

Juniper.JN0-364.v2026-05-20.q85

|                                                                                                                                                   |                                                               |
|---------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------|
| Exam Code:                                                                                                                                        | JN0-364                                                       |
| Exam Name:                                                                                                                                        | Service Provider Routing and Switching, Specialist (JNCIS-SP) |
| Certification Provider:                                                                                                                           | Juniper                                                       |
| Free Question Number:                                                                                                                             | 85                                                            |
| Version:                                                                                                                                          | v2026-05-20                                                   |
| # of views:                                                                                                                                       | 163                                                           |
| # of Questions views:                                                                                                                             | 850                                                           |
| <a href="https://www.freepdfdumps.com/Juniper.JN0-364.v2026-05-20.q85.html">https://www.freepdfdumps.com/Juniper.JN0-364.v2026-05-20.q85.html</a> |                                                               |

NEW QUESTION: 1

Click the Exhibit button.

```
user@Router-1> show route 172.24/16
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
...
172.24.0.0/24      *[OSPF/150] 01:31:31, metric 0, tag 0
                   > to 172.20.0.2 via ge-0/0/2.0
                   to 172.20.1.2 via ge-0/0/3.0
...
user@Router-1> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
...
172.24.0.0/24     user      0          172.20.0.2    ulst 262142 4
                  172.20.1.2    ucst  551    2 ge-0/0/2.0
                  172.20.1.2    ucst  552    2 ge-0/0/3.0
...

```

Referring to the exhibit, which two statements are true? (Choose two.)

- A. This router will choose both next hops in the routing table.
- B. The router is performing default route load-balancing behavior.
- C. The default route load-balancing behavior of this router has been modified.
- D. This router will only choose the next hop with a > next to it in the routing table.


Answer: A,B (LEAVE A REPLY)

The routing table shows two equal-cost next hops for the prefix and both are installed in the forwarding table, which means the router is selecting both next hops for forwarding traffic. This indicates equal-cost multipath forwarding.

The forwarding table entry shows a load-balancing group containing both next hops, which reflects the default Junos behavior of performing load balancing across equal-cost paths when multiple next hops exist for the same route.

### NEW QUESTION: 2

Referring to the exhibit. You want the R1 and R3 routers to forward traffic destined to the 203.0.113.0/24 network through R2. Which BGP attribute would you modify to satisfy this requirement?



- A. community
- B. origin
- C. MED
- D. local preference

**Answer:** ([SHOW ANSWER](#))

To influence the path taken through an autonomous system, the Multi-Exit Discriminator (MED) attribute can be used. By setting a lower MED on R2 for routes advertised to R1 and R3, it will make the path through R2 to the 203.0.113.0/24 network more preferable.

### NEW QUESTION: 3

You are troubleshooting two OSPF routers that have an adjacency that remains in the ExStart state. What would cause this problem?

- A. mismatched OSPF hello intervals on the OSPF interfaces
- B. mismatched authentication settings on the OSPF interfaces
- C. mismatched MTU settings on the OSPF interfaces
- D. mismatched subnet settings on the OSPF interfaces

**Answer:** C ([LEAVE A REPLY](#))

The problem occurs most frequently when you attempt to run OSPF between a Cisco router and another vendor router. The problem occurs when the maximum transmission unit (MTU) settings for neighboring router interfaces do not match. If the router with the higher MTU sends a packet larger than the MTU set on the neighboring router, the neighbor router ignores the packet. A mismatched Maximum Transmission Unit (MTU) setting on OSPF interfaces can cause routers to get stuck in the ExStart state. OSPF requires that both sides of a link have the same MTU to form a full adjacency.

**NEW QUESTION: 4**

What is the default export behavior of IS-IS in the Junos OS?

- A. to export nothing
- B. to export all learned prefixes
- C. to export only IPv6 routes
- D. to export only external routes

**Answer: A** ([LEAVE A REPLY](#))

In Junos OS, IS-IS does not automatically export routes from the routing table into the IS-IS protocol. By default, no routes are advertised into IS-IS unless an explicit export policy is configured to control which prefixes should be advertised.

**NEW QUESTION: 5**

What are two types of BGP messages exchanged while in the Established state? (Choose two.)

- A. request
- B. open
- C. update
- D. notification

**Answer: C,D** ([LEAVE A REPLY](#))

Update messages are exchanged between BGP peers in the Established state to advertise new routes, withdraw previously advertised routes, and share path attributes for reachable prefixes.

Notification messages are used to signal errors detected in the BGP session and immediately terminate the session when a protocol error occurs.

**NEW QUESTION: 6**

Referring to the exhibit. Which prefix in the output shown in the exhibit is an external prefix injected by an OSPF router?

`user@R2> show ospf route`

Topology default Route Table:

| Prefix         | Path<br>Type | Route<br>Type | NH<br>Type | Metric | NextHop<br>Interface | NextHop<br>addr/label |
|----------------|--------------|---------------|------------|--------|----------------------|-----------------------|
| 192.168.1.1    | Intra        | AS BR         | IP         | 1      | ge-0/0/3.0           | 172.26.1.1            |
| 192.168.1.3    | Intra        | Area BR       | IP         | 1      | ge-0/0/1.0           | 172.26.2.2            |
| 172.18.1.0/24  | Ext2         | Network       | IP         | 0      | ge-0/0/3.0           | 172.26.1.1            |
| 172.26.1.0/30  | Intra        | Network       | IP         | 1      | ge-0/0/3.0           |                       |
| 172.26.2.0/30  | Intra        | Network       | IP         | 1      | ge-0/0/1.0           |                       |
| 172.26.3.0/30  | Intra        | Network       | IP         | 100    | ge-0/0/2.0           |                       |
| 172.26.4.0/30  | Inter        | Network       | IP         | 2      | ge-0/0/1.0           | 172.26.2.2            |
| 192.168.1.1/32 | Ext2         | Network       | IP         | 1      | ge-0/0/3.0           | 172.26.1.1            |
| 192.168.1.2/32 | Intra        | Network       | IP         | 0      | lo0.0                |                       |
| 192.168.1.3/32 | Intra        | Network       | IP         | 1      | ge-0/0/1.0           | 172.26.2.2            |
| 192.168.1.4/32 | Inter        | Network       | IP         | 2      | ge-0/0/1.0           | 172.26.2.2            |

- A. 192.168.1.3
- B. 172.18.1.0/24
- C. 192.108.1.4
- D. 172.26.4.0730

**Answer:** (SHOW ANSWER)

In the OSPF routing table output, prefixes are marked with different route types. An external prefix injected into OSPF is marked as 'Ext' (External) followed by a number that indicates whether it's an E1 or E2 route. The prefix 172.18.1.0/24 is marked as Ext2, which indicates that it is an external route that has been redistributed into OSPF from another routing protocol or static configuration.

#### NEW QUESTION: 7

A service provider is onboarding a new enterprise customer that operates multiple branch offices, each with its own set of VLANs. The customer requires transparent Layer 2 connectivity between sites while maintaining separation of internal VLANs. The provider must also ensure that customer VLAN identifiers do not conflict with other customers on the shared infrastructure.

Which solution would provide the desired results?


- A. Deliver Layer 3 VPN services using MPLS.
- B. Aggregate customer traffic using GRE tunnels.
- C. Provide Internet access with NAT and firewall services.
- D. Extend customer VLANs using Q-in-Q tunneling.

**Answer:** D (LEAVE A REPLY)

Q-in-Q tunneling (802.1ad) allows a service provider to encapsulate customer VLAN tags inside an additional provider VLAN tag. This enables transparent Layer 2 connectivity between customer sites while preserving the customer's internal VLAN structure. The outer service- provider VLAN tag ensures separation between different customers and prevents VLAN ID conflicts on the shared provider infrastructure.

#### NEW QUESTION: 8

Referring to the exhibit. Which two statements are correct about the service provider MPLS network shown in the exhibit? (Choose two.)



- A. R3 will perform a label pop operation on the transport MPLS label.
- B. Traffic from Network 1 to Network 3 and traffic from Network 1 to Network 4 each need their own unique label-switched path.
- C. Traffic from Network 1 to Network 3 and from Network 1 to Network 4 can share the same label- switched path.
- D. R3 will perform a label swap operation on the transport MPLS label.

**Answer:** C,D (LEAVE A REPLY)

In MPLS, multiple paths can be merged if they share the same egress router. In the given scenario, traffic from Network 1 to Network 3 and Network 4 can be engineered to follow the same label- switched path (LSP) within the MPLS network until they reach the last common point before diverging to their respective destinations.

