Juniper

Exam JN0-346

Enterprise Routing and Switching, Specialist (JNCIS-ENT)

Verson: Demo

[Total Questions: 10]

Question No : 1

Which mechanism is used to share routes between routing tables?

A. filter-based forwarding
B. forwarding instances
C. RIB groups
D. routing instances

Answer: C

Explanation:
A RIB group is a way to have a routing protocol, in most cases, place information in multiple route tables.

Question No : 2

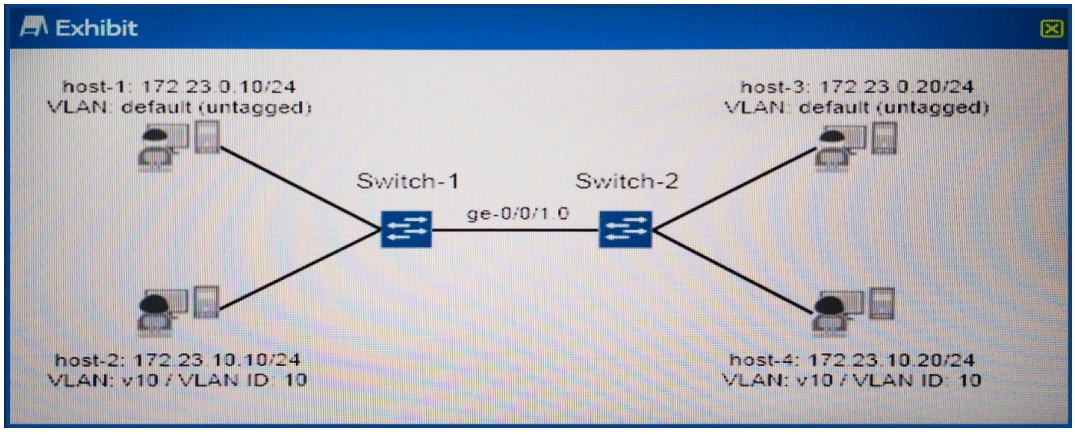

Click the Exhibit button.

Referring to the exhibit, you are asked to ensure that host-1 can communicate with host-3 while also allowing host-2 to communicate with host-4.

What should you do to enable this behavior?

A. Configure thenative-vlan-id defaultstatement under the ge-0/0/1 port settings on Switch-1.

B. Use theallkeyword when defining the member VLANs for the ge-0/0/1 interface on Switch-1.
C. Configure thenative-vlan-id defaultstatement under the ge-0/0/1 port settings on both Switch-1 and Switch-2.
D. Use theallkeyword when defining the member VLANs for the ge-0/0/1 interface on both Switch-1 and Switch-2.

Answer: C

Explanation:
Configuring Mixed Tagging Support for Untagged Packets

For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs, for all MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces configured for 802.1Q flexible VLAN tagging, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, you can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untaggedpackets, include the native-vlan-id statement and the flexible-vlan-tagging statement at the [edit interfaces interface-name] hierarchy level:
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
native-vlan-id number;

References: https://www.juniper.net/documentation/en_US/junos13.3/topics/usage-guidelines/interfaces-enabling-vlan-tagging.html

Question No : 3

What is the default route preference for BGP?

A. 167
B. 170
C. 150
D. 179

Answer: B

Explanation:

BGP has the default preference of 170.

References:https://www.juniper.net/documentation/en_US/junos14.2/topics/reference/general/routing-protocols-default-route-preference-values.html

