Vendor: Juniper

Exam Code: JN0-663

Exam Name: Service Provider Routing and Switching, Professional (JNCIP-SP)

Version: DEMO

QUESTION 1
A customer wants to reduce LSP flooding in their IS-IS network.

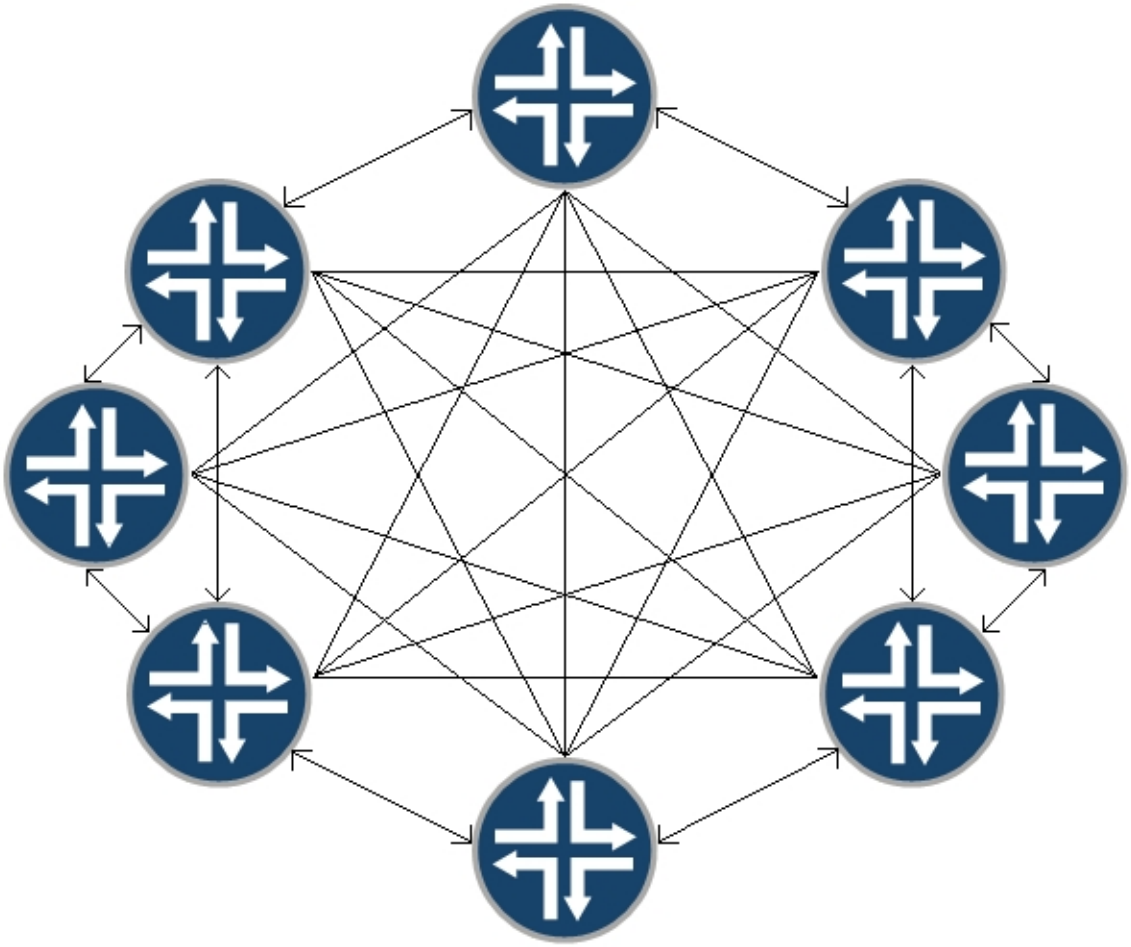

Which parameter should you change to accomplish this task?

A.
```
[edit protocols isis interface <interface-name>]
user@router# set lsp-interval 1000
```
B.
```
[edit protocols isis interface <interface-name>]
user@router# set csnp-interval 65535
```
C.
```
[edit protocols isis interface <interface-name>]
user@router# set mesh-group <mesh-group-number>
```
D.
```
[edit protocols isis]
user@router# set spf-options rapid-runs 5
```

Answer: C
Explanation:
When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.