As for R3 performing label operations, in a typical MPLS network, intermediate routers (like R3) perform label swapping. They replace the incoming label with a new label before forwarding the packet along the LSP. A label pop operation is typically performed by the egress router in the case of an ultimate hop pop (UHP), where it removes the MPLS label before delivering the packet to the final destination outside the MPLS domain.

#### **NEW QUESTION: 9**

You are the administrator for two Junos routers called R1 and R2. These two routers are directly connected to each other. These two routers run IS-IS and BFD. R1 is configured to send BFD packets every 300 milliseconds. R2 is configured to send BFD packets every 400 milliseconds.

In this situation, what is the expected outcome?

- A.** BFD will fail due to the mismatched timers.
- B.** Each router will negotiate to send BFD packets at the fastest of the two rates.
- C.** Each router will negotiate to send BFD packets at the slowest of the two rates.
- D.** Each router will send BFD packets at the rate that has been locally configured.

**Answer:** ([SHOW ANSWER](#))

With BFD, each side advertises its desired transmit interval, but the session uses an effective transmit/receive behavior that must be acceptable to both ends. The receiving side controls how fast it is willing to be polled by advertising its required minimum receive interval, which results in the effective rate converging to the slower (larger interval) value so that neither router is forced to receive BFD control packets faster than it supports.

#### **NEW QUESTION: 10**

Interface ge-0/0/0.0 connects your network to your ISP. You want to advertise this interface address as an Internal route In OSPF without creating a neighbor with your ISP. In this scenario, how is this task accomplished?

- A.** Add ge-0/0/0.0 as a passive interface In OSPF.
- B.** Configure a static route for Interface ge-0/0/0.0.
- C.** Remove interface ge-0/0/0.0 from OSPF.
- D.** Create a generated route for Interface ge-0/0/0.0.

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 11**

Click the Exhibit button.

```

user@R1> show configuration protocols bgp
group IBGP_CORE {
    type internal;
    local-address 192.168.1.1;
    export NEXT_HOP_SELF;
    neighbor 192.168.1.5;
}
group CUSTOMER_A {
    type external;
    peer-as 65101;
    neighbor 10.1.101.101;
}
user@R1> show configuration policy-options
policy-statements NEXT_HOP_SELF {
    term SELF {
        then {
            next-hop self;
        }
    }
}

```

Which statement is correct about the configuration shown in the exhibit?

- A. The NEXT\_HOP\_SELF policy must also be applied to the group called CUSTOMER\_A for bidirectional communication to succeed.
- B. This router will take prefixes learned from external BGP and change the BGP next hop to itself before re-advertising them to IBGP peers.
- C. This policy will reject all prefixes because it does not end with a then accept statement.
- D. This router will take prefixes learned from internal BGP and change the BGP next hop to itself before re-advertising them to EBGp peers.

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 12

When would you use the qualified-next-hop statement with a static route?

- A. You can use it to install the static route into different routing tables.
- B. You can use it to send unwanted traffic to a null route.
- C. You can use it to specify multiple next hops with different preferences.
- D. You can use it to resolve the next hop if the next hop is not directly connected.

**Answer: C** ([LEAVE A REPLY](#))

Qualified next hops allow you to associate one or more properties with a particular next-hop address. You can set an overall preference for a particular static route and then specify a different preference for the qualified next hop. For example, suppose two next-hop addresses (10.10.10.10 and 10.10.10.7) are associated with the static route 192.168.47.5/32. A general preference is assigned to the entire static route, and then a different preference is assigned to only the qualified next-hop address 10.10.10.7. For example:

The qualified-next-hop statement with a static route is used to specify multiple next hops for a static route with different preferences (priorities). This allows for more granular control over the path selection process in the event that the primary next hop becomes unreachable.

**NEW QUESTION: 13**

The segment routing SRGB start label is 10,000 and the SRGB index range is 500. In this scenario, which two statements are correct? (Choose two.)

- A. The first usable label is 10,001.
- B. The last usable label is 10,501.
- C. The last usable label is 10,499.
- D. The first usable label is 10,000.

**Answer: A,C** ([LEAVE A REPLY](#))

The Segment Routing Global Block (SRGB) defines a range of MPLS labels reserved for segment routing. With an SRGB starting label of 10,000 and an index range of 500, the first usable label is indeed 10,001 because label 10,000 is typically reserved and not used for forwarding. The last usable label would be 10,499 because the range includes 500 labels starting from 10,001.

**NEW QUESTION: 14**

By default, which routing table contains a list of all ingress LSPs?

- A. inet.0
- B. inet.1
- C. inet.2
- D. inet.3

**Answer: A** ([LEAVE A REPLY](#))

Ingress LSPs originate on the router and are installed as routes in the main IPv4 unicast routing table. This table contains the destination prefixes and next-hop information used for normal IP forwarding. When an LSP is configured as the forwarding path, the route in this table points to the LSP as the next hop, making it the table that lists ingress LSP routes by default.

**NEW QUESTION: 15**

How does a Junos device learn about MAC addresses when it is first connected to an Ethernet LAN?

- A. The device sends out a network broadcast message asking for all devices and MAC addresses on the network and stores this information in addition to the interface from which the response was received.
- B. The device learns the destination MAC addresses from traffic in the network and stores this MAC address in addition to the interface from which the traffic was received.
- C. The device learns the source MAC addresses from traffic in the network and stores this MAC address in addition to the interface from which the traffic was received.
- D. The device sends out a network multicast message asking for all devices and MAC addresses on the network and stores this information in addition to the interface from which the response was received.

**Answer:** ([SHOW ANSWER](#))

Junos devices learn about MAC addresses by examining the source MAC address of frames received on an interface and then storing these MAC addresses in the MAC address table (also known as the forwarding database) along with the interface on which the traffic was received.

**NEW QUESTION: 16**

The MPLS Label Information Base (LIB) is stored in which table?

- A. inet6.0
- B. mpls.0
- C. inet.0
- D. inet.3

**Answer: B** ([LEAVE A REPLY](#))

The MPLS Label Information Base is maintained in the mpls.0 routing table. This table stores the mappings between MPLS labels and their corresponding forwarding actions, allowing the router to perform label switching for MPLS packets.

**Valid JN0-364 Dumps** shared by Actual4test.com for Helping Passing JN0-364 Exam! Actual4test.com now offer the **newest JN0-364 exam dumps**, the Actual4test.com JN0-364 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com JN0-364 dumps with Test Engine here:

[https://www.actual4test.com/JN0-364\\_examcollection.html](https://www.actual4test.com/JN0-364_examcollection.html) (180 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

#### NEW QUESTION: 17

During MPLS forwarding, which operation describes the act of imposing a label onto an unlabeled packet?

- A. label swap
- B. label pop
- C. label push
- D. label rewrite

**Answer: (**[SHOW ANSWER](#)**)**

Label push is the MPLS operation in which a label is added to an unlabeled packet as it enters the MPLS domain. The ingress label switching router assigns the label and places it on the packet so that subsequent routers can forward it based on the MPLS label instead of the IP header.

#### NEW QUESTION: 18

Which two interface types are used as tunnel endpoints? (Choose two.)

- A. ae
- B. ip
- C. ge
- D. gr

**Answer: B,D** ([LEAVE A REPLY](#))

```
tunnel-end-point name {  
  ipv4 {  
    source-address 10.255.1.1;  
    destination-address 10.255.2.0/25;  
  }  
  gre {  
    key 9;  
  }  
}
```

}

The ip and gr interface types are used as endpoints for various tunneling protocols. ip is a generic term that can refer to any IP-based tunnel, while gr specifically refers to Generic Routing Encapsulation (GRE) tunnel interfaces in Junos.

#### NEW QUESTION: 19

You are implementing traffic engineering in your MPLS network. You must ensure that the MPLS routes are used to traverse your network. Your solution should not affect IGP routes in your route tables.

In this scenario, which traffic engineering setting will accomplish this behavior?

- A. bgp-igp-both-ribs
- B. mpls-forwarding
- C. bgp-igp
- D. bgp


**Answer:** ([SHOW ANSWER](#))

To ensure that MPLS traffic-engineered routes are used without affecting IGP routes in the route tables, the mpls-forwarding traffic engineering setting should be used. This setting allows MPLS to influence forwarding without changing the IGP route selection.