Question No : 4

Exhibit

```
Nov  3 15:39:56.388955 SPF post spf cleanup
finished
Nov  3 15:39:56.388959  Cleanup elapsed time
0.000064s
Nov  3 15:39:56.388965    Total elapsed time
0.003092s
Nov  3 15:39:56.388967 Finished full SPF refresh
for topology default
Nov  3 15:39:56.388969 task_job_delete: delete
background job Route recalc timer for task OSPF
Nov  3 15:39:56.388971 background dispatch
completed job Route recalc timer for task OSPF
Nov  3 15:40:02.900115 task_process_events: recv
ready for OSPF I/O./var/run/ppmd_control
Nov  3 15:40:02.900227 task_process_events: recv
ready for OSPF I/O./var/run/ppmd_control
Nov  3 15:40:02.900242 task_timer_uset: timer
OSPF I/O./var/run/ppmd_control_PPM Hold
<Touched> set to offset 2:00 at 15:42:02
Nov  3 15:40:02.900244 OSPF packet ignored: area
mismatch (0.0.0.1) from 192.168.150.254 on intf
ge-0/0/1.0 area 1.0.0.0
Nov  3 15:40:02.900246 OSPF rcvd Hello
192.168.150.254 -> 224.0.0.5 (ge-0/0/1.0 IFL 72
area 1.0.0.0)
Nov  3 15:40:02.900344   Version 2, length 44,
ID 10.254.254.254, area 0.0.0.1
Nov  3 15:40:02.900346   checksum 0x8a7a,
authtype 0
Nov  3 15:40:02.900348   mask 255.255.255.0,
hello_ivl 10, opts 0x12, prio 128
Nov  3 15:40:02.900350   dead_ivl 40, DR
192.168.150.254, BDR 0.0.0.0
Nov  3 15:40:02.900374 task_timer_uset: timer
OSPF_internal timer <Touched> set to offset 5 at
15:40:07
Nov  3 15:40:04.225141 task_process_events: recv
ready for OSPF I/O./var/run/ppmd_control
Nov  3 15:40:04.225293 task_process_events: recv
ready for OSPF I/O./var/run/ppmd_control
Nov  3 15:40:04.225350 task_timer_uset: timer
OSPF I/O./var/run/ppmd_control_PPM Hold
<Touched> set to offset 2:00 at 15:42:04
Nov  3 15:40:04.225352 OSPF periodic xmit from
192.168.150.253 to 224.0.0.5 (IFL 72 area
1.0.0.0)
Nov  3 15:40:06.025582 task_process_events: recv
ready for OSPF I/O./var/run/ppmd_control
Nov  3 15:40:06.025685 task_process_events: recv
ready for OSPF I/O./var/run/ppmd_control
Nov  3 15:40:06.025713 task_timer_uset: timer
OSPF I/O./var/run/ppmd_control_PPM Hold
<Touched> set to offset 2:00 at 15:42:06
Nov  3 15:40:06.025715 OSPF periodic xmit from
172.16.128.253 to 224.0.0.5 (IFL 71 area
1.0.0.0)
```

Click the Exhibit button.

Based on the traceoptions output shown in the exhibit, what is the problem with the adjacency?

A. connectivity
B. authentication mismatch
C. area mismatch
D. MTU mismatch

Answer: C

Explanation:

From the exhibit we see:

OSPF packet ignored: area mismatch

Question No : 5

What is reviewed first in the BGP route selection process?

A. the peerwith the lowest IP address
B. the route with an origin of incomplete
C. the path with no MED value
D. the next-hop resolution

Answer: D

Explanation:

Understanding BGP Path Selection

The algorithm for determining the active route is as follows:

Etc.

References:https://www.juniper.net/documentation/en_US/junos12.3/topics/reference/general/routing-ptotocols-address-representation.html

Question No : 6

Which two statements are true about OSPF not-so-stubby areas? (Choose two.)

A. The ASBR originates Type 7 LSAafor redistributed external routes.
B. Type 5 LSAs are translated by the ASBR into Type 7 LSAs.
C. The ASBR originates Type 5 LSAs for redistributed external routes.
D. Type 7 LSAs are translated by the ABR into Type 5 LSAs.

Answer: A,D

Explanation:

Redistribution intoan NSSA area creates a special type of link-state advertisement (LSA) known as type 7, which can only exist in an NSSA area. An NSSA autonomous system boundary router (ASBR) generates this LSA and an NSSA area border router (ABR) translates it into a type5 LSA, which gets propagated into the OSPF domain.

References:http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/6208-nssa.html