QUESTION 2
Which two statements about wide and narrow metrics used in IS-IS are correct? (Choose two)

A. Wide metrics are sent by default and use 24 bits in TLVs to send information
B. Narrow metrics are enabled by default and use 8 bits in TLVs to send information
C. Disabling narrow metrics results in external routes being leaked from L1 to L2 areas automatically
D. Wide metrics are enabled with the wide-metrics-only parameter under protocols isis hierarchy.

Answer: AC
Explanation:
With these TLVs, IS-IS metrics can have values up to 16,777,215 (224 – 1).
By default, Junos OS supports the sending and receiving of wide metrics.
https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/isis-wide-metrics.html

QUESTION 3
Referring to the exhibit, which OSPFv3 configuration is implemented on router R1?

```
user@R1> show ospf3 interface
Interface     State       Area         DR ID        BDR ID       Nbrs
ge-0/0/0.0    DR          0.0.0.0      172.16.1.2   172.16.1.1      1
ge-0/0/0.0    PtToPt      0.0.0.1      0.0.0.0      0.0.0.0         1
ge-0/0/1.0    BDR         0.0.0.1      172.16.1.1   172.16.1.2      1

user@R1> show ospf3 neighbor
ID           Interface    State   Pri     Dead
172.16.1.1   ge-0/0/0.0   Full    128       39
   Neighbor-address fe80::20c:29ff:fef9:7f7b
   Area 0.0.0.0
172.16.1.1   ge-0/0/0.0   Full    128       37
   Neighbor-address fe80::20c:29ff:fef9:7f7b
   Area 0.0.0.1
172.16.1.1   ge-0/0/1.0   Full    128       37
   Neighbor-address fe80::20c:29ff:fef9:7f85
   Area 0.0.0.1
```

A.
```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0
set protocols ospf3 area 0.0.0.1 virtual-link neighbor-id 172.16.1.2
```
B.
```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 interface-type p2p
```
C.
```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 secondary
```
D.
```
set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0
set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0
```

Answer: C
Explanation:
By default, a single interface can belong to only one OSPF area. You can configure a single interface to belong in multiple OSPF areas. Doing so allows the corresponding link to be

considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. When configuring a secondary interface, consider the following:
For OSPFv2, you cannot configure point-to-multipoint and nonbroadcast multiaccess (NBMA) network interfaces as a secondary interface because secondary interfaces are treated as a point-to-point unnumbered link.
Secondary interfaces are supported for LAN interfaces (the primary interface can be a LAN interface, but any secondary interfaces are treated as point-to-point unnumbered links over the LAN). In this scenario, you must ensure that there are only two routing devices on the LAN or that there are only two routing devices on the LAN that have secondary interfaces configured for a specific OSPF area.
https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/topic-map/configuring-ospf-areas.html#id-example-configuring-multiarea-adjacency-for-ospf

QUESTION 4
Which two statements regarding Ethernet segments (ES) are correct? (Choose two)

A. The Type-1 EVPN route will indicate if the ES is all-active or single-active.
B. The Type-4 EVPN route will be used to elect the designated forwarder for the ES.
C. The Type-2 EVPN route will indicate if there is a designated forwarder on the ES.
D. The Type-3 EVPN route will be used for the aliasing function to load-balance to the ES

Answer: AB
Explanation:
A designated forwarder (DF) is required when customer edge devices (CEs) are multihomed to more than one provider edge (PE) device. Without a designated forwarder, multihomed hosts would receive duplicate packets. Designated forwarders are chosen for an Ethernet segment identifier (ESI) based on type 4 route advertisements.
https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/ref/statement/designated-forwarder-election-hold-time-evpn.html

QUESTION 5
The link between CE1 and PE1 has a history of flapping. To avoid the impact that flapping causes to the network you decide to use route damping
Which statement is correct in this scenario?

A. Dampened routes decay at a sliding rate known as half-life
B. Routes become dampened when the configured max-suppress value is reached
C. Dampening is enabled on interfaces
D. Dampened routes become active when their figure of merit drops below the reuse value.