#### NEW QUESTION: 20

Click the Exhibit button. You must configure the router called ROUTER\_1 to take all valid prefixes learned from internal BGP peers in AS 64523, and then re-advertise them to other internal BGP peers in the same autonomous system.

Referring to the exhibit, which configuration must you deploy on ROUTER\_1 to accomplish this task?



- A. Configure ROUTER\_1 to belong to a different autonomous system than the other BGP routers in your network.
- B. Configure a routing policy on ROUTER\_1 that removes the no-exportBGP community from all received prefixes.
- C. Configure ROUTER\_1's internal BGP group with a routing policy that exports prefixes learned from internal BGP.
- D. Configure ROUTER\_1's internal BGP group with the keyword cluster, followed by a unique 32-bit number.

**Answer:** D ([LEAVE A REPLY](#))

To re-advertise routes learned from iBGP peers to other iBGP peers within the same AS, ROUTER\_1 must act as a route reflector. On Junos, enabling route reflection is done by configuring a cluster ID under the internal BGP group. The cluster ID uniquely identifies the route- reflector cluster and allows ROUTER\_1 to reflect iBGP-learned routes to its other iBGP neighbors.

#### NEW QUESTION: 21

You want to share routes between two routing instances that you have configured? What are two ways to accomplish this task? (Choose two.)

- A. Use a non-forwarding instance.
- B. Configure an instance import policy
- C. Create a forwarding instance.
- D. Use a RIB group.

**Answer: B,D** ([LEAVE A REPLY](#))

Static route with a next-hop of next-table pointing to the appropriate routing table which contains more accurate information rib-groups to mirror routing information from one route-table to another. However, in many cases, in order to make this work, interface-routes also need to be mirrored. RIB Group policy can be used to constrain the routing information instance-import and instance-export statements configured within the individual routing-instances to leak routes from one table to another. Again, policy can be used here to constrain the routing information. This method is more straightforward than the rib-group method A final approach is to use physical interfaces or logical- tunnels to stitch routing-instances and use a routing protocol or static routes across this connection between the two routing-instances.

To share routes between two routing instances on a Junos device, you can configure an instance import policy in one or both instances to import routes from the other instance. Alternatively, a RIB (Routing Information Base) group can be used to share routes between instances.

#### NEW QUESTION: 22

You are configuring BGP for IPv6 operations.

In this scenario, which two statements are correct? (Choose two.)

- A. The Autonomous System Number (ASN) can be either a 32-bit or 64-bit value.
- B. The router ID uses a 128-bit identifier value.
- C. The router ID uses a 32-bit identifier value.
- D. The Autonomous System Number (ASN) must be a 64-bit value.

**Answer: (**[SHOW ANSWER](#)**)**

BGP uses a router ID that is always a 32-bit value, even when BGP is operating for IPv6 address families.

BGP supports both 16-bit and 32-bit Autonomous System Numbers through the use of four-byte ASN extensions, allowing ASNs to be represented using a 32-bit value.

#### NEW QUESTION: 23

Referring to the exhibit. You configured interface ge-0/070.0 to run IS-IS. but this interface does not appear in the output of the show isis adjacency command as shown in the exhibit.

What is the problem in this scenario?

```

root@R1> show configuration protocols isis
interface ge-0/0/0.0 {
}
interface ge-0/0/1.0 {
}
interface lo0.0 {
level 1 disable;
level 2 wide-metrics-only;
reference-bandwidth 100g;
root@R1> show configuration interfaces ge-0/0/0
unit 0 {
    family inet {
        address 10.1.2.1/30;
    }
    family inet6;
    family mpls;
}
root@R1> show isis adjacency
Interface      System      L State      Hold (secs) SNPA
ge-0/0/1.0     R6          2 Up          19

```

- A. This is a Gigabit Ethernet interface, that is incompatible with the reference-bandwidth 100g statement.
- B. The family iso statement must be added to the logical interface.
- C. The router at the other end of the link is not sending any IS-IS Hello messages.
- D. The router at the other end of the link is a Level 1 only router.

**Answer: B** ([LEAVE A REPLY](#))

For IS-IS to form adjacencies, the 'family iso' statement must be configured on the interface. This is required because IS-IS operates over the connectionless network service (CLNS) and not over IP. Without the 'family iso' configuration, IS-IS cannot form adjacencies over the interface.

#### NEW QUESTION: 24

Which two LSA types are permitted in OSPF totally stubby areas? (Choose two.)

- A. Type 3
- B. Type 7
- C. Type 1
- D. Type 5

**Answer: A,D** ([LEAVE A REPLY](#))

#### NEW QUESTION: 25

Click the Exhibit button.

```

sereki> show route 10.16.2.0/23 exact detail
net.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
0.16.2.0/23 (1 entry, 1 announced)
  *Aggregate Preference: 130
    Next hop type: Reject
    Address: 0x8f3fd44
    Next-hop reference count: 2
    State: <Active Int Ext>
    Age: 1:39:21
    Task: Aggregate
    Announcement bits (1): 0-KRT
    AS path: I (LocalAgg)
    Flags:                               Depth: 0           Active
    AS path list:
    AS path: I Refcount: 2
    Contributing Routes (2):
      10.16.2.0/24 proto Direct
      10.16.3.0/24 proto Direct

```

Which destination IP address will be matched by the aggregate route shown in the exhibit?

- A. packets destined to 10.16.1.214
- B. packets destined to 10.16.0.4
- C. packets destined to 10.16.4.183
- D. packets destined to 10.16.3.79


**Answer:** ([SHOW ANSWER](#))

The aggregate route is 10.16.2.0/23. A /23 prefix covers the address range from 10.16.2.0 through 10.16.3.255. Therefore, any destination within that range will match the aggregate route.

The address 10.16.3.79 falls within this range, so packets destined for that address match the aggregate route.

## NEW QUESTION: 26

Click the Exhibit button.



Referring to the exhibit, R1 and R2 are advertising the same prefix 203.0.113.0/24 to R3 and R4 over EBGP. R3 and R4 both advertise this prefix to R5.

Which advertisement does R5 choose to install in its routing table?

- A. The advertisement from R3 is chosen.
- B. The advertisements from both R3 and R4. but R4 is chosen for forwarding.
- C. The advertisement from R4 is chosen.
- D. The advertisements from both R3 and R4. but R3 is chosen for forwarding.

**Answer: (SHOW ANSWER)**

R3 and R4 both advertise the same prefix into iBGP toward R5, but they assign different local preference values. R3 sets a local preference of 150, while R4 sets a local preference of 200.

Since BGP selects routes with the highest local preference before considering other attributes, R5 chooses the route received from R4 and installs it in its routing table.

#### NEW QUESTION: 27

Which two statements are correct when using LDP? (Choose two.)

- A. The Inet.3 table will contain only the paths explicitly defined.
- B. The inet.3 table will contain a full mesh of label-switched paths to other LDP-enabled routers.
- C. LDP label-switched paths are created by configuring LDP on at least one physical router interface.
- D. LDP label-switched paths are created by configuring LDP on the loopback Interface.

**Answer: B,C (LEAVE A REPLY)**

With Label Distribution Protocol (LDP), a full mesh of label-switched paths (LSPs) is automatically created between all LDP-enabled routers, and these paths are reflected in the inet.3 table. LDP LSPs are set up when LDP is enabled on interfaces connecting routers in an MPLS network.

#### NEW QUESTION: 28

Which two statements are correct about IS-IS? (Choose two.)

- A. A level 1 only router can never form an adjacency with a level 2 only router.
- B. For level 2 adjacencies, the area IDs can be different.
- C. For level 2 adjacencies, the area IDs must be the same.
- D. A level 1 only router can form an adjacency with a level 2 only router.

**Answer: A,B** ([LEAVE A REPLY](#))


A Level 1 router can become adjacent with the Level 1 and Level 1-2 (L1/L2) router. A Level 2 router can become adjacent with Level 2 or Level 1-2 (L1/L2) router. There is no adjacency between L1 only and L2 only router. HOWEVER: If two routers are in different areas, they can only form a Level 2 adjacency. As such, two routers in different areas can NOT form a Level 1 adjacency. If you want two routers to form a Level 1 adjacency, they have to be in the same area.

IS-IS (Intermediate System to Intermediate System) operates at two levels: Level 1 and Level 2.