Question No : 7

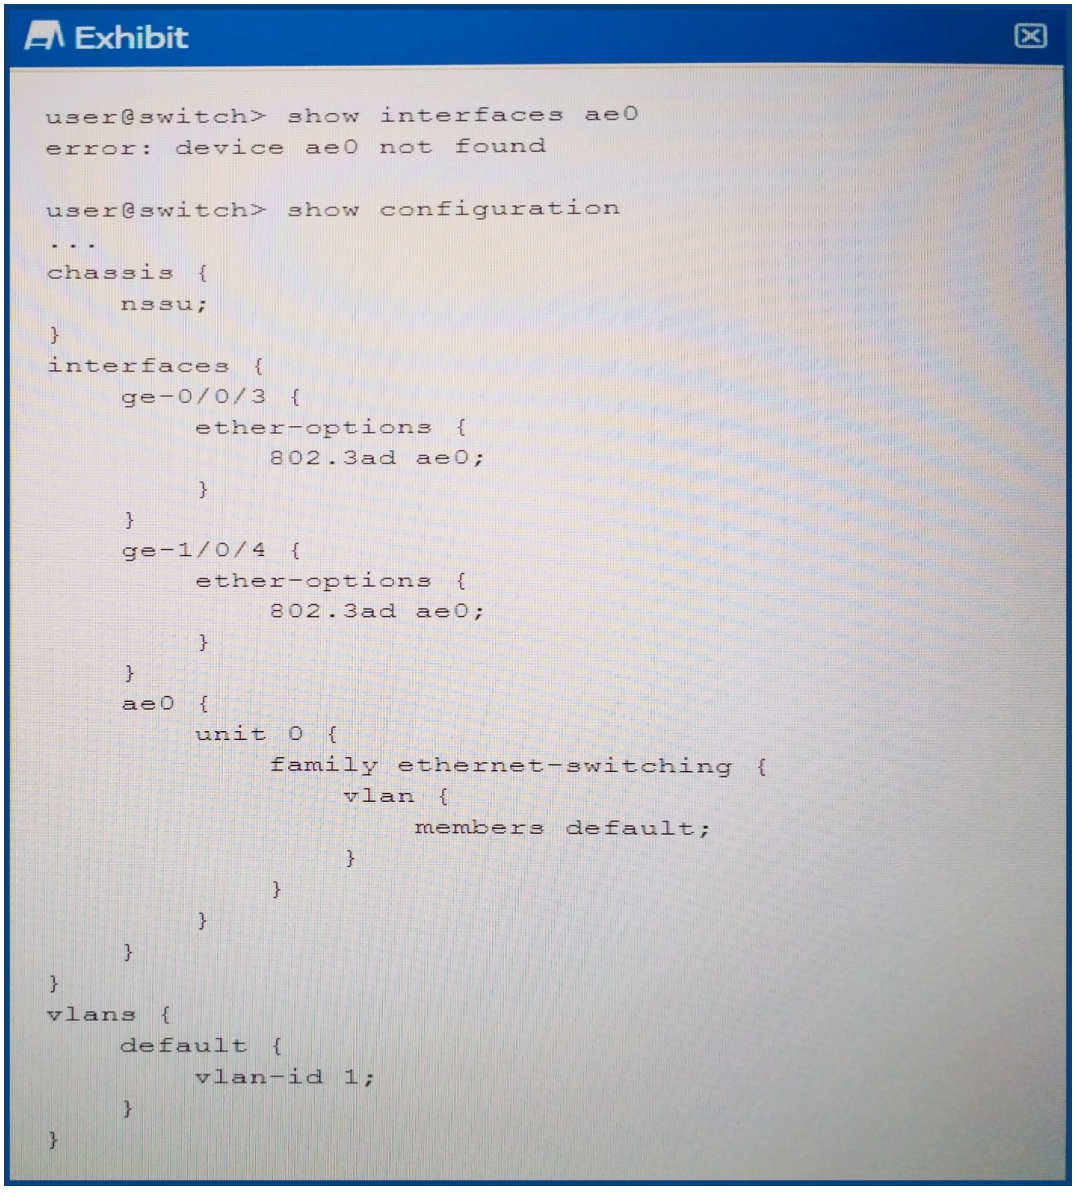

Click the Exhibit button.

Referring to the exhibit, what is the problem?

A. LAG requires more than two member links.
B. LACP is required for LAG to work.
C. Aggregated interfaces must be defined under the chassis stanza.
D. The LAG member interfaces are configured across different line cards.

Answer: C

Explanation:

Use the link aggregation feature to aggregate one or more links to form a virtual link or link aggregation group (LAG).

To configure aggregated Ethernet interfaces, using the CLI:

