



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



JN0-649 Dumps
JN0-649 Braindumps
JN0-649 Real Questions
JN0-649 Practice Test
JN0-649 Actual Questions



killexams.com

Juniper

JN0-649

Enterprise Routing and Switching Professional (JNCIP-ENT)

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/JN0-649>



Question: 541

You are configuring a multicast network with PIM-SM and Auto-RP. The mapping agent configuration on Router R1 is:

```
ip pim send-rp-discovery Loopback0 scope 16
interface Loopback0
ip address 10.1.1.1 255.255.255.255
ip pim sparse-mode
```

A candidate RP (R2) is configured for group 239.10.10.10, but other routers show no RP mapping. The show ip pim rp mapping on R1 is empty. What is the most likely issue?

- A. The scope value is too low
- B. Auto-RP messages are filtered
- C. The candidate RP is not sending announcements
- D. PIM is disabled on R1's interfaces

Answer: B

Explanation: Auto-RP relies on the mapping agent (R1) receiving RP announcements from candidate RPs (R2) via 224.0.1.39 and distributing mappings via 224.0.1.40. If show ip pim rp mapping is empty, R1 is not receiving or processing these announcements. A common issue is a multicast boundary or access list filtering Auto-RP messages (224.0.1.39/40), preventing R1 from learning the RP. The scope value (16) is sufficient for campus networks, and PIM on Loopback0 is enabled. If the candidate RP were not sending announcements, only R2's groups would be affected, but an empty mapping suggests a broader issue. Thus, filtered Auto-RP messages are the most likely cause.

Question: 542

You are troubleshooting a connectivity issue in a data center where a Juniper QFX5100 switch is configured with access and trunk ports. Interface ge-0/0/10 is an access port in VLAN 50, and ge-0/0/11 is a trunk port carrying VLANs 50 and 60. A host connected to ge-0/0/10 cannot communicate with a server on VLAN 60 via ge-0/0/11. The configuration is correct, but the issue persists. What is the most likely cause?

- A. The trunk port is not tagging VLAN 60 traffic
- B. The access port is sending tagged frames

- C. An IRB interface is missing for VLAN 60
- D. The server is not configured to handle tagged traffic

Answer: D

Explanation: Since ge-0/0/10 is an access port in VLAN 50, it sends untagged frames, and ge-0/0/11 is a trunk port carrying VLANs 50 and 60, the switch configuration appears correct. For the host in VLAN 50 to communicate with the server in VLAN 60, the server must be configured to handle tagged traffic for VLAN 60, as the trunk port sends tagged frames. An IRB interface is only needed for inter-VLAN routing, not direct VLAN communication.

Question: 543

You are troubleshooting a performance issue on a Juniper QFX5100 switch where multicast traffic on interface xe-0/0/20.0 is experiencing drops. You use monitor traffic to capture 300 IGMP packets (protocol 2) and save them to "igmp_capture.pcap". Which command is correct?

- A. monitor traffic interface xe-0/0/20.0 matching "ip proto 2" count 300 write-file igmp_capture.pcap
- B. monitor traffic interface xe-0/0/20.0 matching "proto igmp" count 300 write-file igmp_capture.pcap
- C. monitor traffic interface xe-0/0/20.0 matching "ip igmp" count 300 write-file igmp_capture.pcap
- D. monitor traffic interface xe-0/0/20.0 matching "proto 2" count 300 write-file igmp_capture.pcap

Answer: A

Explanation: IGMP uses IP protocol 2. The monitor traffic command uses matching "ip proto 2" to capture IGMP packets, with count 300 and write-file igmp_capture.pcap to save 300 packets. Incorrect options use invalid match conditions (proto igmp, ip igmp, or proto 2 without ip).