Answer: D
Explanation:
Reuse merit
Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/command/show-policy-damping.html

QUESTION 6
Referring to the exhibit, the local BGP router is receiving IPv4 routes from the BGP neighbor, but

it is not receiving L3 VPN routes from the BGP neighbor.

```
user@router> show bgp neighbor 192.168.100.2
Peer: 192.168.100.2+50862 AS 65512 Local: 192.168.100.1+179 AS 65512
  Group: INT                    Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Refresh>
  Options: <GracefulShutdownRcv>
  Local Address: 192.168.100.1 Holdtime: 90 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Number of flaps: 0
  Peer ID: 192.168.100.2    Local ID: 192.168.100.1    Active Holdtime: 90
  Keepalive Interval: 30      Group index: 0     Peer index: 0      SNMP index: 3
  I/O Session Thread: bgpio-0     State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet-vpn-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
  Restart flag received from the peer: Notification
  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 does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65512)
  Peer does not support Addpath
  NLRI(s) enabled for color nexthop resolution: inet-unicast
...
```

Which two actions should you take to solve this problem? (Choose two.)

A. Configure the family inet-vpn unicast statement on the local BGP router.
B. Configure the family inet unicast statement on the local BGP router.
C. Configure the family inet unicast statement on the BGP neighbor
D. Configure the family inet-vpn unicast statement on the BGP neighbor.

Answer: AB
Explanation:
It is clearly shown in the exhibit that inet-unicast is the only NLRI enabled locally, so inet-vpn-unicast needs to be configured too. However, when you do this, it overrides the default NLRI for a peer which is inet-unicast, resulting in only inet-vpn-unicast being advertised to the peer. To continue using inet-unicast this must now also be configured explicitly.

QUESTION 7
A Layer 3 VPN exists in the provider network and the CE devices are connecting to the PE devices using BGP. The PE devices are receiving BGP routes from the CE devices and the PE devices have the CE BGP routes in their respective routing tables. However, the remote CE devices are not receiving the BGP routes.
Referring to the exhibit, what is the problem?

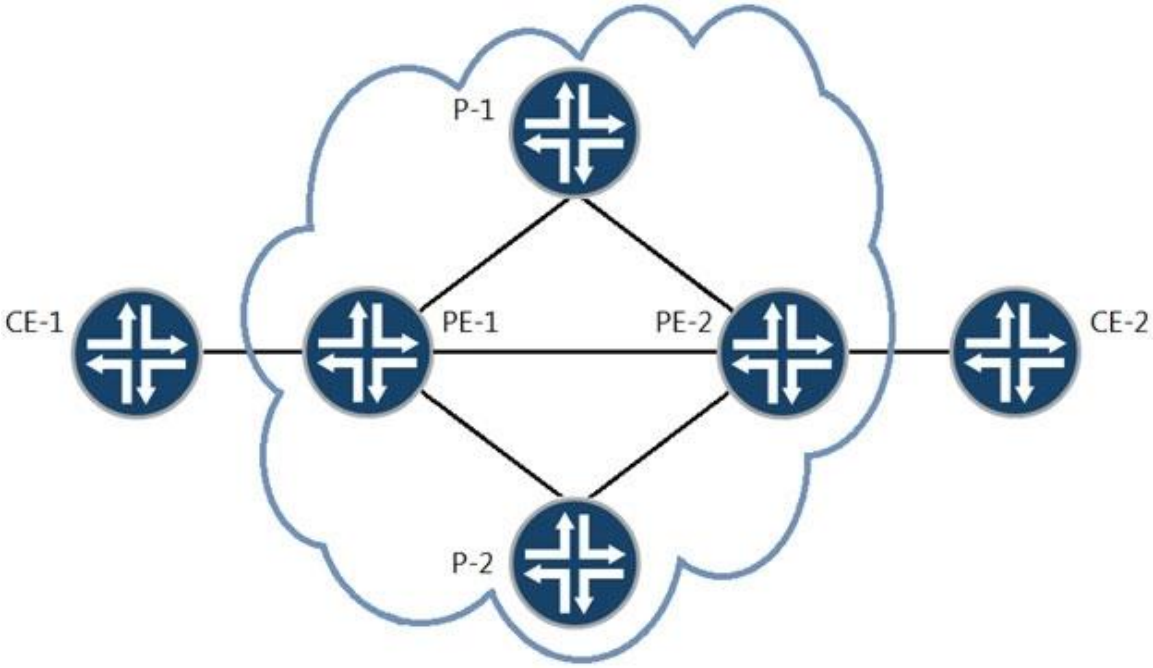

A. The CE devices are detecting an AS loop
B. A VRF target community mismatch exists.
C. A route distinguisher mismatch exists.
D. The PE devices are detecting an AS loop.

