



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

Click the Exhibit button.

You are troubleshooting a BGP connection.

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

```
user@router> show log messages | match notification
Dec 22 19:22:29 router rpd[7394]: bgp_process_open:4185: NOTIFICATION sent to
192.168.1.4 (Internal AS 65000): code 2 (Open Message Error) subcode 2 (bad peer AS
number), Reason: peer 192.168.1.4 (Internal AS 65000) claims 65100, 65000 configured
Dec 22 19:22:33 router rpd[7394]: bgp_pp_rcv:4798: NOTIFICATION sent to 192.168.1.4+
56774 (proto): code 2 (Open Message Error) subcode 2 (bad peer AS number), Reason: no
group for 192.168.1.4+56774 (proto) from AS 65100 found (peer as mismatch)in master
(ge-0/0/1.0), dropping him
Dec 22 19:23:29 router kernel: tcp_auth_ok: Packet from 192.168.1.5:64047 missing MD5
digest
Dec 22 19:23:30 router kernel: tcp_auth_ok: Packet from 192.168.1.6:56201 missing MD5
digest
---(more)---
```

- A. Packet fragmentation is preventing the session from establishing.
- B. The 192.168.1.5 peer has a misconfigured MD5 key.
- C. The ge-0/0/1 interface is disabled.
- D. The 192.168.1.4 peer has a misconfigured autonomous system number.

Suggested Answer: AC

Community vote distribution

BD (100%)

 **CiscoTest** Highly Voted 2 years, 7 months ago

It is B and D.

upvoted 10 times

 **kim1276** Most Recent 7 months, 2 weeks ago

B and D

check the configuration carefully

upvoted 1 times

 **song820902** 1 year, 5 months ago

IT IS B and D

upvoted 1 times

 **pawel1025** 1 year, 7 months ago

Has anyone taken the JN0-649 exam recently? Are those question valid? Thanks

upvoted 3 times

 **tsukasa123** 1 year, 6 months ago

About half are still valid.

There are 30 new questions, but if you understand the content here, you can pass.

upvoted 2 times

 **somanquestions** 1 year, 8 months ago

Selected Answer: BD

it's clearly in the logs


upvoted 1 times

 **aconejop202** 1 year, 10 months ago

Selected Answer: BD

B and D

upvoted 2 times

 **jmbecerra** 1 year, 10 months ago

Selected Answer: BD

B and D

upvoted 2 times

🗨️ 👤 **alliki** 2 years, 1 month ago

Any body completed JN0-649 by these 65 questions?

upvoted 2 times

🗨️ 👤 **machine327** 1 year, 11 months ago

Taking it next week; you take it yet?

upvoted 1 times

🗨️ 👤 **alliki** 1 year, 11 months ago

Have you taken exam yet?

upvoted 2 times

🗨️ 👤 **hbstyleboy** 2 years, 3 months ago

Don't try to smoke us, Answer is definitely B & D

upvoted 3 times

🗨️ 👤 **JoeSun** 2 years, 6 months ago

Selected Answer: BD

B and D. The log message show it.

upvoted 3 times

🗨️ 👤 **dragossky** 2 years, 6 months ago

B&D obviously ...

upvoted 3 times

🗨️ 👤 **mohdema** 2 years, 6 months ago

Selected Answer: BD

You just have to read the output to figure out it's B&D, can an admin correct this please ?

upvoted 2 times


Click the Exhibit button.

Referring to the exhibit, anycast RP is implemented to ensure multicast service availability. The source is currently sending multicast traffic using group 239.1.1.1 and R3 is receiving PIM register messages, but R2 does not have active source information.

In this scenario, what are two methods to receive the active source information on R2? (Choose two.)

```
user@R1> show pim statistics | match "(PIM Message type) | (V2 Register)"
```

PIM Message type	Received	Sent	Rx errors
V2 Register	0	857	0
V2 Register Stop	0	0	0



```
user@R3> show pim statistics | match "(PIM Message type) | (V2 Register)"
```

PIM Message type	Received	Sent	Rx errors
V2 Register	857	0	0
V2 Register Stop	0	0	0

```
user@R5> show pim join
```

```
...
Group: 239.1.1.1
Source: 10.222.3.2
Flags: sparse,spt
Upstream interface: ge-0/0/12.0
```

- A. Configure an RP set in PIM on R1, allowing R1 to forward PIM register messages to R2 and R3 in the set.
- B. Configure an MSDP protocol between R2 and R3.
- C. Configure an RP set in PIM on R2 and R3, allowing the RPs to forward PIM register messages to the other RPs in the set.
- D. Configure an MSDP protocol between R1 and R2.

