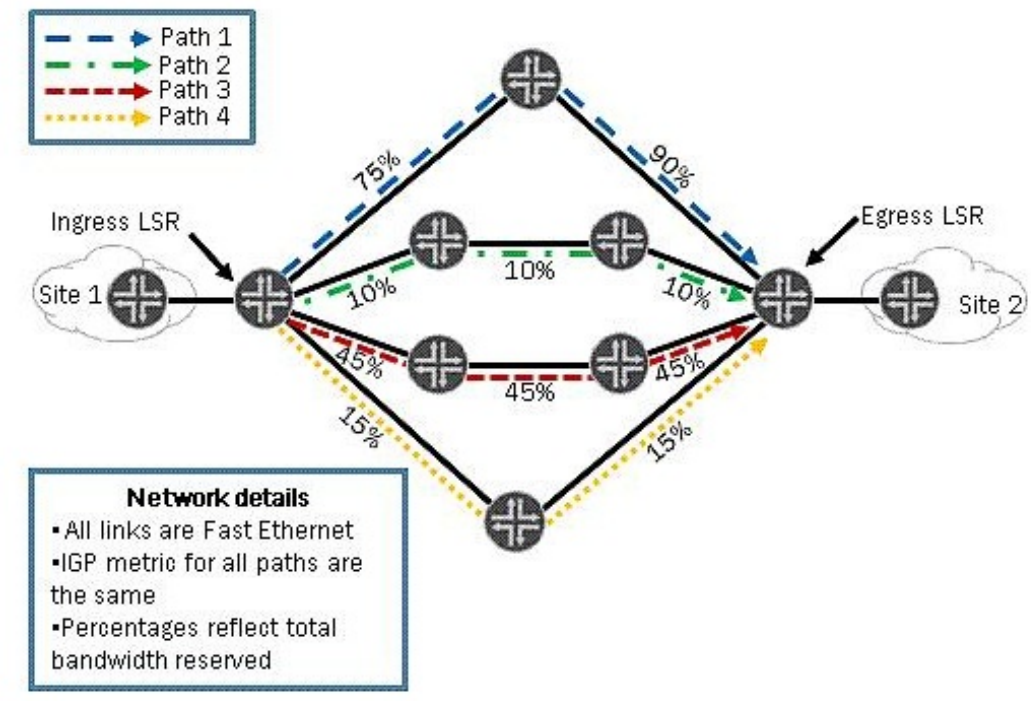

You have an MPLS network and you have configured most-fill as a CSPF tiebreaker. Using the information in the exhibit, which path will be used to signal a new LSP requiring 12 Mbps?

A. Path 1
B. Path 2
C. Path 3
D. Path 4

Answer: D

Question No : 59 - (Topic 1)

Click the Exhibit button.

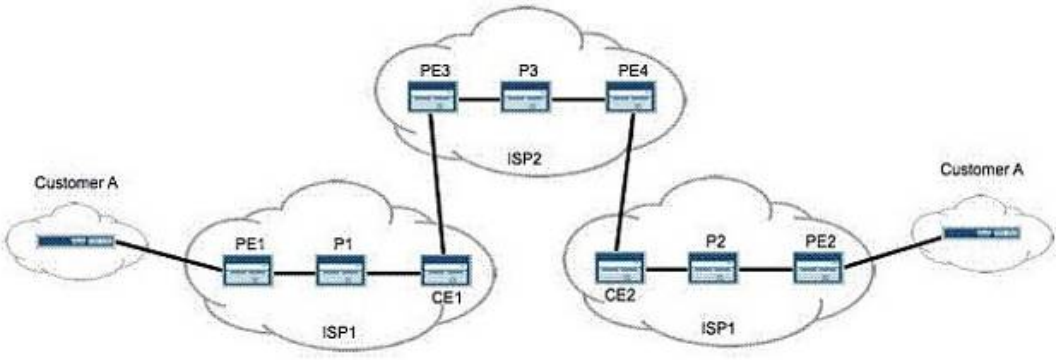

Referring to the exhibit, PE2 requires the loopback of PE1 to appear in the inet.3 routing table as a labeled route. Which configuration parameter is specifically required to support this?

A. resolve-vpn
B. family inet-vpn
C. traffic-engineering bgp-igp
D. traffic-engineering mpls-forwarding

Answer: A

Question No : 60 - (Topic 1)

An OSPF network has been designed with multiple areas to improve scalability. Which two statements are true? (Choose two.)