Answer: D
Explanation:
By default, Junos OS does not advertise the routes learned from one EBGP peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGP peers that are in the same AS as the originating peer, regardless of the routing instance.
https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/statement/advertise-peer-as-edit-protocols-bgp.html

QUESTION 8
What occurs when a router running IS-IS receives an LSP with the overload bit set?

A. The LSP is not added to the link-state database.
B. The LSP's metric will be set to 65535.
C. The LSP is ignored during SPF calculation.
D. The LSP's metric will be set to 16777215.

Answer: B
Explanation:
In overload mode, the routing device advertisement is originated with all the transit routing device links (except stub) set to a metric of 0xFFFF = 65535

QUESTION 9
You have a Layer 3 VPN established between PE-1 and PE-2 as well as between PE-1 and PE-3. You are using a route reflector (RR-1) to distribute VPN routes to your IBGP peers. You are

asked to ensure that only relevant routes are sent from RR-1 to each of the PE routers. Referring to the exhibit, which statement is correct?

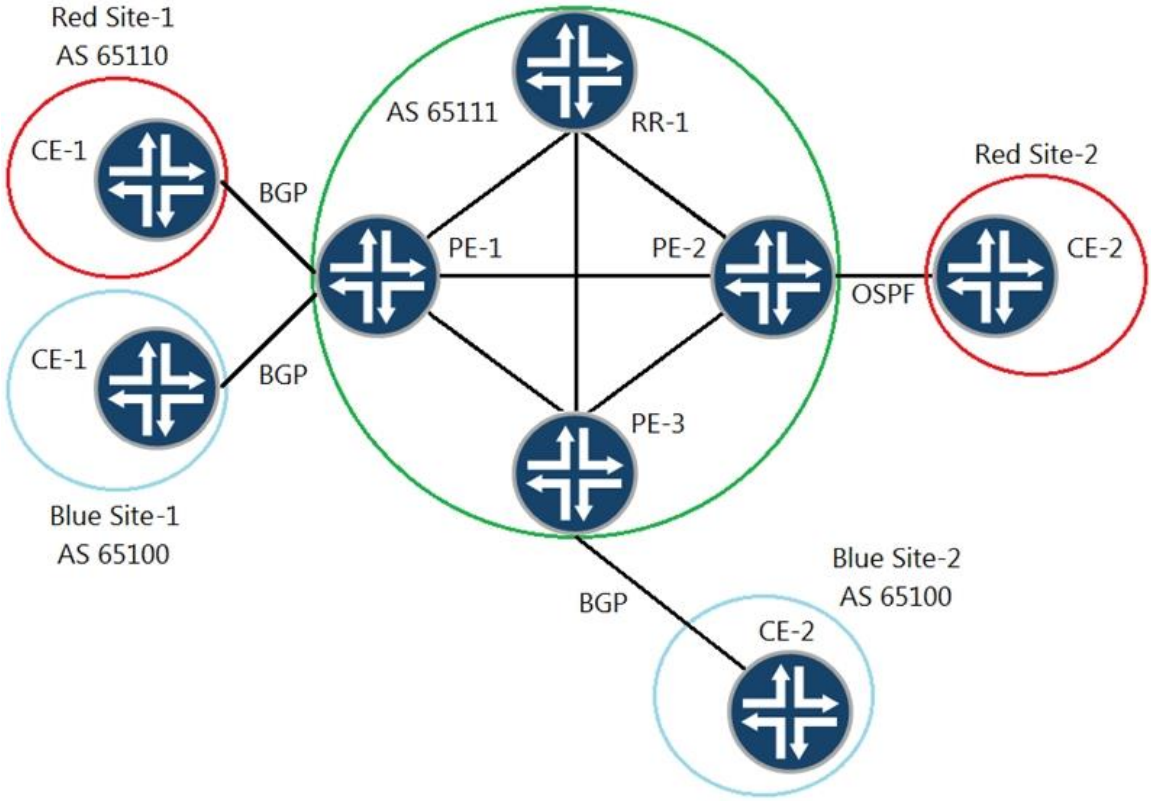

A. You should use VRF export policies on RR-1 to control which routes are sent to each PE router.
B. You should use route target filtering on RR-1 and all the PE devices to control which routes are sent to each PE router.
C. You should use firewall filtering on RR-1 and all the PE devices to control which routes are sent to each PE router.
D. You should use route target filtering only on RR-1 to control which routes are sent to each PE router.

Answer: B
Explanation:
Route Target needs to be configured on PE and RR. PE will advertise to RR required target: communities and RR will advertise 0/0 to request all VPN routes from PE.