Suggested Answer: AC

Community vote distribution

BC (100%)

proptz 5 months, 3 weeks ago

Selected Answer: AB

There is no doubt the B is correct answer. I am not sure why others selected C as it doesn't make sense. Why would you create an RP set on the RP itself to magically learn anything new from R1? R1 needs to be aware who the RPs are. Hence A is correct.

upvoted 1 times

kim1276 7 months, 2 weeks ago

b and c

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/topic-map/mcast-pim-anycast-rp.html>

upvoted 1 times

song820902 1 year, 5 months ago

b and c

upvoted 1 times

somanyquestions 1 year, 8 months ago

Selected Answer: BC

Anycast RP uses Multicast Source Discovery Protocol (MSDP) to discover and maintain a consistent view of the active sources.

include the anycast-pim statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the rp-set statement.

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/topic-map/mcast-pim-anycast-rp.html>

upvoted 2 times

sanalainen 2 years, 6 months ago

Selected Answer: BC

MSDP or anycast-PIM with rp-set should be configured on all anycast RP routers.

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/topic-map/mcast-pim-anycast-rp.html>

upvoted 4 times

  **penguin02007** 2 years, 6 months ago

Should be B and C.

upvoted 2 times

  **mohdema** 2 years, 6 months ago

Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.

The remaining statements are explained separately. See CLI Explorer.

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/ref/statement/rp-set-edit-protocols-pim.html>

upvoted 1 times

You are asked to establish interface level authentication for users connecting to your network. You must ensure that only corporate devices, identified by MAC addresses, are allowed to connect and authenticate. Authentication must be handled by a centralized server to increase scalability.

Which authentication method would satisfy this requirement?

- A. MAC RADIUS
- B. captive portal
- C. 802.1X with single-secure supplicant mode
- D. 802.1X with multiple supplicant mode

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **kim1276** 7 months, 2 weeks ago

MAC RADIUS

upvoted 2 times

🗳️ 👤 **root9211** 2 years, 1 month ago

Selected Answer: A

A is correct.

upvoted 1 times

🗳️ 👤 **root9211** 2 years, 1 month ago

Selected Answer: A

AAAAAAAAAAAAA

upvoted 1 times

🗳️ 👤 **mohdema** 2 years, 6 months ago

Selected Answer: A

<https://www.juniper.net/documentation/us/en/software/junos/user-access/topics/topic-map/mac-radius-authentication-switching-devices.html>

You can configure MAC RADIUS authentication on an interface that also allows 802.1X authentication, or you can configure either authentication method alone.

If both MAC RADIUS and 802.1X authentication are enabled on the interface, the switch first sends the host three EAPoL requests to the host. If there is no response from the host, the switch sends the host's MAC address to the RADIUS server to check whether it is a permitted MAC address. If the MAC address is configured as permitted on the RADIUS server, the RADIUS server sends a message to the switch that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.

upvoted 2 times

🗳️ 👤 **penguin02007** 2 years, 6 months ago

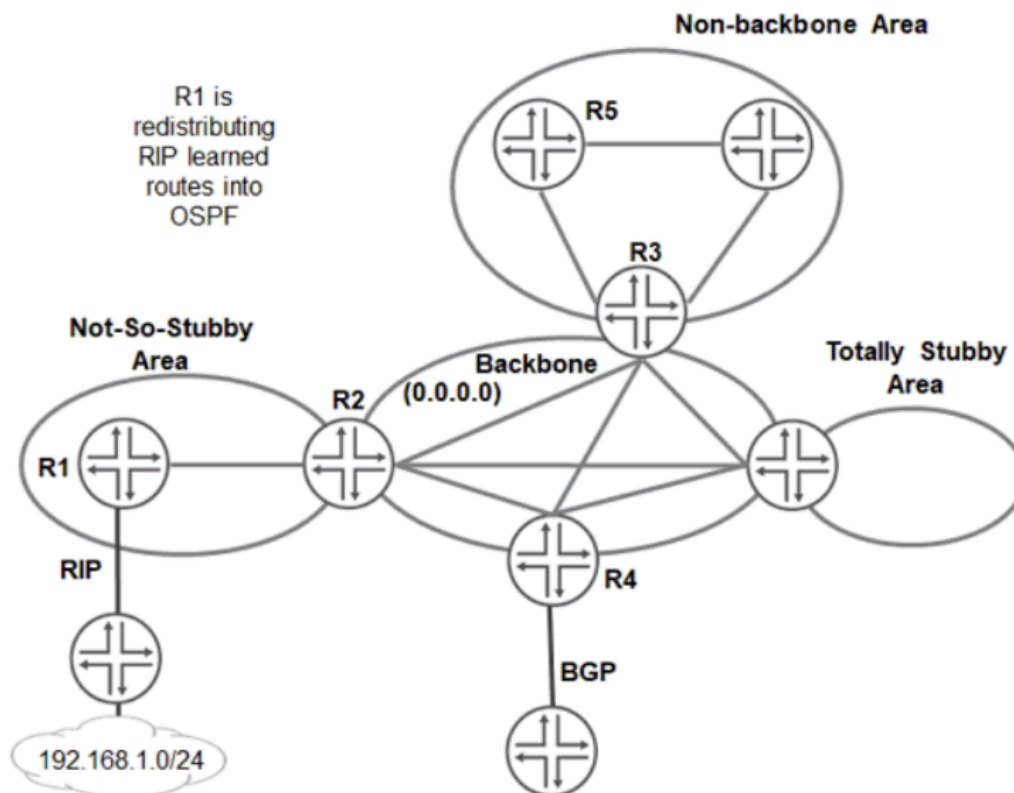
A is correct.