Level 1 routers are only aware of their own area's topology, while Level 2 routers have knowledge of the topology across areas. A Level 1 router cannot form an adjacency with a Level 2 router unless the Level 2 router is also operating as a Level 1 router (Level 1-2 router). Level 2 routers can form adjacencies regardless of their area IDs because Level 2 operates at the domain level and is used to interconnect different IS-IS areas.

#### NEW QUESTION: 29

You are asked to configure the OSPF environment to prevent the DRother routers from participating in DR/BDR election. Referring to the exhibit, which command will accomplish this task?



- A. set protocols ospf area 0.0.0.0 interface ge-0/0/1 priority 255
- B. set protocols ospf area 0.0.0.0 interface ge-0/0/1 priority
- C. set protocols ospf area 0.0.0.0 interface ge-0/0/1 interface-type nbma
- D. set protocols ospf area 0.0.0.0 interface ge--0/0/1 interface-type p2p

**Answer: B** ([LEAVE A REPLY](#))

To prevent a router from participating in the OSPF designated router (DR) and backup designated router (BDR) election process, you set the OSPF priority to 0 on the interfaces of that router. This tells OSPF that the router on those interfaces is not to be considered for election as either DR or BDR, effectively making them DROther routers. The exhibit indicates that this is the desired configuration, and therefore the correct command to issue on the Juniper device would be to set the priority to 0 on the relevant interfaces.

#### NEW QUESTION: 30

You are evaluating BGP between two Juniper routers and the BGP session is stuck in the Idle state. What would cause this behavior?

- A. The BGP hold time is too short.
- B. The BGP group type is set to internal instead of external.
- C. The local AS number is missing.
- D. The peer IP address is incorrect.

**Answer: D** ([LEAVE A REPLY](#))

The Idle state indicates that the BGP process is waiting to initiate a connection but cannot establish a TCP session with the configured peer. If the peer IP address is incorrect, the router attempts to connect to a non-existent or wrong neighbor, preventing the TCP session on port 179 from being established and keeping the BGP session in the Idle state.

#### NEW QUESTION: 31

Which two statements are correct about default BGP route propagation? (Choose two.)

- A. UEBGP speakers advertise IBGP-learned routes to other EBGp peers.
- B. EBGp speakers advertise IBGP-learned routes to other EBGp peers by using a policy.
- C. IBGP speakers advertise IBGP-learned routes to other IBGP peers.
- D. IBGP speakers advertise EBGp-learned routes to other IBGP peers.


**Answer: A,D** ([LEAVE A REPLY](#))

**Valid JN0-364 Dumps** shared by Actual4test.com for Helping Passing JN0-364 Exam! Actual4test.com now offer the **newest JN0-364 exam dumps**, the Actual4test.com JN0-364 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com JN0-364 dumps with Test Engine here:

[https://www.actual4test.com/JN0-364\\_examcollection.html](https://www.actual4test.com/JN0-364_examcollection.html) (180 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

#### NEW QUESTION: 32

Referring to the exhibit, what is the minimum number of LSPs required to support all four networks?



- A. 1
- B. 2
- C. 8
- D. 4

**Answer: A** ([LEAVE A REPLY](#))

Given the RSVP Signaled Service Provider MPLS Core network shown in the exhibit, only a single Label Switched Path (LSP) is required to support all four networks because MPLS can use label stacking to multiplex multiple LSPs over a single physical path. Therefore, one LSP can be established through the MPLS core, and all networks can use this LSP with different label stacks.

#### NEW QUESTION: 33

Which new field is added to an IPv6 header as compared to IPv4?

- A. version
- B. checksum
- C. fragment offset
- D. flow label

**Answer: D** ([LEAVE A REPLY](#))

The IPv6 header includes a new field that is not found in the IPv4 header, called the flow label.

The flow label in IPv6 is used to identify packets that require special handling by routers for quality of service (QoS) or other reasons, allowing these packets to be handled efficiently as they move through the network.

#### NEW QUESTION: 34

What are two requirements for a unified in-service software upgrade? (Choose two.)


- A. The device must have dual Routing Engines.
- B. Bidirectional Forwarding Detection must be enabled on the device.
- C. Nonstop active routing must be enabled on the device.
- D. The device must be part of the chassis cluster.

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 35

Click the Exhibit button. You have configured IPv4 and IPv6 in your network and all OSPF neighbors are established. You apply the configuration shown in the exhibit.

Which statement is true in this scenario?



```
[edit]
user@R2# show protocols
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
    interface ge-0/0/1.0;
  }
}
ospf3 {
  realm ipv4-unicast {
    area 0.0.0.0 {
      interface ge-0/0/0.0;
      interface ge-0/0/1.0;
      interface lo0.0;
    }
  }
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
    interface lo0.0;
  }
}
```

- A. There will not be a route in R1 for network 172.16.2.0/24.
- B. There will only be an OSPFv3 entry in R1 for network 172.16.2.0/24.
- C. There will only be an OSPFv2 entry in R1 for network 172.16.2.0/24.
- D. There will be an OSPFv2 and OSPFv3 entry in R1 for network 172.16.2.0/24.

**Answer: D** ([LEAVE A REPLY](#))

The network is being advertised into OSPFv2 because the LAN-facing interface is enabled under OSPF. It is also being advertised into OSPFv3 for IPv4 because OSPFv3 is configured with the ipv4-unicast realm on the same interface. As a result, R1 learns 172.16.2.0/24 through both protocol instances and will have both protocol entries for that prefix.

#### NEW QUESTION: 36

Which two solutions are commonly used to send IPv6 traffic over IPv4 networks? (Choose two.)

- A. VXLAN
- B. NAT6
- C. 6to4
- D. 6over4

**Answer: C,D** ([LEAVE A REPLY](#))

6to4 is an IPv6 transition mechanism that encapsulates IPv6 packets inside IPv4 packets, allowing IPv6 traffic to traverse an IPv4 network infrastructure.

6over4 is another tunneling mechanism that transports IPv6 packets over an IPv4 network by using IPv4 as the underlying transport, enabling IPv6 communication across IPv4-only networks.

#### NEW QUESTION: 37

In an OSPF network, what is a purpose of a designated router?

- A. to flood routes to all other OSPF devices in the entire domain
- B. to forward traffic within the configured subnet
- C. to reduce OSPF traffic on the OSPF segment

**D.** to assign an OSPF router ID to all routers in the OSPF segment

**Answer: C** ([LEAVE A REPLY](#))

A designated router is elected on multiaccess OSPF networks to minimize the number of adjacencies and reduce the amount of OSPF control traffic. Instead of every router forming a full adjacency with every other router on the segment, routers form adjacencies with the designated router, which then manages the distribution of link-state updates on that segment.

#### **NEW QUESTION: 38**

What are two types of SIDs used in segment routing? (Choose two.)

**A.** adjacency

**B.** node

**C.** link

**D.** interface

**Answer: A,B** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 39**

You want to ensure that access switches on your network cannot transition to a root bridge state, while maintaining your carefully designed and implemented spanning-tree topology. Which feature would you deploy on interfaces connected to your access switches?

**A.** BPDU protection

**B.** root protection

**C.** loop protection

**D.** storm control

**Answer: (**[SHOW ANSWER](#)**)**

#### **NEW QUESTION: 40**

You are configuring LDP in a service provider network. After enabling LDP on core interfaces, you notice that labels are being advertised for every loopback IPv4 address that is in your IGP.

Which label distribution mode is being used in this scenario?

**A.** conservative retention

**B.** ordered control

**C.** downstream unsolicited

**D.** downstream on demand

**Answer: (**[SHOW ANSWER](#)**)**

Downstream unsolicited label distribution means that an LSR advertises labels to its neighbors without receiving a specific request. In this mode, routers automatically advertise labels for all known FECs, such as loopback prefixes learned through the IGP, which results in labels being distributed for every reachable prefix.

#### **NEW QUESTION: 41**

You are deploying link aggregation groups. By default, what are two considerations in this scenario? (Choose two.)

**A.** There should only be four member links per LAG.

**B.** All the ports must have the same speed.

- C. Member links are required to be contiguous ports.
- D. Member links can reside on different members within an MC-LAG.

**Answer: B,D** ([LEAVE A REPLY](#))

When deploying Link Aggregation Groups (LAGs), two key default considerations are:

All member ports must have the same speed. This ensures consistent performance and proper aggregation since mismatched speeds can cause issues in traffic distribution and link stability.

Member links can reside on different members within a Multi-Chassis LAG (MC-LAG). MC-LAG allows member links of the aggregation group to be spread across different physical chassis, providing redundancy and load sharing across devices.