A. Each router in the OSPF network runs the shortest-path-first algorithm to determine paths through the network.
B. The Area Border Router for each area runs the shortest-path-first algorithm and floods its results through the area.
C. Each area must have at least one link connecting it to each of the other areas of the OSPF network.
D. OSPF provides loop-free routing within an OSPF routing domain, but does not guarantee symmetrical routing.

Answer: A,D

Question No : 61 - (Topic 1)

```
user@router> show bgp neighbor
Peer: 2.2.2.2+59344 AS 89       Local: 1.1.1.1+179 AS 89
  Type: Internal State: Established Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Refresh>
  Local Address: 1.1.1.1 Holdtime: 90 Preference: 170
  Number of flaps: 1
  Last flap event: RecvNotify
  Error: 'Hold Timer Expired Error' Sent: 0 Recv: 1
  Peer ID: 2.2.2.2    Local ID: 1.1.1.1     Active Holdtime: 90
  Keepalive Interval: 30   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 89)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:       0
    Received prefixes:     0
    Accepted prefixes:     0
    Suppressed due to damping:  0
    Advertised prefixes: 0
  Last traffic (seconds): Received 23 Sent 12 Checked 66
  Input messages: Total 13292 Updates 1R efreshes 0O ctets 252592
  Output messages: Total 13301U pdates 0R efreshes 0O ctets 252782
  Output Queue[0]: 0
```

Which two scenarios would explain the error shown in the exhibit? (Choose two.)

A. Packets were dropped in the receive path of the local router.
B. The local router received a route withdrawal from the peer.
C. Packets were dropped in the receive path of the peer router.
D. The local router egress interface towards the peer was congested.

Answer: C,D

Question No : 62 - (Topic 1)

Click the Exhibit button.

Microsoft Exams List

70-246 Dump PDF VCE	70-485 Dump PDF VCE	70-742 Dump PDF VCE	98-366 Dump PDF VCE
70-247 Dump PDF VCE	70-486 Dump PDF VCE	70-743 Dump PDF VCE	98-367 Dump PDF VCE
70-331 Dump PDF VCE	70-487 Dump PDF VCE	70-744 Dump PDF VCE	98-368 Dump PDF VCE
70-332 Dump PDF VCE	70-488 Dump PDF VCE	70-761 Dump PDF VCE	98-369 Dump PDF VCE
70-333 Dump PDF VCE	70-489 Dump PDF VCE	70-762 Dump PDF VCE	98-372 Dump PDF VCE
70-334 Dump PDF VCE	70-490 Dump PDF VCE	70-765 Dump PDF VCE	98-373 Dump PDF VCE
70-339 Dump PDF VCE	70-491 Dump PDF VCE	70-768 Dump PDF VCE	98-374 Dump PDF VCE
70-341 Dump PDF VCE	70-492 Dump PDF VCE	70-980 Dump PDF VCE	98-375 Dump PDF VCE