<https://www.juniper.net/documentation/us/en/software/junos/user-access/topics/topic-map/mac-radius-authentication-switching-devices.html>

upvoted 1 times

Click the Exhibit button.

Referring to the exhibit, which LSA type is used to advertise 192.168.1.0/24 to R5?



- A. Type 5
- B. Type 4
- C. Type 3
- D. Type 7

Suggested Answer: B

Community vote distribution

A (100%)

mohdema Highly Voted 2 years, 6 months ago

Selected Answer: A

Area-1 has no external connections. However, Area-1 has static route (172.16.31.0/24) that are not internal OSPF route. You can limit the external route advertisements to the area and advertise the static routes by designating the area an NSSA. In an NSSA, the ASBR (vMX1) generates NSSA external (Type 7) LSAs and floods them into the NSSA, where they are contained.

Type-7 LSAs allow an NSSA to support the presence of ASBR and their corresponding external routing information. The ABR (vMX2) converts Type-7 LSAs into Type-5 External LSAs and leaks them to the other areas, but external routes from other areas are not advertised within the NSSA.

An admin should check this and change it

<https://www.packetswitch.co.uk/configuring-junos-ospf-stub-and-nssa-areas/>

<https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/ref/statement/nssa-edit-protocols-ospf.html>
upvoted 7 times

carlitox Highly Voted 2 years, 7 months ago

Should it be type 5 instead?

upvoted 7 times

🗄️ 👤 **kim1276** Most Recent 7 months, 2 weeks ago

Type 5

upvoted 2 times

🗄️ 👤 **fotavid775** 11 months, 2 weeks ago

Correct Answer: B. Type 4 : Passed JN0-649 with flying colors! Next goal: Pursuing JNCIS-SP to deepen my networking knowledge.

upvoted 1 times

🗄️ 👤 **dragosky** 2 years, 6 months ago

OSPF AS SCOPE link state database

Type ID Adv Rtr Seq Age Opt Cksum Len

Extern 10.10.10.10 10.10.10.2 0x80000001 1591 0x22 0xcb7c 36

Extern 10.10.10.11 10.10.10.2 0x80000001 1591 0x22 0xb790 36

Extern 192.168.1.0 10.10.10.2 0x80000001 1591 0x22 0xd234 36

upvoted 1 times

🗄️ 👤 **dragosky** 2 years, 6 months ago

Extern 192.168.1.0 10.10.10.2 0x80000001 1224 0x22 0xd234 36

mask 255.255.255.0

Topology default (ID 0)

Type: 2, Metric: 2, Fwd addr: 10.10.10.1, Tag: 0.0.0.0

Aging timer 00:39:35

Installed 00:20:22 ago, expires in 00:39:36, sent 00:20:22 ago

Last changed 00:20:22 ago, Change count: 1

upvoted 1 times

🗄️ 👤 **dragosky** 2 years, 6 months ago

ASBR Summary LSA (Type 4) is generated by ABR (Area Border Router) to inform its areas about how to reach the ASBR (Autonomous System Border Router). ASBR Summary LSA (Type 4) includes ASBR's Router ID.

upvoted 2 times

🗄️ 👤 **carlitox** 2 years, 7 months ago

Each of the NSSA external LSAs advertised by the ASBR router is translated by the ABR into an AS external LSA. These Type 5 LSAs are then advertised to the rest of the network using a

domain-flooding scope. These "new" Type 5 LSAs, in addition to the appropriate ASBR summary LSA, are seen on the router in area 2

upvoted 3 times

You enable the Multiple VLAN Registration Protocol (MVRP) to automate the creation and management of virtual LANs. Which statement is correct in this scenario?