#### NEW QUESTION: 42

What are two bridging concepts that are used to maintain an Ethernet switching table? (Choose two.)

- A. learning
- B. exporting
- C. aging
- D. timing

**Answer: A,C** ([LEAVE A REPLY](#))

In Ethernet switching, learning and aging are two fundamental concepts for maintaining a dynamic Ethernet switching table (also known as a MAC address table). Learning is the process by which switches listen for frames and learn the MAC addresses by associating them with the incoming port. Aging is the mechanism that ensures the switch updates its table by removing MAC addresses that have not been seen for a certain amount of time (aging time).

#### NEW QUESTION: 43

Click the Exhibit button. You have a network of ten routers that have all been configured with an identical SRGB. The exhibit shows the IS-IS configuration from a router called R10. The other nine routers do not yet have an IPv4 shortest-path SR-MPLS LSP to this router.

Which missing part of the configuration must you add on R10 to solve this problem?

```
user@R10> show configuration protocols isis
interface ge-0/0/1.0 {
  point-to-point;
}
interface ge-0/0/2.0 {
  point-to-point;
}
interface lo0.0;
source-packet-routing {
  srgb start-label 300000 index-range 10000;
}
level 1 disable;
level 2 wide-metrics-only;
reference-bandwidth 100g;
```

- A. R10 must tag its internal IPv4 BGP prefixes with a BGP prefix SID
- B. R10 must be configured with an explicit IPv4 node SID.
- C. R10 must be configured with an explicit binding SID.
- D. R10 must be configured with explicit IPv4 adjacency SID.

**Answer: B** ([LEAVE A REPLY](#))

For other routers to build an IPv4 shortest-path SR-MPLS LSP to R10, R10 must advertise a reachable SR node segment for its IPv4 loopback/prefix. Configuring an explicit IPv4 node SID ensures R10's loopback prefix is associated with a SID in the SRGB and is advertised in IS-IS, allowing the rest of the network to compute and label-switch traffic to R10 using SR.

#### **NEW QUESTION: 44**

Which two statements about the BGP path selection process are correct? (Choose two.)

- A. BGP selects the advertisement with the numerically lowest local preference.
- B. BGP selects the advertisement with the numerically lowest ME
- C. BGP selects the advertisement with the numerically highest ME
- D. BGP selects the advertisement with the numerically highest local preference.

**Answer: B,D** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 45**

You are building a new network using the IS-IS protocol. The network design calls for a 1GbE interface to have a cost of 100, and a 10GbE interface to have a cost of 10. Which statement is correct?

- A. The links between two 10GbE interfaces must have Level 1 disabled.
- B. All routers in the network must be configured for Level 2 only.
- C. Wide metrics must be enabled on all IS-IS routers.
- D. Jumbo frames must be enabled on all interfaces.

**Answer: C** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 46**

What is the correct order of BGP attributes for active route selection?

- A. next hop -> local preference -> AS path -> MED -> origin
- B. next hop -> AS path -> local preference -> origin -> MED
- C. next hop -> local preference -> AS path -> origin -> MED
- D. next hop -> origin -> local preference -> AS path -> MED

**Answer: C** ([LEAVE A REPLY](#))

BGP selects the active route based on an ordered list of attributes. The correct order is:

Weight (Cisco proprietary, not listed in the options)

Local Preference

Network (route) originated from the BGP router itself

Shortest AS Path

Lowest Origin Type

Lowest MED

eBGP over iBGP paths

Closest IGP Neighbor

Lowest Router ID

The next hop is checked for reachability but is not part of the BGP decision process for selecting the best path.

**Valid JN0-364 Dumps** shared by Actual4test.com for Helping Passing JN0-364 Exam! Actual4test.com now offer the **newest JN0-364 exam dumps**, the Actual4test.com JN0-364 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com JN0-364 dumps with Test Engine here:

[https://www.actual4test.com/JN0-364\\_examcollection.html](https://www.actual4test.com/JN0-364_examcollection.html) (180 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

#### **NEW QUESTION: 47**

An OSPF router does not have a router ID configured.

In this scenario, which statement is correct about the router ID?

- A.** The Junos OS will use the IP address assigned to the interface with the lowest MAC address.
- B.** A router ID will not be assigned until it is manually configured.
- C.** The Junos OS will use the IP address assigned to the loopback interface for the router ID.
- D.** The Junos OS will use the IP address assigned to the Interface with the highest priority.

**Answer: C** ([LEAVE A REPLY](#))

The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used. In OSPF, if a router ID is not manually configured, Junos OS will automatically select the router ID based on the highest IP address of any of the router's loopback interfaces. If no loopback interface exists, the highest IP address of any active interface will be used. Therefore, answer C is correct.

#### **NEW QUESTION: 48**

You are troubleshooting a Level 1 IS-IS router that has an adjacency with a Level 1/2 router.

Which two statements are correct in this scenario? (Choose two.)

- A.** The Level 1/2 router merges Level 1 and Level 2 into one complete topology.
- B.** The Level 1/2 router sees the Level 1 network and the Level 2 network as two separate topologies.
- C.** The Level 1 router will only learn the topology of the Level 1 network.
- D.** The Level 1 router will learn the full topology of the Level 2 network.

**Answer: B,C** ([LEAVE A REPLY](#))

A Level 1/2 router maintains separate link-state databases for Level 1 and Level 2, treating them as two distinct topologies rather than merging them into a single topology.

A Level 1 router only participates in the Level 1 area and therefore only learns the topology information contained within the Level 1 link-state database.

#### **NEW QUESTION: 49**

You are adding an IPv6 configuration to an Interface on a Junos device. In this scenario, which statement is correct?

- A.** The link local address must be manually configured within the fd00::/8 prefix range.
- B.** The link local address must be manually configured within the fe80::/10 prefix range.
- C.** The link local address is automatically created using the MAC address within the fe80::/10 prefix range.
- D.** The link local address is automatically created using the MAC address within the fd00::/8 prefix range.

**Answer: C** ([LEAVE A REPLY](#))

IPv6 link-local addresses are automatically generated for each interface and have a prefix of fe80::/10. The interface's MAC address is typically used as part of the process to create the Interface Identifier (IID) in the link-local address, following the EUI-64 format.

#### **NEW QUESTION: 50**

You are designing an MPLS network and want to ensure that traffic traverses an LSP between PE routers that follow an explicit path through the core.

Which protocol would accomplish this task?

- A.** BGP
- B.** LDP
- C.** IS-IS
- D.** RSVP

**Answer: D** ([LEAVE A REPLY](#))

RSVP-TE is used to signal MPLS traffic-engineered LSPs and supports explicit path configuration. This allows you to define the exact sequence of core hops the LSP must follow between PE routers, ensuring traffic traverses the desired path rather than simply following the IGP shortest path.

#### **NEW QUESTION: 51**

Which BGP attribute is used to detect routing loops?

- A.** AS path
- B.** MED
- C.** local preference
- D.** next hop

**Answer: A** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 52**

What information is determined by using the AS path attribute included in the BGP update message? (Choose two.)

- A.** the presence of a routing loop
- B.** the shortest AS path to reach a prefix
- C.** the origin of a route from IGP or EGP
- D.** the total number of next-hop devices to reach a prefix

**Answer:** ([SHOW ANSWER](#))

The AS path attribute lists the autonomous systems that a route advertisement has traversed.

Routers use this attribute to detect routing loops by checking if their own AS number appears in the path.

The AS path length is also used in the BGP best-path selection process, where routes with shorter AS paths are preferred because they represent fewer AS hops to reach the destination prefix.

#### **NEW QUESTION: 53**

Referring to the exhibit. Which two statements are correct about the information shown in the exhibit? (Choose two.)

```

user@switch> show spanning-tree bridge
STP bridge parameters
Context ID : 0
Enabled protocol : RSTP
Root ID : 8192.59.c5:8d:ae:db:41
Hello time : 10 seconds
Maximum age : 40 seconds
Forward delay : 30 seconds
Message age : 0
Number of topology changes : 6
Time since last topology change : 781 seconds
Topology change initiator : ge-0/0/14.0
Topology change last recvd. from : 2c:6b:f5:31:06:0b
Local parameters
Bridge ID : 8192.50.c5:8d:ae:db:41
Extended system ID : 0
Internal instance ID : 0

```


- A. The root bridge is reachable using the ge-0/0/14 interface.
- B. This switch is the root bridge for this spanning tree topology.
- C. This switch has a bridge priority of 8k.
- D. The root bridge's priority is 4k.