70-342 Dump PDF VCE	70-494 Dump PDF VCE	70-981 Dump PDF VCE	98-379 Dump PDF VCE
70-345 Dump PDF VCE	70-496 Dump PDF VCE	70-982 Dump PDF VCE	MB2-700 Dump PDF VCE
70-346 Dump PDF VCE	70-497 Dump PDF VCE	74-343 Dump PDF VCE	MB2-701 Dump PDF VCE
70-347 Dump PDF VCE	70-498 Dump PDF VCE	74-344 Dump PDF VCE	MB2-702 Dump PDF VCE
70-348 Dump PDF VCE	70-499 Dump PDF VCE	74-409 Dump PDF VCE	MB2-703 Dump PDF VCE
70-354 Dump PDF VCE	70-517 Dump PDF VCE	74-678 Dump PDF VCE	MB2-704 Dump PDF VCE
70-383 Dump PDF VCE	70-532 Dump PDF VCE	74-697 Dump PDF VCE	MB2-707 Dump PDF VCE
70-384 Dump PDF VCE	70-533 Dump PDF VCE	77-420 Dump PDF VCE	MB2-710 Dump PDF VCE
70-385 Dump PDF VCE	70-534 Dump PDF VCE	77-427 Dump PDF VCE	MB2-711 Dump PDF VCE
70-410 Dump PDF VCE	70-640 Dump PDF VCE	77-600 Dump PDF VCE	MB2-712 Dump PDF VCE
70-411 Dump PDF VCE	70-642 Dump PDF VCE	77-601 Dump PDF VCE	MB2-713 Dump PDF VCE
70-412 Dump PDF VCE	70-646 Dump PDF VCE	77-602 Dump PDF VCE	MB2-714 Dump PDF VCE
70-413 Dump PDF VCE	70-673 Dump PDF VCE	77-603 Dump PDF VCE	MB2-715 Dump PDF VCE
70-414 Dump PDF VCE	70-680 Dump PDF VCE	77-604 Dump PDF VCE	MB2-716 Dump PDF VCE
70-417 Dump PDF VCE	70-681 Dump PDF VCE	77-605 Dump PDF VCE	MB2-717 Dump PDF VCE
70-461 Dump PDF VCE	70-682 Dump PDF VCE	77-881 Dump PDF VCE	MB2-718 Dump PDF VCE
70-462 Dump PDF VCE	70-684 Dump PDF VCE	77-882 Dump PDF VCE	MB5-705 Dump PDF VCE
70-463 Dump PDF VCE	70-685 Dump PDF VCE	77-883 Dump PDF VCE	MB6-700 Dump PDF VCE
70-464 Dump PDF VCE	70-686 Dump PDF VCE	77-884 Dump PDF VCE	MB6-701 Dump PDF VCE
70-465 Dump PDF VCE	70-687 Dump PDF VCE	77-885 Dump PDF VCE	MB6-702 Dump PDF VCE
70-466 Dump PDF VCE	70-688 Dump PDF VCE	77-886 Dump PDF VCE	MB6-703 Dump PDF VCE
70-467 Dump PDF VCE	70-689 Dump PDF VCE	77-887 Dump PDF VCE	MB6-704 Dump PDF VCE
70-469 Dump PDF VCE	70-692 Dump PDF VCE	77-888 Dump PDF VCE	MB6-705 Dump PDF VCE
70-470 Dump PDF VCE	70-695 Dump PDF VCE	77-891 Dump PDF VCE	MB6-884 Dump PDF VCE
70-473 Dump PDF VCE	70-696 Dump PDF VCE	98-349 Dump PDF VCE	MB6-885 Dump PDF VCE
70-480 Dump PDF VCE	70-697 Dump PDF VCE	98-361 Dump PDF VCE	MB6-886 Dump PDF VCE
70-481 Dump PDF VCE	70-698 Dump PDF VCE	98-362 Dump PDF VCE	MB6-889 Dump PDF VCE
70-482 Dump PDF VCE	70-734 Dump PDF VCE	98-363 Dump PDF VCE	MB6-890 Dump PDF VCE
70-483 Dump PDF VCE	70-740 Dump PDF VCE	98-364 Dump PDF VCE	MB6-892 Dump PDF VCE
70-484 Dump PDF VCE	70-741 Dump PDF VCE	98-365 Dump PDF VCE	MB6-893 Dump PDF VCE