- A. The forbidden mode does not register or declare VLANs.
- B. When enabled, MVRP affects all interfaces.
- C. Timers dictate when link state changes are propagated.
- D. MVRP works with RSTP and VSTP.

Suggested Answer: B

Community vote distribution

A (100%)

 **mohdema** Highly Voted 2 years, 6 months ago

Selected Answer: A

A. The forbidden mode does not register or declare VLANs.

You can change the registration mode of a specific interface to forbidden. An interface in forbidden registration mode does not participate in MVRP even if MVRP is enabled on the switch.

B. When enabled, MVRP affects all interfaces.

NOTE: Only trunk interfaces can be enabled for MVRP.

C. Timers dictate when link state changes are propagated.

The timers in MVRP define the amount of time an interface waits to join or leave MVRP or to send or process the MVRP information for the router or switch after receiving an MVRP PDU

D. MVRP works with RSTP and VSTP

MVRP works with Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP), but not with VLAN Spanning Tree Protocol (VSTP).

<https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/topic-map/mvrp.html>

upvoted 5 times

 **Garrysh** Most Recent 5 months, 1 week ago


Selected Answer: C

Answer A is incorrect because, you can specifically configure an interface to forbidden registration mode but by default is normal mode.

When you enable the MVRP, affects only trunk interfaces. So answer B is incorrect.

Answer C is correct "The timers define when MVRP PDUs can be sent and when MVRP information can be updated on a switch."

upvoted 1 times

 **kim1276** 7 months, 2 weeks ago

C. Timers dictate when link state changes are propagated.

upvoted 1 times

 **App431** 1 year, 3 months ago

In the scenario of enabling the Multiple VLAN Registration Protocol (MVRP) to automate VLAN creation and management, the correct statement is:

B. When enabled, MVRP affects all interfaces.

Explanation:

MVRP, when enabled, operates on all interfaces of the network device. It dynamically manages VLANs by exchanging control messages between neighboring switches to automate the creation and management of VLANs.

Option A is true but it's not the only correct statement. The forbidden mode in MVRP prevents a port from participating in the VLAN registration process.

Option C is incorrect. Timers in MVRP are used for various purposes related to the protocol operation, but they are not directly related to link state

changes propagation.

Option D is incorrect. MVRP (Multiple VLAN Registration Protocol) is independent of spanning tree protocols like RSTP (Rapid Spanning Tree Protocol) and VSTP (VLAN Spanning Tree Protocol). It focuses on VLAN management and does not directly interact with spanning tree protocols.

upvoted 1 times

🗲️ 👤 **DinoVercotti** 1 year, 6 months ago

Selected Answer: A

A is right

upvoted 1 times

🗲️ 👤 **somanyquestions** 1 year, 8 months ago

Selected Answer: A

An interface in forbidden registration mode does not participate in MVRP even if MVRP is enabled on the switch.

upvoted 1 times

🗲️ 👤 **Rudy6969** 2 years, 5 months ago

A all the way

upvoted 1 times

🗲️ 👤 **dragosky** 2 years, 6 months ago

Selected Answer: A

A for sure

upvoted 1 times

🗲️ 👤 **penguin02007** 2 years, 6 months ago

Should be A.

<https://www.juniper.net/documentation/us/en/software/junos/mvrp/multicast-l2/topics/concept/mvrp-mx-series-understanding.html>

upvoted 1 times

🗲️ 👤 **CiscoTest** 2 years, 7 months ago

Wrong answer. Correct answer is A

MVRP is disabled by default on the switches and, when enabled, affects only trunk interfaces. Once you enable MVRP, all VLAN interfaces on the switch belong to MVRP (the default normal registration mode) and those interfaces accept PDU messages and send their own PDU messages.

forbidden—The interface does not register or declare VLANs (except statically configured VLANs).

upvoted 3 times

Which address range is used for source-specific multicast?