Question: 544

In a complex OSPF topology, you are tasked with summarizing routes in Area 1 to reduce the LSDB size in Area 0. Router R1 is an Area Border Router (ABR) connecting Area 1 to Area 0. You configure route summarization on R1 for the prefix 172.16.0.0/16, but the summarized route is not appearing in Area 0. The exhibit shows the OSPF configuration on R1:

Exhibit:

```
protocols {
  ospf {
    area 0.0.0.1 {
      area-range 172.16.0.0/16;
      interface ge-0/0/1.0;
    }
    area 0.0.0.0 {
      interface ge-0/0/0.0;
    }
  }
}
```

What is the most likely reason the summarized route is not appearing in Area 0?

- A. The area-range command is applied to the wrong area
- B. The summarized prefix is not present in the R1 routing table
- C. The area-range command requires an explicit metric
- D. Area 1 is configured as a stub area, preventing summarization

Answer: A

Explanation: The area-range command for route summarization must be applied to the area where the routes originate (Area 1) but advertised into the backbone (Area 0). In the configuration, the area-range is incorrectly applied under Area 1, meaning it attempts to summarize routes within Area 1 rather than advertising the summary to Area 0. The summarized prefix must be present in the routing table, but this is not indicated as the issue. The area-range command does not require an explicit metric, and stub areas do not inherently prevent summarization unless misconfigured.

Question: 545

You are configuring MAC RADIUS authentication on an EX Series switch running Junos OS 21.2R2 for a device on interface ge-0/0/6 with MAC address 00:33:44:55:66:77. The RADIUS server is at 192.168.30.10, and you want to assign authenticated devices to VLAN 500. The exhibit shows the configuration:

```
set access radius-server 192.168.30.10 secret "macpass"  
set access profile mac-profile authentication-order radius  
set vlans vlan500 vlan-id 500
```

Which command enables MAC RADIUS with dynamic VLAN assignment?

- A. set protocols dot1x authenticator interface ge-0/0/6 mac-radius
- B. set protocols dot1x authenticator interface ge-0/0/6 vlan-assignment vlan500
- C. set protocols dot1x authenticator interface ge-0/0/6 static 00:33:44:55:66:77
- D. set services captive-portal interface ge-0/0/6 authentication-profile-name mac-profile

Answer: A

Explanation: MAC RADIUS authentication is enabled with the mac-radius option, and dynamic VLAN assignment is supported via RADIUS VSAs. The command set protocols dot1x authenticator interface ge-0/0/6 mac-radius enables MAC RADIUS authentication, allowing the RADIUS server to assign VLAN 500. The vlan-assignment command is for static VLANs, static bypasses authentication, and captive portal is unrelated.

Question: 546

A network engineer is configuring an OSPF network with a stub area (Area 10) and observes that

external routes redistributed by an ASBR in Area 0 are not appearing in the routing table of routers within Area 10. The ASBR is advertising a Type 5 LSA for the external prefix 192.168.1.0/24 with a metric of 100. The ABR connecting Area 0 to Area 10 is configured with the command `set protocols ospf area 0.0.0.10 stub default-metric 10`. The LSDB of a router in Area 10 shows a default route via the ABR but no Type 5 LSAs. What is the most likely reason for this behavior, and what configuration change would allow the external routes to appear in Area 10's routing table?

- A. Change the area type to NSSA using `set protocols ospf area 0.0.0.10 nssa`
- B. Remove the stub configuration with `delete protocols ospf area 0.0.0.10 stub`
- C. Add a summary LSA with `set protocols ospf area 0.0.0.10 area-range 192.168.1.0/24`
- D. Increase the default metric using `set protocols ospf area 0.0.0.10 stub default-metric 200`

Answer: A

Explanation: Stub areas do not allow Type 5 LSAs (external routes) to be flooded into them, which explains why the 192.168.1.0/24 prefix is absent in Area 10's routing table. Instead, the ABR injects a default route, as seen in the LSDB. Configuring Area 10 as a Not-So-Stubby Area (NSSA) allows external routes to be advertised as Type 7 LSAs within the area, which can be translated to Type 5 LSAs by the ABR for flooding into Area 0. Removing the stub configuration would make it a regular area, allowing Type 5 LSAs but also other LSA types, which may not be desired. Area-range is for summarization, not enabling external routes, and changing the default metric does not affect Type 5 LSA propagation.