**Answer: A,C (LEAVE A REPLY)**

The exhibit shows the output of the command show spanning-tree bridge, which provides information about the spanning tree status of the switch. From the output, we can see that the switch has a bridge ID different from the root ID, which implies that this switch is not the root bridge. The "Topology change initiator" field shows ge-0/0/14, which indicates that the last topology change occurred on this interface, and this is also the interface used to reach the root bridge.

#### NEW QUESTION: 54

Referring to the exhibit, how do you verify the status of the tunnel from R1?



- A. Issue the ping 172.20.111.10 source 172.20.110.1 command.
- B. Issue the ping 172.20.111.10 source 198.51.100.1 command.
- C. Issue the ping 172.20.111.10 source 203.0.113.2 command.
- D. Issue the ping 172.20.111.10 command.

**Answer: (SHOW ANSWER)**

To verify the status of the tunnel from R1, you would issue a ping from the source address that is assigned to R1's end of the tunnel. In the exhibit, we can see that the tunnel interface (ge-0/0/0.0) has the IP address 198.51.100.1 on R1. Therefore, to test the tunnel's status, you should ping the IP address at the other end of the tunnel (which is likely the address on User B's interface or another interface on R2) from R1's tunnel source address.

#### NEW QUESTION: 55

Which feature allows Junos OS to perform recursive lookups for static route next hops?

- A. discard
- B. resolve
- C. reject
- D. next-table

**Answer: B** ([LEAVE A REPLY](#))

The resolve option enables recursive next-hop resolution for static routes in Junos OS. When configured, the router performs a lookup in the routing table to determine how to reach the specified next-hop address, allowing the next hop to be resolved through another route instead of requiring it to be directly connected.

#### NEW QUESTION: 56

Click the Exhibit button. All devices in the network are configured correctly and the path requirements are valid. Referring to the exhibit, which two statements are correct? (Choose two.)

```
[edit]
user@R1# show protocols mpls
label-switched-path R1-to-R6 {
    to 10.1.1.6;
    primary via-R2-R4;
    secondary any-path;
}
path via-R2-R4 {
    10.1.1.2 strict;
    10.1.1.4 strict;
}
path any-path;
interface ge-0/0/0.0;
interface ge-0/0/1.0;
```


- A. The primary LSP will be signaled, and its state will be up.
- B. The secondary LSP will not be signaled, and its state will be down.
- C. The secondary LSP will be signaled, and its state will be up.
- D. The primary LSP will not be signaled, and its state will be down.

**Answer: A,C** ([LEAVE A REPLY](#))

According to the exhibit, the primary LSP is configured with a strict path via R2 to R4. Since the configuration shows valid next-hops and there is no indication of any issues, the primary LSP will be signaled and its state will be up. The secondary LSP with any-path is also configured and will be signaled as a backup; therefore, its state will be up as well, ready to take over if the primary fails.

### NEW QUESTION: 57

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



- A. Prefixes in Level 1 will be redistributed to Level 2.
- B. Prefixes in Level 2 will be not redistributed to Level 1.
- C. Prefixes in Level 1 will not be redistributed to Level 2.
- D. Prefixes in Level 2 will be redistributed to Level 1.

**Answer: A,D** ([LEAVE A REPLY](#))

In IS-IS, Level 1 routes are usually contained within the same area and Level 2 routes are used to interconnect different areas. By default, routes from Level 1 are redistributed into Level 2, and vice versa, to ensure reachability between areas.

By default, IS-IS protocol leaks routing information from a Level 1 area to a Level 2 area.

However, to leak routing information from a Level 2 area to a Level 1 area, an export policy must be explicitly configured.

### NEW QUESTION: 58

You are asked to configure interfaces on Juniper devices to support dual VLAN tags.

In this scenario, which two interface statements would accomplish this task? (Choose two.)

- A. vlan-tagging
- B. flexible-vlan-tagging
- C. stacked-vlan-tagging
- D. gigether-options

**Answer: B,C** ([LEAVE A REPLY](#))

Flexible VLAN tagging enables an interface to support multiple VLAN tagging formats, including QinQ configurations that use two VLAN tags.

Stacked VLAN tagging enables the interface to process double-tagged Ethernet frames, allowing an outer service VLAN tag to be added to an inner customer VLAN tag for QinQ deployments.

### NEW QUESTION: 59

Which BGP attribute is used to detect routing loops?

- A. AS path
- B. MED
- C. local preference
- D. next hop

**Answer: A** ([LEAVE A REPLY](#))

The AS path is a BGP attribute that lists the autonomous systems (AS) that routing information has passed through to get to a destination. It is used to detect and prevent routing loops by checking if the router's own AS number is in the path.

#### NEW QUESTION: 60

A BGP router receives two routes to the same prefix. One route has a higher local preference, while the other has a shorter AS path.

In this scenario, which route would be selected?

- A. The route with the lower origin code.
- B. The route with the higher local preference.
- C. The route with the shorter AS path.
- D. The route with the lowest MED value.

**Answer: B** ([LEAVE A REPLY](#))

Local preference is evaluated before the AS path length in the BGP best-path selection process.

A route with a higher local preference is therefore preferred even if another route has a shorter AS path.

#### NEW QUESTION: 61

Which LSA type does an OSPF ABR use to advertise external routes generated by an NSSAASBR into the backbone?

- A. Type 5
- B. Type 7
- C. Type 3
- D. Type 1

**Answer: B** ([LEAVE A REPLY](#))

In OSPF, an Area Border Router (ABR) that connects to a Not-So-Stubby Area (NSSA) will convert the Type 7 LSAs generated by an NSSA Autonomous System Boundary Router (ASBR) into Type 5 LSAs to advertise them into the OSPF backbone area (Area 0). Type 7 LSAs are specific to NSSAs and are used to advertise external routes within the NSSA.

**Valid JN0-364 Dumps** shared by Actual4test.com for Helping Passing JN0-364 Exam! Actual4test.com now offer the **newest JN0-364 exam dumps**, the Actual4test.com JN0-364 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com JN0-364 dumps with Test Engine here:

[https://www.actual4test.com/JN0-364\\_examcollection.html](https://www.actual4test.com/JN0-364_examcollection.html) (180 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

#### NEW QUESTION: 62

Click the Exhibit button.



```
user@R1> show configuration protocols mpls
label-switched-path to-r3 {
    to 192.168.100.3;
}
interface ge-0/0/0.0;
user@R1> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
}
user@R1> show route 192.168.100.3
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.100.3/32    *[OSPF/10] 00:05:39, metric 2
                   > to 172.16.1.2 via ge-0/0/0.0
user@R1> show mpls lsp detail
Ingress LSP: 1 sessions
192.168.100.3
  From: 192.168.100.1, State: Dn, ActiveRoute: 0, LSPname: to-r3
  ActivePath: (none)
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Will be enqueued for recomputation in 27 second(s).
    17 Sep 14 20:29:00.840 CSPF: could not determine self
Total 1 displayed, Up 0, Down 1
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
user@R1> show configuration interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 172.16.1.1/24;
        }
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
```

```

        address 10.0.1.11/24;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.100.1/32;
        }
    }
}

```

You have configured an MPLS LSP to 192.168.100.3. However, the LSP is in the down state. Referring to the exhibit, which two actions would solve this problem? (Choose two.)

- A.** Issue the set protocols ospf traffic-engineering command and commit.
- B.** Issue the set interfaces lo0 family mpls command on router R1 and commit.
- C.** Issue the set protocols mpls label-switched-path to-r3 no-cspf command and commit.
- D.** Issue the set routing-options rib inet.3 static route 192.168.100.1 command and commit.

**Answer: A,C** ([LEAVE A REPLY](#))

The LSP is using CSPF and the output indicates CSPF cannot build a path because the traffic-engineering database is not available. Enabling OSPF traffic-engineering allows the router to populate the TE database used by CSPF. Alternatively, disabling CSPF forces the LSP to be computed using the normal IGP shortest path, which removes the dependency on the TE database and allows the LSP to come up.

#### NEW QUESTION: 63

You are deploying an MC-LAG connection at the edge of your network to a new customer. This solution must ensure that traffic uses all available links when forwarding traffic between the two networks.

In this scenario, which two statements are correct? (Choose two.)