- A. 239.0.0.0/8
- B. 233.0.0.0/8
- C. 232.0.0.0/8
- D. 224.2.0.0/16

Suggested Answer: C

Community vote distribution

C (100%)

🗲️ 👤 **kim1276** 7 months, 2 weeks ago

232.0.0.0/8 for sure
upvoted 2 times

🗲️ 👤 **fecaca7957** 11 months, 2 weeks ago

Correct answer is C. 232.0.0.0/8. Passed JN0-649 Looking forward to earning JNCIS-Cloud to specialize further in Juniper's cloud solutions.
upvoted 2 times

🗲️ 👤 **mohdema** 2 years, 6 months ago

Selected Answer: C

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

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/concept/multicast-pim-ssm.html>

upvoted 3 times

Which three configuration parameters must match on all switches within the same MSTP region? (Choose three.)

- A. VLAN to instance mapping
- B. revision level
- C. configuration name
- D. bridge priority
- E. region name

Suggested Answer: BCE

Community vote distribution



mohdema Highly Voted 2 years, 6 months ago

Selected Answer: ABE

I think it's ABE

When enabling MSTP, you define one or more MSTP regions. An MSTP region defines a logical domain where multiple spanning-tree instances (MSTIs) can be administered independently of MSTIs in other regions, setting the boundary for bridge protocol data units (BPDUs) sent by one MSTI. An MSTP region is a group of switches that is defined by three parameters:

Region name—User-defined alphanumeric name for the region.

Revision level—User-defined value that identifies the region.

Mapping table—Numerical digest of VLAN-to-instance mappings.

<https://www.juniper.net/documentation/us/en/software/junos/stp-l2/topics/topic-map/spanning-tree-configuring-mstp.html>

upvoted 14 times

MiguelAMZ 1 year, 1 month ago

Correct

upvoted 1 times

JoeSun 2 years, 6 months ago

Relate the juniper doc which you shared, in the MSTP CLI

[edit ... protocols mstp]

user@switch# set configuration-name configuration-name

Therefore if we said the CLI should be "configuration-name". However, Maybe the name of this parameters call "region name"?

upvoted 5 times

harrypogi Highly Voted 1 year, 9 months ago

Selected Answer: ABC

Region name and configuration name are basically the same.

But the question is looking for specifically for config parameter.

In config, the word used is configuration-name.

upvoted 5 times

Alex_NYC Most Recent 3 months, 2 weeks ago

Selected Answer: ABE

When enabling MSTP, you define one or more MSTP regions. An MSTP region defines a logical domain where multiple spanning-tree instances (MSTIs) can be administered independently of MSTIs in other regions, setting the boundary for bridge protocol data units (BPDUs) sent by one MSTI. An MSTP region is a group of switches that is defined by three parameters:

Region name—User-defined alphanumeric name for the region.

Revision level—User-defined value that identifies the region.

Mapping table—Numerical digest of VLAN-to-instance mappings.

upvoted 1 times

🗳️ 👤 **Alex_NYC** 3 months, 3 weeks ago

Selected Answer: ABC

VLAN to instance mapping ✓

All switches in the MSTP region must have the same VLANs mapped to the same MST instances.

This ensures that MSTP properly calculates the spanning tree for each instance.

Revision level ✓

The revision number acts as a version control mechanism.

If it does not match, the switches may consider themselves in different MST regions.

Configuration name (also called region name) ✓

The region name (or configuration name) must be identical on all switches in the MSTP region.

If the region name does not match, the switches will treat each other as separate MSTP regions.

upvoted 1 times

🗳️ 👤 **Rafrafi** 4 months ago

Selected Answer: ABE

An MSTP region is a group of switches that is defined by three parameters:

- Region name—User-defined alphanumeric name for the region.
- Revision level—User-defined value that identifies the region.
- Mapping table—Numerical digest of VLAN-to-instance mappings.

<https://www.juniper.net/documentation/us/en/software/junos/stp-l2/topics/topic-map/spanning-tree-configuring-mstp.html#:~:text=An%20MSTP%20region%20is%20a%20group%20of%20switches%20that%20is%20defined%20by%20three%20parameters%3A>

Region name—User-defined alphanumeric name for the region.

Revision level—User-defined value that identifies the region.

Mapping table—Numerical digest of VLAN-to-instance mappings.

upvoted 1 times

🗳️ 👤 **bn1996** 6 months, 1 week ago

Selected Answer: ABE

When enabling MSTP, you define one or more MSTP regions. An MSTP region defines a logical domain where multiple spanning-tree instances (MSTIs) can be administered independently of MSTIs in other regions, setting the boundary for bridge protocol data units (BPDUs) sent by one MSTI.

An MSTP region is a group of switches that is defined by three parameters:

Region name—User-defined alphanumeric name for the region.

Revision level—User-defined value that identifies the region.