Question: 547

You are configuring IGMP snooping in a Layer 2 network to optimize multicast traffic for a video streaming application using group 239.7.7.7. The switch connects to a PIM router via interface ge-0/0/1 and to receivers via ge-0/0/2. The configuration is: `set protocols igmp-snooping vlan 200 interface ge-0/0/1.0`. Receivers send IGMPv2 join messages, but the snooping table shows no entries, and traffic floods all ports in VLAN 200. The PIM router is sending IGMP queries. What is the most likely cause of the issue?

- A. IGMP snooping is disabled for VLAN 200
- B. The PIM router's IGMP version is incompatible
- C. The switch lacks an IGMP snooping querier
- D. The interface ge-0/0/2.0 is not IGMP snooping-enabled

Answer: D

Explanation: IGMP snooping requires all relevant interfaces in the VLAN to be configured for snooping to build the group membership table. The configuration only includes ge-0/0/1.0 (connected to the PIM router), omitting ge-0/0/2.0 (connected to receivers). As a result, the switch does not process IGMP joins from ge-0/0/2.0, causing the snooping table to remain empty and traffic to flood all ports in VLAN 200. IGMP snooping is enabled for VLAN 200, and the PIM router's queries indicate compatibility. A separate querier is unnecessary since the PIM router provides queries.

Question: 548

In a data center network, you are implementing ECMP load balancing on a Juniper QFX switch to distribute traffic across four equal-cost paths to the destination network 10.20.30.0/24. The switch uses a hash algorithm that includes Layer 3 and Layer 4 information. Which configuration under [edit forwarding-options] ensures that traffic is balanced based on source/destination IP addresses and TCP/UDP port numbers?

- A. enhanced-hash-key { family inet { layer-3; layer-4; } }
- B. load-balance { family inet { layer-3; layer-4; } }
- C. hash-key { family inet { layer-3; } }
- D. enhanced-hash-key { family inet { layer-3; } }

Answer: A

Explanation: ECMP load balancing in Junos OS uses a hash algorithm to distribute traffic across equal-cost paths. To include both Layer 3 (source/destination IP) and Layer 4 (TCP/UDP ports) information in the hash, the enhanced-hash-key configuration under [edit forwarding-options] is used with layer-3 and layer-4 options enabled for the inet family. Option A correctly configures this requirement.

Question: 549

You are designing a high-availability campus network with two MX960 routers configured for Virtual Router Redundancy Protocol (VRRP). Router R1 is the primary with VRRP priority 200, and Router R2 is the backup with priority 100. The VRRP group is configured on interface ge-0/0/0 with virtual IP 192.168.1.254. The configuration on R1 includes: set interfaces ge-0/0/0 unit 0 family inet address 192.168.1.1/24 vrrp-group 1 virtual-address 192.168.1.254 priority 200 preempt. During a network outage, R2 becomes primary, but when R1 recovers, it does not reclaim the primary role despite the higher priority. Which configuration change is required on R2 to allow R1 to reclaim the primary role, and how can you verify the VRRP state?

- A. Configure set interfaces ge-0/0/0 unit 0 family inet address 192.168.1.2/24 vrrp-group 1 virtual-address 192.168.1.254 priority 100 preempt on R2
- B. Verify VRRP state with show vrrp detail on both routers
- C. Remove the preempt knob from R2's VRRP configuration
- D. Check interface status with show interfaces ge-0/0/0 terse to confirm IP addressing

Answer: A, B

Explanation: For R1 to reclaim the primary VRRP role upon recovery, both routers must have the preempt option configured, allowing the router with the higher priority to take over. On R2, adding preempt to the VRRP configuration ensures this behavior. The show vrrp detail command verifies the VRRP state, showing the current primary, priority, and preemption settings on both routers. Removing

the preempt knob from R2 would prevent preemption entirely, which is not desired. Checking interface status confirms IP addressing but does not verify VRRP-specific states.