QUESTION 10
You must deploy an interprovider VPN option that ensures that the ASBRs do not need to store any VPN routes.
In this scenario, which interprovider VPN option should you choose?

A. option A
B. option B
C. option C
D. option D

Answer: C
Explanation:
Labeled IPv4 (not VPN-IPv4) routes are exchanged by the ASBRs to support MPLS. An MP-EBGP session between the end PE routers is used for the announcement of VPN-IPv4 routes. In this manner, VPN connectivity is provided while keeping VPN-IPv4 routes out of the core network.
https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-interprovider.html

QUESTION 11
Which three statements are true about the show pim join output shown in the exhibit? (Choose three.)

```
user@host> show pim join 234.100.0.1 extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 234.100.0.1
    Source: 192.168.100.2
    Flags: sparse, spt
    Active upstream interface: ge-1/0/0.0
    Active upstream neighbor: 192.168.101.2
    MoFRR Backup upstream interface: ge-1/0/1.0
    MoFRR Backup upstream neighbor: 192.168.102.2
    Upstream state: Join to Source, No Prune to RP
    Keepalive timeout: 300
    Uptime: 00:00:15
    Downstream neighbors:
        Interface: ge-1/2/0.0
            192.168.103.2 State: Join Flags: S  Timeout: Infinity
            Uptime: 00:00:15 Time since last Join: 00:00:15
Number of downstream interfaces: 1
```

A. This is a source-specific multicast stream.
B. The multicast stream does not have an RP.
C. The multicast stream has been configured with a backup path to allow for fast reroute.
D. The shortest path to the source is through the RP.
E. The multicast receiver is still using the RP to receive the stream.

Answer: CDE
Explanation:
PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in Table 1.

QUESTION 12

Referring to the exhibit, you need to implement VPLS between CE-A1 and CE-A2. You must ensure that no loops are created due to the multihoming of the connection from CE-A2 to PE2 and PE3.

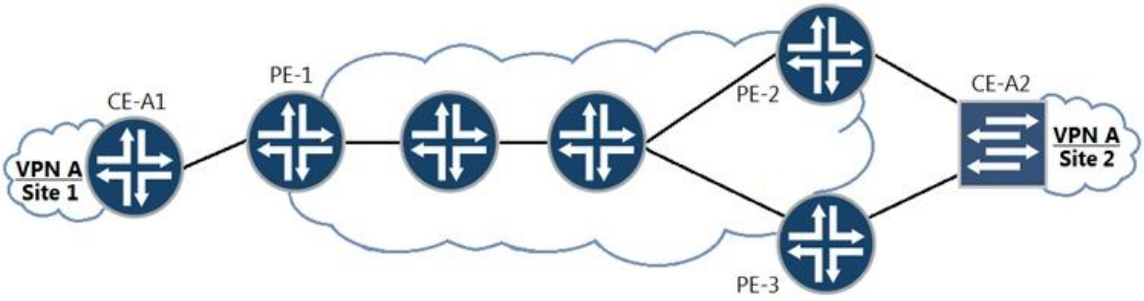

Based on the type of VPLS, which two solutions will satisfy this requirement? (Choose two.)

A. In a BGP VPLS, configure a primary and backup neighbor.
B. In an LDP VPLS. configure multihoming and local preference on PE-2 and PE-3
C. In an LDP VPLS, configure a primary and backup neighbor.
D. In a BGP VPLS, configure multihoming and local preference on PE-2 and PE-3.

Answer: AD
Explanation:
If you want to enable multihoming for a VPLS routing instance, you cannot also enable LDP signaling. You can only enable BGP signaling.
https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpns-configuring-vpls-multihoming.html

QUESTION 13
Which two statements are true about IS-IS adjacency formation? (Choose two.)

A. Level 1 only routers never form an adjacency with Level 2 only routers.
B. Level 1 only routers always form an adjacency with Level 2 only routers.
C. For Level 1 adjacencies, area IDs must be the same.
D. For Level 2 adjacencies, area IDs must be the same.

Answer: AC
Explanation:
An L1 router does not form an adjacency with an L2 router, regardless of area.
L1 routers form L1 adjacencies with L1 and L1-L2 routers in their area.
L2 routers form L2 adjacencies with L2 and L1-L2 routers in their area or another area.
http://www.ciscopress.com/articles/article.asp?p=730191&seqNum=7

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

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

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**