Mapping table—Numerical digest of VLAN-to-instance mappings.

upvoted 1 times

🗳️ 👤 **Warrior0000** 6 months, 2 weeks ago

Selected Answer: ABC

According to AJES training (Revision 21A), page 104

upvoted 1 times

🗳️ 👤 **kim1276** 7 months, 2 weeks ago

a,b,c sure

upvoted 1 times

🗳️ 👤 **OkoJun** 8 months, 3 weeks ago



The correct answer is ABC by using the elimination method

see :

user

https://supportportal.juniper.net/s/article/EX-series-Switch-and-MSTP-Multiple-Spanning-Tree-Protocol?language=en_US

upvoted 1 times

  **OkoJun** 8 months, 3 weeks ago

The region name should be the same for several switches in the same region but not for all MSTP switches.

upvoted 1 times

  **davidrsr** 10 months, 1 week ago

Selected Answer: ABE

As per AJEX training.

upvoted 1 times

  **e359166** 1 year, 5 months ago

Selected Answer: ABC

https://supportportal.juniper.net/s/article/EX-series-Switch-and-MSTP-Multiple-Spanning-Tree-Protocol?language=en_US

upvoted 3 times

  **wassupkay** 1 year, 5 months ago

A B and C - Juniper doesn't have an option for region name.

VLAN TO instance mapping

REvision

and

COnfiguration name

upvoted 1 times

  **somanyquestions** 1 year, 8 months ago

Selected Answer: ABC

tricky one, since configuration-name and region name refer to the same thing, since they mention config parameter, it choose ABC

upvoted 4 times

  **Iulian1995** 2 years, 2 months ago

ABE is the correct answer

upvoted 1 times

  **abhi999210** 2 years, 3 months ago

ABE is the correct answer

upvoted 1 times

  **dragossky** 2 years, 6 months ago

Selected Answer: ABC

ABC valid answers

upvoted 4 times

  **dragossky** 2 years, 6 months ago

edit protocols mstp

set Configuration-name region1

set revision-level 1

set msti 1 vlan [1-4094]

set protocols mstp interface all

upvoted 3 times

Which two statements are correct about the deployment of EVPN-VXLAN on QFX Series devices? (Choose two.)

- A. Type 1 route advertisements always have the single-active flag set to 1.
- B. Junos OS supports underlay replication for BUM traffic forwarding.
- C. Junos OS supports ingress replication for BUM traffic forwarding.
- D. Type 1 route advertisements always have the single-active flag set to 0.

Suggested Answer: BC

Community vote distribution

CD (82%)

BC (18%)

 **CiscoTest** Highly Voted 2 years, 7 months ago


C and D is correct

BUM traffic handling is given to the overlay network by configuring VXLAN ingress-node-replication instead of multicast overlay

In an EVPN-VXLAN environment, EVPN multihoming active-active mode is used instead of MC-LAG for redundant connectivity between hosts and leaf devices.

In active-active mode, each of the multihomed PE device advertises a mandatory autodiscovery route per Ethernet segment as in the active-standby state. However, in the active-active state, the autodiscovery route per Ethernet segment is modified such that the active-standby bit carried in the MPLS extended community is cleared to indicate that the active-active mode is in operation.

upvoted 6 times

 **Jafet0990** Highly Voted 1 year, 6 months ago

Selected Answer: CD

CD

If the Single-Active flag is set to 1, that means only one link associated with the Ethernet segment can be used for forwarding. If the Single-Active flag is set to 0, that means that all links associated with the Ethernet segment can be used for forwarding data (we call this active/active forwarding).

Juniper QFX devices only support active/active forwarding and so this flag is always set to 0 for these devices.

upvoted 5 times

 **kim1276** Most Recent 7 months, 2 weeks ago

answer will be C & D


upvoted 2 times

 **somanquestions** 1 year, 8 months ago

Selected Answer: CD

Type 1 route advertisements are used for multi-homed hosts for route-withdrawels and Aliasing.

upvoted 2 times

 **JoeSun** 2 years, 6 months ago

Selected Answer: CD

EVPN is a overlay protocol most likely not change the underlay tag.

single-active represented by binary 1 or active-active represented by binary 0.

upvoted 3 times

 **mohdema** 2 years, 6 months ago

BUM Traffic Forwarding

Junos devices that use MPLS encapsulation for EVPNs can only use ingress replication at this time.

Ingress replication means, to flood traffic to remote PE routers, the traffic has to be replicated, once for each remote PE router.

The EVPN label for this BUM traffic is learned per PE router from the route type 3, inclusive multicast Ethernet tag route.

This table shows the format of the inclusive multicast Ethernet tag route.