Question: 550

In a data center running Contrail Enterprise Multicloud, you are implementing a YANG-based configuration management system using NETCONF to manage QFX switches. The YANG model defines a custom RPC to retrieve EVPN MAC table information. After deploying the RPC, you notice that the NETCONF client receives incomplete data, missing some MAC addresses. What is the most likely cause of this issue?

- A. The YANG model lacks a list statement for the MAC table entries
- B. The NETCONF session is using an outdated Junos OS version
- C. The RPC is not filtering the MAC table by VNI
- D. The Contrail Controller is overriding the MAC table updates

Answer: A

Explanation: In YANG, a list statement is used to define repeating elements, such as MAC table entries. If the YANG model does not include a list for MAC table entries, the RPC may return incomplete or incorrect data. The other options are less likely to cause missing MAC addresses in the NETCONF response.

Question: 551

In an enterprise network, you are troubleshooting a BGP session that is in the OpenConfirm state. The network uses a confederation (AS 65000, sub-AS 65001) and includes flap damping and graceful restart. The exhibit shows the BGP configuration. What could be causing the issue?

[Exhibit: BGP Configuration]

```
protocols {
  bgp {
    group CONFED {
      type external;
      neighbor 10.1.1.2 {
        peer-as 65002;
      }
    }
  }
}
```

- A. A firewall is blocking keepalives
- B. The peer AS is incorrect

- C. Flap damping is suppressing the session
- D. The local router ID is not configured

Answer: A

Explanation: A BGP session in the OpenConfirm state is waiting for a keepalive or update message to transition to Established. A firewall blocking keepalives can prevent this transition. An incorrect peer AS would cause the session to fail in OpenSent. Flap damping affects route advertisement, not session establishment. A missing router ID would affect the OpenSent state.

Question: 552

A Juniper EX9200 switch is configured with Multiple Spanning Tree Protocol (MSTP) to prevent loops in a network with VLANs 10, 20, and 30. The MSTP configuration includes two instances: MSTI 1 for VLAN 10 and MSTI 2 for VLANs 20 and 30. The switch is experiencing unexpected traffic drops due to incorrect MSTP convergence. The configuration is shown below. What is the likely cause of the issue?

```
set protocols mstp configuration-name region1
set protocols mstp msti 1 vlan 10
set protocols mstp msti 2 vlan [20 30]
set protocols mstp bridge-priority 4096
```

- A. The bridge priority is too high, causing the switch to lose the root election
- B. The configuration-name is inconsistent across switches in the region
- C. VLANs 20 and 30 should be in separate MSTIs for better load balancing
- D. The MSTP protocol is not enabled on all trunk interfaces

Answer: B

Explanation: In MSTP, all switches in the same region must have the same configuration-name, revision level, and VLAN-to-MSTI mappings. If the configuration-name region1 is not identical across all switches, they form separate MST regions, leading to incorrect spanning tree calculations and potential traffic drops. The bridge priority, VLAN mappings, and interface enablement are secondary concerns if the region configuration is misaligned.

Question: 553

In a multi-tenant data center, you are configuring PIM Sparse Mode with Source-Specific Multicast (SSM) for a secure application using group 232.1.1.1. Receivers send IGMPv3 include-mode join messages specifying the source 192.168.30.30. The mroute table on the receiver's router R2 shows no (S, G) entry, despite correct IGMP joins. The configuration on R2 includes: set protocols pim ssm-groups 232.0.0.0/8. The unicast route to 192.168.30.30 is valid, and PIM is enabled on all relevant interfaces. What is the most likely reason for the missing mroute entry?

- A. The SSM group range is misconfigured on R2
- B. The receivers are using an incorrect IGMP version
- C. The source is not sending traffic to the group

D. The RPF interface is not PIM-enabled

Answer: C

Explanation: In SSM, receivers explicitly join a (S, G) channel using IGMPv3, and the router builds an (S, G) mroute entry only when traffic from the specified source is received. If the mroute table lacks an (S, G) entry despite valid IGMP joins and correct unicast routing, the most likely cause is that the source (192.168.30.30) is not sending traffic to the group (232.1.1.1). The SSM group range (232.0.0.0/8) is correct, as 232.1.1.1 falls within it. IGMPv3 is required for SSM and is confirmed by the include-mode joins. The RPF interface must be PIM-enabled for joins to be processed, which is implied by the valid setup.