Question No : 8

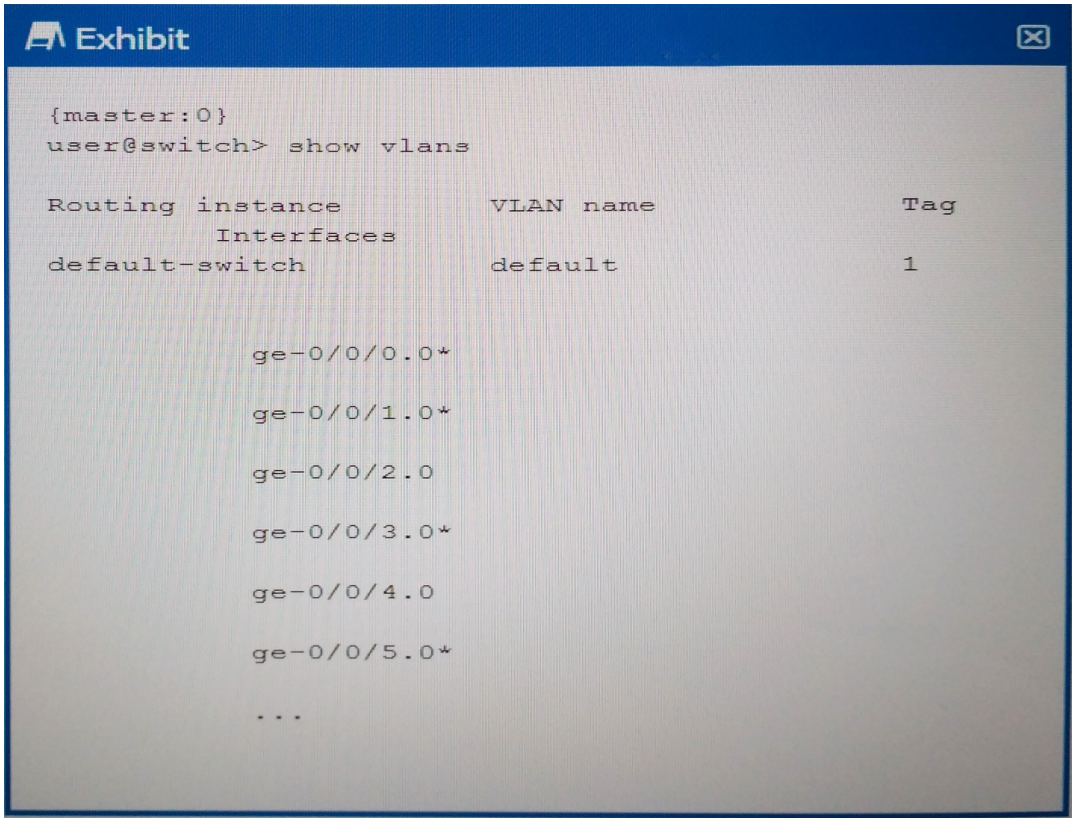

Click the Exhibit button.

Referring to the exhibit, what does the asterisk (*) following the ge-0/0/5.0 interface indicate?

A. It indicates the interface is a trunk port.
B. It indicates the interface is not active.
C. It indicates the interface is an access port.
D. It indicates the interface is active.

Answer: D

Explanation:

An asterisk (*) beside the interface indicates that the interface is UP.

References:
http://www.juniper.net/documentation/en_US/junos14.1/topics/reference/command-summary/show-vlans-bridging-qfx-series.html

Question No : 9

You configured a GRE tunnel that traverses a path using default MTU settings. You want to ensure that packets are not dropped or fragmented.

In this scenario, what is the maximum packet size that would traverse the GRE tunnel?

A. 1476
B. 1500
C. 1400
D. 1524

Answer: A

Explanation:
The default Ethernet MTU is 1500. There is a 24 byte GRE overhead, so there remain 1476 bytes for the data packet.

References:https://kb.juniper.net/InfoCenter/index?page=content&id=KB7848&actp=search

Question No : 10

Exhibit

```
[edit]
user@Router-1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.10.33/24;
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 10.1.0.254/24;
    }
    family iso {
      address 49.0003.0192.0168.0113.00;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.1.11/32;
    }
    family iso {
      address 49.0002.0192.0168.0111.00;
    }
  }
}

[edit]
user@Router-1# show protocols
isis {
  overload;
  level 2 disable;
  interface all;
}

[edit]
user@Router-2# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.10.34/24;
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 10.1.0.1/16;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.1.12/32;
    }
    family iso {
      address 49.0001.0192.0168.0112.00;
    }
  }
}

[edit]
user@Router-2# show protocols
isis {
  interface all;
}
```

Click the Exhibit button.

Referring to the exhibit, Router-1 and Router-2 are failing to form an IS-IS adjacency.

What should you do to solve the problem?

A. Change the IP subnet masks to match on the ge-0/0/2 interfaces of both routers.
B. Change the ISO areas on the Io0 interfaces to match on both routers.
C. Remove the ISO address fromge-0/0/2 on Router-1
D. Remove theoverloadedstatement from Router-1.

Answer: C

Explanation:

There are two interfaces with ISO addresses on Router-1, and they have different area IDs, 002 and 003. Only one interface on Router-1 need to have an ISO address.