All-Active Redundancy (4)

This diagram shows the format of the type 1 route, A-D route per ES. The split horizon label is advertised as part of an extended community attached to the type 1 route. The split horizon label is also called the ESI label. The extended community also indicates what type of redundancy mode is used for this given ESI: single-active represented by binary 1 or active-active represented by binary 0.

upvoted 1 times

🗨️ 👤 **mohdema** 2 years, 6 months ago

Selected Answer: BC

See my comments below

upvoted 2 times

🗨️ 👤 **mohdema** 2 years, 6 months ago

Answer is actually CD I mistyped

upvoted 1 times

🗨️ 👤 **mohdema** 2 years, 6 months ago

Single-Active Mode: When a device or a network is multi-homed to two or more PEs and when only a single PE in such redundancy group can forward traffic to/from the multi-homed device or network for a given VLAN, then such multi-homing or redundancy is referred to as "Single-Active".

All-Active: When a device is multi-homed to two or more PEs and when all PEs in such redundancy group can forward traffic to/from the multi-homed device for a given VLAN, then such multi-homing or redundancy is referred to as "All-Active".

If "Single-Active" flag in "ESI Label Extended

Community" is set, the ES is operating in Single-Active redundancy mode. Otherwise, it is operating in All-Active redundancy mode. The Ethernet A-D per EVI route can be used for Aliasing and Backup-Path, aliasing is used for all-active mode, backup-path is for singleactive mode.

upvoted 1 times

🗨️ 👤 **mohdema** 2 years, 6 months ago

The bit in the active-standby flag field in the ESI label extended community is used for signaling the active-standby mode (bit set).

upvoted 1 times

🗨️ 👤 **mohdema** 2 years, 6 months ago

In active-standby mode, the designated forwarder (DF) advertises the autodiscovery route per Ethernet segment with an ESI MPLS label extended community that has the standby bit set to 1. The autodiscovery route is advertised per ESI, and the ESI label is set to 0 when active-standby mode is in operation.

upvoted 1 times

🗨️ 👤 **penguin02007** 2 years, 6 months ago

C and D is correct. QFX only support active-active forwarding so flag is always set to zero.

upvoted 2 times




Your enterprise network is running BGP VPNs to support multitenancy. Some of the devices with which you peer BGP do not support the VPN NLRI. You must ensure that you do not send BGP VPN routes to the remote peer. Which two configuration steps will satisfy this requirement? (Choose two.)

- A. Configure an import policy on the remote peer to reject the routes when they are received.
- B. Configure an export policy on the local BGP peer to reject the VPN routes being sent to the remote peer.
- C. Configure a route reflector for the VPN NLRI.
- D. Configure the apply-vpn-export feature on the local BGP peer.

Suggested Answer: BD

Community vote distribution

BD (100%)

  **mohdema**  2 years, 6 months ago

Selected Answer: BD

pply both the VRF export and BGP group or neighbor export policies (VRF first, then BGP) before routes from the vrf or l2vpn routing tables are advertised to other PE routers.

upvoted 8 times

  **mohdema** 2 years, 6 months ago

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/statement/vpn-apply-export-edit-protocols-bgp-vp.html>

upvoted 1 times

  **kim1276**  7 months, 2 weeks ago

Answer: B. Configure an export policy on the local BGP peer to reject the VPN routes being sent to the remote peer.
and

D. Configure the apply-vpn-export feature on the local BGP peer.

upvoted 1 times

  **harrypogi** 1 year, 9 months ago

Selected Answer: BD

bd is correct

upvoted 1 times

You want to create an OSPF area that only contains intra-area route information in the form of Type 1 and Type 2 LSAs. In this scenario, which area is needed to accomplish this task?

- A. totally non-to-stubby area
- B. totally stubby area
- C. stub area
- D. non-to-stubby area

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **kim1276** 7 months, 2 weeks ago

B. totally stubby area

A totally stubby area is a more restrictive type of OSPF stub area that blocks not only Type 4 and 5 LSAs but also Type 3 LSAs from other OSPF areas.
upvoted 2 times

🗳️ 👤 **harrypogi** 1 year, 9 months ago

Selected Answer: B

b is correct

upvoted 2 times

🗳️ 👤 **mohdema** 2 years, 6 months ago

A totally stubby area (TSA) is a stub area in which summary link-state advertisement (type 3 LSAs) are not sent. A default summary LSA, with a prefix of 0.0. 0.0/0 is originated into the stub area by an ABR, so that devices in the area can forward all traffic for which a specific route is not known, via ABR.

upvoted 3 times

🗳️ 👤 **carlitox** 2 years, 7 months ago

B should be correct

upvoted 2 times

You are implementing the route summarization feature of OSPF.
Which two results do you achieve in this scenario? (Choose two.)