Question: 554

An IS-IS network has a Level 2 router redistributing a static route 172.16.4.0/24 with a metric of 50. The command show isis database detail on a neighboring router shows the prefix with a metric of 60. The link between the routers has a default metric of 10. What configuration change would ensure the neighboring router sees the metric as 50?

- A. Configure set protocols isis interface ge-0/0/0.0 level 2 metric 0
- B. Enable wide metrics with set protocols isis level 2 wide-metrics-only
- C. Modify the redistribution policy to set an internal metric
- D. Disable adjacency with set protocols isis interface ge-0/0/0.0 level 2 disable

Answer: A

Explanation: The metric of 60 includes the redistributed metric (50) plus the link metric (10). Setting the link metric to 0 ensures the neighboring router sees only the redistributed metric of 50. Wide metrics don't eliminate link costs, and changing to an internal metric doesn't address link metric accumulation. Disabling the adjacency would prevent all communication.

Question: 555

In a high-availability enterprise network running Junos OS, you are configuring Graceful Routing Engine Switchover (GRES) on a dual Routing Engine system to ensure minimal disruption during a switchover. The system uses MX480 routers with Routing Engine 0 as primary and Routing Engine 1 as backup. You have enabled GRES and synchronized the configuration, but during a manual switchover test, you observe that some OSPF adjacencies briefly drop before re-establishing. The network topology includes multiple OSPF areas with area 0 as the backbone, and the router is configured with the following: set chassis redundancy graceful-switchover and set routing-options nonstop-routing. Which additional configuration is required to prevent OSPF adjacency drops during the GRES switchover, and what is the correct sequence of steps to verify the GRES state post-switchover?

- A. Configure set protocols ospf graceful-restart to enable OSPF graceful restart

- B. Verify GRES readiness with show chassis routing-engine and check for "Backup" state on Routing Engine 1
- C. Enable set system commit synchronize to ensure configuration synchronization between Routing Engines
- D. Check GRES synchronization with show system switchover on the backup Routing Engine

Answer: A, D

Explanation: To prevent OSPF adjacency drops during a GRES switchover, enabling OSPF graceful restart is necessary to maintain neighbor relationships by allowing the router to inform neighbors it is undergoing a restart, preserving adjacency states. The configuration set protocols ospf graceful-restart achieves this. Additionally, verifying GRES synchronization is critical post-switchover. The show system switchover command on the backup Routing Engine confirms that the kernel state and forwarding state are synchronized, ensuring GRES is functioning correctly. The show chassis routing-engine command shows the state of Routing Engines but does not specifically verify GRES synchronization. Configuration synchronization via set system commit synchronize is already implied as enabled for GRES to work but is not directly related to preventing OSPF drops.

Question: 556

To secure a Layer 2 network on a Juniper EX9200 switch, you configure storm control and 802.1X authentication on interface ge-0/0/4. The configuration is:

```
set interfaces ge-0/0/4 unit 0 family ethernet-switching storm-control bandwidth-percentage 10
set protocols dot1x authenticator interface ge-0/0/4 supplicant single
```

During a broadcast storm, the interface exceeds the storm control threshold, and a device fails 802.1X authentication. Which two outcomes occur?

- A. The interface drops excess broadcast traffic.
- B. The device is denied network access.
- C. The interface is shut down due to storm control.
- D. The device is placed in a guest VLAN.

Answer: A, B

Explanation: Storm control limits broadcast, unknown unicast, and multicast traffic to 10% of the interface bandwidth, dropping excess traffic without shutting down the interface unless explicitly configured (e.g., action shutdown). The dot1x configuration with supplicant single requires 802.1X authentication; a failed authentication denies network access unless a guest VLAN is configured, which is not indicated here. Thus, excess broadcast traffic is dropped, and the unauthenticated device is blocked.