- A.** Transit traffic sent between the two PE devices will use the bundle's member interfaces.
- B.** You must use active/standby mode when defining the MC-LAG bundle.
- C.** Transit traffic sent between the two PE devices will use the ICL interface.
- D.** You must use active/active mode when defining the MC-LAG bundle.

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 64

You are asked to configure filter-based forwarding on a Junos device. Which two statements are correct in this scenario? (Choose two.)

- A.** You must create a routing policy.
- B.** You must create a route target.
- C.** You must create and apply a match filter.
- D.** You must create a routing instance.

**Answer: A,C** ([LEAVE A REPLY](#))

Filter-based forwarding in Junos requires the creation of a routing policy to specify the match conditions and actions and the application of a match filter to classify traffic to use the policy.

#### NEW QUESTION: 65

You are asked to configure an LSP which uses the OSPF link state database for path computations. Which two statements are correct in this scenario? (Choose two.)

- A. You must use the no-cspf parameter in the label-switched-path configuration.
- B. Traffic engineering extensions are enabled by default in OSPF.
- C. Traffic engineering extensions are not enabled by default in OSPF.
- D. You must use the policing parameter in the label-switched-path configuration.

**Answer: A,C** ([LEAVE A REPLY](#))

In Junos OS, traffic engineering extensions for OSPF, which are required for an LSP to use the OSPF link state database for path computations, are not enabled by default. They must be explicitly enabled in the OSPF configuration. Therefore, answer C is correct. Answer B is incorrect and contradicts C. The 'no-cspf' command would disable CSPF (Constrained Shortest Path First), which is used for path computations in MPLS traffic engineering, so it should not be used if you want the LSP to utilize OSPF's link state database for path computations. Therefore, answer A is incorrect. The 'policing' parameter is not relevant to enabling traffic engineering extensions in OSPF, so answer D is incorrect.

#### NEW QUESTION: 66

Which statement about RSVP-signaled LSPs is correct?

- A. CSPF is used to calculate the path for a traffic-engineered LSP.
- B. CSPF is not required for LSPs using admin-groups.
- C. The paths used by LSPs are always calculated using the SRGB.
- D. The paths used by LSPs are always calculated using the TED.

**Answer: A** ([LEAVE A REPLY](#))

RSVP traffic-engineered LSPs use constraint-based shortest path first (CSPF) to compute a path that satisfies TE constraints such as bandwidth, metrics, and administrative groups. The resulting path is then signaled hop-by-hop using RSVP.

#### NEW QUESTION: 67

Referring to the exhibit, what does the configuration achieve on the ge-0/0/3 interface?

```
[edit interfaces ge-0/0/3 unit 0]
user@R1# set family inet6 address 2001:db8:0:1::/64 eui-64
```

- A. This interface will send a request to a DHCP server, requesting an EUI-64 address in this prefix.
- B. This interface will assign itself an IPV6 address within this prefix, based on a 64-bit hash of the loopback IPv6 address.
- C. This interface will assign itself an IPv6 address within this prefix, based on its MAC address.
- D. This interface will send a request to an EUI server, requesting a 64-bit interface address within this 64-bit prefix.

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 68

Which two protocols are capable of distributing labels for segment routing? (Choose two.)

- A. RSVP

- B. LDP
- C. IS-IS
- D. OSPF

**Answer: C,D** ([LEAVE A REPLY](#))

Segment routing labels can be distributed by link-state routing protocols such as IS-IS and OSPF, which have been extended to support segment routing capabilities.

#### NEW QUESTION: 69

Which two statements are correct about the community BGP attribute on a Junos device?

(Choose two.)

- A. The community attribute is a mandatory BGP attribute.
- B. If the community attribute is present, it is ignored and deleted in the BGP updates.
- C. If the community attribute is present, it should be passed unchanged in the BGP updates.
- D. The community attribute is an optional BGP attribute.

**Answer: C,D** ([LEAVE A REPLY](#))

The community attribute in BGP is an optional transitive attribute. It is not mandatory for BGP operations, but when it is present, it should be passed unchanged to other BGP peers unless explicitly modified by a routing policy. The community attribute is used to group destinations in a certain way, often to apply routing policies to those destinations collectively.

#### NEW QUESTION: 70

Referring to the exhibit, what will happen to untagged frames?

```
[edit]
user@switch# show interfaces ge-0/0/1
native-vlan-id 20;
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members [ 10 20 ];
        }
    }
}
```



- A. The untagged frames are dropped.
- B. The untagged frames are load balanced between VLAN 10 and VLAN 20.
- C. The untagged frames are associated with VLAN 10.
- D. The untagged frames are associated with VLAN 20.

**Answer: (SHOW ANSWER)**

In the exhibit, the interface is configured with the native-vlan-id 20 parameter. This means that any untagged frames arriving on this interface will be associated with VLAN 20. The native VLAN serves as the default VLAN for untagged frames on a trunk port.

#### NEW QUESTION: 71

Which BGP message type is used to re-advertise routes that have already been sent to a peer and acknowledged using TCP?

- A. update
- B. keepalive
- C. notification
- D. refresh

**Answer: A** ([LEAVE A REPLY](#))

In BGP (Border Gateway Protocol), the Update message is used to advertise new routes to a peer or to withdraw previously advertised routes. If there's a change in the routing information or if a new route needs to be advertised, BGP uses an Update message. This message is crucial for maintaining the most current and accurate routing information among BGP peers.

The other message types serve different purposes:

Keepalive: This message type is used to maintain the connection between BGP peers and ensure that the link is still active.

Keepalive messages are periodically sent between peers when there is no other BGP activity to maintain the session.

Notification: This message is used to indicate errors or other significant events to a BGP peer. If a BGP error condition is detected, a Notification message is sent and the BGP session is typically closed.

Refresh: This is not a standard BGP message type. However, in the context of BGP, Route Refresh is a capability (not a message type) that allows a BGP router to request a full or partial re-advertisement of routes from a peer. This helps in re-synchronizing routing tables without tearing down the BGP session.

The BGP Update message is used to advertise routable destinations and includes route advertisements, route withdrawals, and path attribute modifications. Update messages are used to re-advertise routes that have already been acknowledged.

#### NEW QUESTION: 72

You are a network architect designing a brand new network. You want to deploy RSVP LSPs in this network. You are currently in the process of choosing whether to run OSPF or IS-IS as your interior gateway protocol.

In this scenario, which two statements are correct about IGP traffic engineering extensions in an RSVP network? (Choose two.)

- A. You must explicitly configure IS-IS to carry traffic engineering extensions.
- B. You must explicitly configure OSPF to carry traffic engineering extensions.
- C. In OSPF, traffic engineering extensions are enabled by default.
- D. In IS-IS, traffic engineering extensions are enabled by default.

**Answer: A,B** ([LEAVE A REPLY](#))

When using IS-IS in an RSVP traffic engineering environment, traffic engineering extensions must be explicitly enabled so that IS-IS advertises TE information such as link bandwidth and administrative groups to build the traffic engineering database.

When using OSPF, traffic engineering extensions must also be explicitly configured. Enabling OSPF traffic engineering allows the protocol to advertise TE attributes required by RSVP-TE and CSPF for calculating constrained paths.

#### NEW QUESTION: 73


When considering BGP route propagation, which two statements are correct? (Choose two.)

- A. IBGP speakers send routes learned from IBGP neighbors to IBGP neighbors.
- B. IBGP speakers send routes learned from EBGp neighbors to IBGP neighbors.
- C. EBGp speakers send routes learned from EBGp neighbors to IBGP neighbors.
- D. EBGp speakers send routes learned from IBGP neighbors to IBGP neighbors.

**Answer: B,D** ([LEAVE A REPLY](#))

#### NEW QUESTION: 74

Click the Exhibit button.



You have configured an MPLS LSP that begins on R1 and terminates on R5 using the Junos default settings. Referring to the exhibit, which router will perform only label swap operations?

- A. R1
- B. R3
- C. R4
- D. R5

**Answer:** ([SHOW ANSWER](#))

With default MPLS behavior, the ingress router pushes a label, transit routers perform label swaps, the penultimate router pops the label (PHP), and the egress router receives the packet without the transport label. In the path R1-R5, R1 pushes the label, R4 performs the pop as the penultimate hop, and R5 is the egress. Therefore, the only router performing purely label swap operations is the intermediate transit router R3.

#### NEW QUESTION: 75

An established RSVP-signaled LSP suffers a link failure. What is one action taken by the local RSVP-enabled router?