Cisco Exams List

010-151 Dump PDF VCE	350-018 Dump PDF VCE	642-737 Dump PDF VCE	650-667 Dump PDF VCE
100-105 Dump PDF VCE	352-001 Dump PDF VCE	642-742 Dump PDF VCE	650-669 Dump PDF VCE
200-001 Dump PDF VCE	400-051 Dump PDF VCE	642-883 Dump PDF VCE	650-752 Dump PDF VCE
200-105 Dump PDF VCE	400-101 Dump PDF VCE	642-885 Dump PDF VCE	650-756 Dump PDF VCE
200-120 Dump PDF VCE	400-151 Dump PDF VCE	642-887 Dump PDF VCE	650-968 Dump PDF VCE
200-125 Dump PDF VCE	400-201 Dump PDF VCE	642-889 Dump PDF VCE	700-001 Dump PDF VCE
200-150 Dump PDF VCE	400-251 Dump PDF VCE	642-980 Dump PDF VCE	700-037 Dump PDF VCE
200-155 Dump PDF VCE	400-351 Dump PDF VCE	642-996 Dump PDF VCE	700-038 Dump PDF VCE
200-310 Dump PDF VCE	500-006 Dump PDF VCE	642-997 Dump PDF VCE	700-039 Dump PDF VCE
200-355 Dump PDF VCE	500-007 Dump PDF VCE	642-998 Dump PDF VCE	700-101 Dump PDF VCE
200-401 Dump PDF VCE	500-051 Dump PDF VCE	642-999 Dump PDF VCE	700-104 Dump PDF VCE
200-601 Dump PDF VCE	500-052 Dump PDF VCE	644-066 Dump PDF VCE	700-201 Dump PDF VCE
210-060 Dump PDF VCE	500-170 Dump PDF VCE	644-068 Dump PDF VCE	700-205 Dump PDF VCE
210-065 Dump PDF VCE	500-201 Dump PDF VCE	644-906 Dump PDF VCE	700-260 Dump PDF VCE
210-250 Dump PDF VCE	500-202 Dump PDF VCE	646-048 Dump PDF VCE	700-270 Dump PDF VCE
210-255 Dump PDF VCE	500-254 Dump PDF VCE	646-365 Dump PDF VCE	700-280 Dump PDF VCE
210-260 Dump PDF VCE	500-258 Dump PDF VCE	646-580 Dump PDF VCE	700-281 Dump PDF VCE
210-451 Dump PDF VCE	500-260 Dump PDF VCE	646-671 Dump PDF VCE	700-295 Dump PDF VCE
210-455 Dump PDF VCE	500-265 Dump PDF VCE	646-985 Dump PDF VCE	700-501 Dump PDF VCE
300-070 Dump PDF VCE	500-275 Dump PDF VCE	648-232 Dump PDF VCE	700-505 Dump PDF VCE
300-075 Dump PDF VCE	500-280 Dump PDF VCE	648-238 Dump PDF VCE	700-601 Dump PDF VCE
300-080 Dump PDF VCE	500-285 Dump PDF VCE	648-244 Dump PDF VCE	700-602 Dump PDF VCE
300-085 Dump PDF VCE	500-290 Dump PDF VCE	648-247 Dump PDF VCE	700-603 Dump PDF VCE
300-101 Dump PDF VCE	500-801 Dump PDF VCE	648-375 Dump PDF VCE	700-701 Dump PDF VCE
300-115 Dump PDF VCE	600-199 Dump PDF VCE	648-385 Dump PDF VCE	700-702 Dump PDF VCE
300-135 Dump PDF VCE	600-210 Dump PDF VCE	650-032 Dump PDF VCE	700-703 Dump PDF VCE
300-160 Dump PDF VCE	600-211 Dump PDF VCE	650-042 Dump PDF VCE	700-801 Dump PDF VCE
300-165 Dump PDF VCE	600-212 Dump PDF VCE	650-059 Dump PDF VCE	700-802 Dump PDF VCE
300-180 Dump PDF VCE	600-455 Dump PDF VCE	650-082 Dump PDF VCE	700-803 Dump PDF VCE
300-206 Dump PDF VCE	600-460 Dump PDF VCE	650-127 Dump PDF VCE	810-403 Dump PDF VCE
300-207 Dump PDF VCE	600-501 Dump PDF VCE	650-128 Dump PDF VCE	820-424 Dump PDF VCE
300-208 Dump PDF VCE	600-502 Dump PDF VCE	650-148 Dump PDF VCE	840-425 Dump PDF VCE
300-209 Dump PDF VCE	600-503 Dump PDF VCE	650-159 Dump PDF VCE	
300-210 Dump PDF VCE	600-504 Dump PDF VCE	650-281 Dump PDF VCE	
300-320 Dump PDF VCE	640-692 Dump PDF VCE	650-393 Dump PDF VCE	
300-360 Dump PDF VCE	640-875 Dump PDF VCE	650-472 Dump PDF VCE	
300-365 Dump PDF VCE	640-878 Dump PDF VCE	650-474 Dump PDF VCE	
300-370 Dump PDF VCE	640-911 Dump PDF VCE	650-575 Dump PDF VCE	
300-375 Dump PDF VCE	640-916 Dump PDF VCE	650-621 Dump PDF VCE	
300-465 Dump PDF VCE	642-035 Dump PDF VCE	650-663 Dump PDF VCE	
300-470 Dump PDF VCE	642-732 Dump PDF VCE	650-665 Dump PDF VCE	
300-475 Dump PDF VCE	642-747 Dump PDF VCE	650-754 Dump PDF VCE	

HOT EXAMS

Cisco

100-105 Dumps VCE PDF
200-105 Dumps VCE PDF
300-101 Dumps VCE PDF
300-115 Dumps VCE PDF
300-135 Dumps VCE PDF
300-320 Dumps VCE PDF
400-101 Dumps VCE PDF
640-911 Dumps VCE PDF
640-916 Dumps VCE PDF

Microsoft

70-410 Dumps VCE PDF
70-411 Dumps VCE PDF
70-412 Dumps VCE PDF
70-413 Dumps VCE PDF
70-414 Dumps VCE PDF
70-417 Dumps VCE PDF
70-461 Dumps VCE PDF
70-462 Dumps VCE PDF
70-463 Dumps VCE PDF
70-464 Dumps VCE PDF
70-465 Dumps VCE PDF
70-480 Dumps VCE PDF
70-483 Dumps VCE PDF
70-486 Dumps VCE PDF
70-487 Dumps VCE PDF

CompTIA

220-901 Dumps VCE PDF
220-902 Dumps VCE PDF
N10-006 Dumps VCE PDF
SY0-401 Dumps VCE PDF