Question: 557

You are implementing DHCP snooping on an EX Series switch running Junos OS 20.4R3 in VLAN 1100. The DHCP server is on interface ge-0/0/6, and clients are on ge-0/0/7 to ge-0/0/10. The exhibit shows the configuration:

```
set vlans vlan1100 vlan-id 1100
set ethernet-switching-options dhcp-snooping vlan vlan1100
```

Which command ensures the DHCP server's messages are processed correctly?

- A. set ethernet-switching-options dhcp-snooping vlan vlan1100 interface ge-0/0/6 trusted
- B. set ethernet-switching-options dhcp-snooping vlan vlan1100 no-option-82
- C. set interfaces ge-0/0/6 unit 0 family ethernet-switching dhcp-trusted
- D. set ethernet-switching-options dhcp-snooping vlan vlan1100 examine-dhcp disable

Answer: A

Explanation: The DHCP server interface must be trusted to allow its messages to populate the snooping database. The command `set ethernet-switching-options dhcp-snooping vlan vlan1100 interface ge-0/0/6 trusted` achieves this. Disabling option-82 or DHCP inspection is unnecessary, and `dhcp-trusted` is not a valid command.

Question: 558

You are tasked with setting up BGP in a network that includes both iBGP and eBGP peers. You need to ensure that routing information is correctly propagated within the AS while also adhering to best practices. Which of the following statements accurately describe the rules for iBGP and eBGP peering and the use of route reflectors for scalability?

- A. iBGP requires a full mesh of peers, while eBGP does not.
- B. Route reflectors can break the full mesh requirement of iBGP by allowing route advertisement among clients.
- C. eBGP peers must be directly connected.
- D. Route reflectors can only be used within the same AS.

Answer: A, B, D

Explanation: iBGP typically requires a full mesh to avoid routing loops, while eBGP does not have this restriction. Route reflectors allow the elimination of the full mesh requirement and can operate within the same AS, making them essential for scalability in larger networks.

Question: 559

In an OSPF network, you are configuring route redistribution on an ASBR (R1) to inject BGP routes into OSPF. The BGP routes include a prefix 203.0.113.0/24 with a community tag 65000:100. You want to ensure that only routes with this community are redistributed into OSPF as Type 5 LSAs with a metric

of 50. The OSPF domain includes Area 0 and Area 1, with R1 in Area 0. Which configuration on R1 achieves this requirement?

- A. set policy-options policy-statement redistrib term 1 from community 65000:100 then metric 50 accept
- B. set protocols ospf export metric 50 community 65000:100
- C. set protocols ospf area 0 interface lo0.0 community 65000:100
- D. set policy-options community 65000:100 members 65000:100

Answer: A

Explanation: To filter BGP routes for redistribution into OSPF based on a community, a policy-statement is used. The configuration set policy-options policy-statement redistrib term 1 from community 65000:100 then metric 50 accept matches routes with community 65000:100, sets the metric to 50, and accepts them for redistribution as Type 5 LSAs. Other options either misapply communities or lack policy control.

Question: 560

You are managing a BGP environment with multiple paths to the same destination across different ISPs. To optimize traffic distribution without compromising redundancy, you decide to implement BGP multipath. Which of the following configurations are necessary to enable BGP multipath and ensure that load balancing occurs effectively across multiple paths while maintaining optimal path selection based on the BGP path selection process?

- A. Configure the “bgp bestpath multipath” command in the BGP configuration.
- B. Ensure that all paths have the same local preference value.
- C. Enable CEF (Cisco Express Forwarding) to support load balancing.
- D. Set the maximum number of paths to be used in load balancing to a specific value.

Answer: A, B, C, D

Explanation: To enable BGP multipath, the bestpath multipath command must be configured, and it's crucial for paths to have the same local preference for them to be eligible for load balancing. CEF must also be enabled to facilitate load balancing, and setting a maximum number of paths helps control the distribution across multiple paths.



KILLEXAMS.COM

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:



Actual Exam Questions: Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.

Exam Dumps: Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.

Practice Tests: Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success: Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.

Updated Content: Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.

Technical Support: Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.