- A. A ResvTear message is sent downstream.
- B. A ResvTear message is sent upstream.
- C. A ResvErr message is sent downstream.
- D. A ResvErr message is sent upstream.

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 76

In OSPF, which three fields must match between neighbors before forming an adjacency?

(Choose three.)

- A. hello interval
- B. designated router
- C. network mask
- D. dead interval
- E. router priority

**Answer:** ([SHOW ANSWER](#))

For OSPF neighbors to form an adjacency, certain parameters exchanged in Hello packets must match. The hello interval must be identical on both routers so they send Hello packets at the same rate. The network mask must match on broadcast and nonbroadcast networks to ensure both routers are operating on the same subnet. The dead interval must also match so both routers


agree on how long to wait before declaring a neighbor down.

**Valid JN0-364 Dumps** shared by Actual4test.com for Helping Passing JN0-364 Exam! Actual4test.com now offer the **newest JN0-364 exam dumps**, the Actual4test.com JN0-364 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com JN0-364 dumps with Test Engine here:

[https://www.actual4test.com/JN0-364\\_examcollection.html](https://www.actual4test.com/JN0-364_examcollection.html) (180 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

#### NEW QUESTION: 77

Referring to the exhibit, from which device(s) does R3 learn about Route X?



- A. both R2 and R1
- B. R1 only
- C. directly from the router in AS 64502
- D. R2 only

**Answer: D** ([LEAVE A REPLY](#))

R2 can not forward IBGP learned routes to another IBGP neighbor correct R3 will learn about Route X from R2 only. R1 is in a separate AS and Route X is advertised to R2 through eBGP. R2 will then advertise the route to R3 using iBGP. R3 cannot learn routes directly from another AS without having a direct eBGP session with that AS.

#### NEW QUESTION: 78

Which OSPF packet type is used to initiate and maintain neighbor relationships?

- A. Database Description
- B. hello
- C. Link-State Update
- D. Link-State Acknowledgment

**Answer: (**[SHOW ANSWER](#)**)**

Hello packets are used by OSPF routers to discover neighbors and establish adjacency relationships. They are also sent periodically to maintain the neighbor relationship and verify that neighboring routers are still reachable.

#### NEW QUESTION: 79

Which two statements are correct about the way that BGP propagates routes by default?

(Choose two.)

- A. A route learned by EBGP will be re-advertised to IBGP peers.

- B. A route learned by IBGP will not be re-advertised to IBGP peers.
- C. A route learned by EBGP will not be re-advertised to IBGP peers.
- D. A route learned by IBGP will be re-advertised to IBGP peers.

**Answer: A,B** ([LEAVE A REPLY](#))

By default, BGP will propagate routes learned from EBGP peers to other IBGP peers. However, due to the BGP split-horizon rule, routes learned from one IBGP peer are not re-advertised to another IBGP peer. This behavior is designed to prevent routing information loops within the same AS.

#### NEW QUESTION: 80

You are using RSVP in your network to signal MPLS LSPs. You are asked to ensure that the egress MPLS router receives the full label stack, containing both the hop-by-hop and VPN label.

Which configuration will satisfy this requirement?


- A. Apply set protocols rsvp interface no-aggregate to the egress interface.
- B. Use the set protocols mpls expand-loose-hop command.
- C. Use the set protocols rsvp keep-multiplier command.
- D. Use the set protocols mpls explicit-null command.

**Answer: D** ([LEAVE A REPLY](#))

#### NEW QUESTION: 81

Referring to the exhibit. The routers shown in the exhibit are configured for segment routing.

In this scenario, what is the adjacency SID that Router B advertises to Router A?



- A. 12
- B. 10.1.1.1
- C. 10.1.1.2
- D. 11

**Answer: A** ([LEAVE A REPLY](#))

The adjacency SID in segment routing is advertised by one router to inform another about the label that should be used to forward traffic to it over a particular link. In this scenario, Router B would advertise an adjacency SID to Router A, which would be used by Router A to forward traffic to Router B. Based on the exhibit, the adjacency SID that Router B advertises to Router A would be 12.

#### NEW QUESTION: 82

BGP update messages contain which two elements? (Choose two.)

- A. path advertisements
- B. path attributes
- C. path assignments

D. path announcement

Answer: ([SHOW ANSWER](#))

#### NEW QUESTION: 83

Which BGP attribute is optional, transitive, and is passed unchanged to other BGP peers if not recognized?

A. MED

B. Community

C. AS Path


D. Origin

Answer: B ([LEAVE A REPLY](#))

The Community attribute is an optional transitive BGP attribute. Optional transitive attributes are propagated to other BGP peers even if a router does not recognize them. The Community attribute is commonly used to tag routes so that routing policies can apply specific handling to those routes as they propagate through the network.

#### NEW QUESTION: 84

Click the Exhibit button.



```
user@switch1> show spanning-tree interface
Spanning tree interface parameters for instance 0
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| ge-0/0/6.0  | 128:519 | 128:519               | 32768.0019e2552481      | 20000        | FWD   | DESG |
| ge-0/0/7.0  | 64:520  | 64:520                | 32768.0019e2552481      | 20000        | FWD   | DESG |
| ge-0/0/8.0  | 32:521  | 32:521                | 32768.0019e2552481      | 20000        | FWD   | DESG |
| ge-0/0/9.0  | 32:522  | 32:522                | 32768.0019e2552481      | 20000        | FWD   | DESG |
| ge-0/0/11.0 | 32:524  | 32:524                | 32768.0019e2552481      | 20000        | FWD   | DESG |
| ge-0/0/12.0 | 64:525  | 64:525                | 32768.0019e2552481      | 20000        | FWD   | DESG |
| ge-0/0/13.0 | 64:526  | 64:526                | 32768.0019e2552481      | 20000        | FWD   | DESG |

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

A. The switch1 device is using VSTP.

B. The switch1 device is the root bridge.

C. The ge-0/0/8, ge-0/0/9, and ge-0/0/11 interfaces are using the default interface priority.

D. The bridge priority for switch 1 is 32k.

Answer: ([SHOW ANSWER](#))

All listed ports are in the designated role and forwarding state, and the designated bridge ID shown matches the local bridge ID, which indicates the device is acting as the root bridge for this spanning-tree instance.

The bridge ID begins with 32768, which corresponds to the default bridge priority of 32,768 (32k).

#### NEW QUESTION: 85

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

```

user@R2> show ospf interface extensive
Interface State Area DR ID BDR ID Vbrs
ge-0/0/3.0 DR 0.0.0.1 192.168.1.2/24 192.168.1.1 1 Type: LAN, Address: 172.26.1.2, Mask:
255.255.255.252, MTU: 1500, Cost: 1
DR addr: 172.26.1.2, BDR addr: 172.26.1.1, Priority: 128, Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None Topology default (ID 0) -> Cost: 0
ge-0/0/1.0 BDR 0.0.0.0 192.168.1.3 192.168.1.2 1
Type: LAN, Address: 172.26.2.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1
DR addr: 172.26.2.2, BDR addr: 172.26.2.1, Priority: 128, Adj count: 1 Hello: 10,
Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 0

```

- A. The OSPF Interfaces are configured as point-to-point.
- B. The ge-0/0/1.0 Interface is configured as passive.
- C. The R2 device is an ABR.
- D. Junos OS default OSPF hello timers and dead intervals are used on all interfaces.

**Answer: C,D (LEAVE A REPLY)**

You can see area 0.0.0.1 set for one interface and area 0.0.0.0 for the other interface. This means the router communicates in the backbone and area 1, thus it is an ABR.

Default hello timer and dead intervals are configured

hello-interval--Specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. By default, the routing device sends hello packets every 10 seconds.

The range is from 1 through 255 seconds.

dead-interval--Specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. By default, the routing device waits

40 seconds (four times the hello interval). The range is 1 through 65,535 seconds.

C). The R2 device is shown as having interfaces in two different OSPF areas (Area 0 and another area not specified), which makes it an Area Border Router (ABR).

D). The OSPF hello and dead intervals are set to their default values of 10 and 40 seconds, respectively.

**Valid JN0-364 Dumps** shared by Actual4test.com for Helping Passing JN0-364 Exam! Actual4test.com now offer the **newest JN0-364 exam dumps**, the Actual4test.com JN0-364 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com JN0-364 dumps with Test Engine here:

[https://www.actual4test.com/JN0-364\\_examcollection.html](https://www.actual4test.com/JN0-364_examcollection.html) (180 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)