- A. It helps in migrating to future multi-area OSPF network designs.
- B. It reduced the routing table size, enabling devices to store and process less information.
- C. It reduces the impact of topology changes on a device.
- D. It provides optimal routing in the network.

Suggested Answer: BC

Community vote distribution

BC (100%)

🗲️ 👤 **mohdema** Highly Voted 👍 2 years, 6 months ago

Selected Answer: BC

OSPF inter-area route summarization reduces the routing information exchanged between areas and the size of routing tables, and improves routing performance. OSPF inter-area route summarization enables an ABR to summarize contiguous networks into a single network and advertise the network to other areas.

upvoted 5 times

🗲️ 👤 **kim1276** Most Recent ⌚ 7 months, 2 weeks ago

B & C is the correct answer

benefits of route summarization in ospf

Saves memory: routing tables will be smaller which reduces memory requirements.

Saves bandwidth: there are less routes to advertise so we save some bandwidth.

Saves CPU cycles: less packets to process and smaller routing tables to work on.

Stability: Prevents routing table instability due to flapping networks.

so on

upvoted 2 times

🗲️ 👤 **yh511** 1 year, 7 months ago

Selected Answer: BC

BC correct

upvoted 1 times

🗲️ 👤 **harrypogi** 1 year, 9 months ago

Selected Answer: BC

bc is correct


upvoted 1 times

Your organization has recently acquired another company. You must carry all of the company's existing VLANs across the corporate backbone to the existing branch locations without changing addressing and with minimal configuration. Which technology will accomplish this task?

- A. Q-in-Q all-in-one bundling
- B. PVLAN isolated VLAN
- C. MVRP registration normal
- D. EVPN-VXLAN anycast gateway

Suggested Answer: A

Community vote distribution



🗲️ 👤 **carlitox** Highly Voted 👍 2 years, 7 months ago

Selected Answer: A

A looks correct

upvoted 5 times

🗲️ 👤 **kim1276** Most Recent 🕒 7 months, 2 weeks ago

Q-in-Q all-in-one bundling looks perfect

upvoted 2 times

🗲️ 👤 **rfg8366** 1 year ago

Selected Answer: C

<https://www.juniper.net/documentation/us/en/software/junos/mvrp/multicast-l2/topics/concept/mvrp-mx-series-understanding.html>

upvoted 1 times

🗲️ 👤 **rfg8366** 1 year ago

C is correct <https://www.juniper.net/documentation/us/en/software/junos/mvrp/multicast-l2/topics/concept/mvrp-mx-series-understanding.html>

upvoted 1 times

🗲️ 👤 **rfg8366** 1 year ago

C is correct

upvoted 1 times

🗲️ 👤 **rfg8366** 1 year ago

C is correct

upvoted 1 times

🗲️ 👤 **rfg8366** 1 year ago

C is correct

upvoted 1 times

🗲️ 👤 **harrypogi** 1 year, 9 months ago

Selected Answer: A

a is correct


upvoted 1 times

Your enterprise network uses routing instances to support multitenancy. Your Junos devices use BGP to peer to multiple BGP devices. You must ensure that load balancing is achieved within the routing instance. Which two statements would accomplish this task? (Choose two.)

- A. Configure the multipath option at the [edit protocols bgp group <group-name> neighbor] hierarchy.
- B. Configure the multipath option at the [edit protocols bgp group] hierarchy.
- C. Configure a load-balance per-packet policy and apply it at the [edit routing-options forwarding-table] hierarchy.
- D. Configure the multipath option at the [edit routing-instances <instance-name> routing-options] hierarchy.

Suggested Answer: BD

Community vote distribution



🗳️ 👤 **kim1276** 7 months, 2 weeks ago

C and D looks perfect
upvoted 1 times

🗳️ 👤 **Kornrono** 9 months ago

Selected Answer: CD

The question is asking about the load balancing of BGP Protocol within routing-instance (l3vpn) also the correct answer is "C and D"
<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-load-balancing.html#d74e318>
upvoted 1 times

🗳️ 👤 **h4lv4r3** 1 year, 4 months ago

Answer is CD.

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-load-balancing.html>
upvoted 2 times

🗳️ 👤 **somanyquestions** 1 year, 8 months ago

Selected Answer: CD

c and d it is
upvoted 2 times

🗳️ 👤 **harrypogi** 1 year, 9 months ago

Selected Answer: CD

CD
the question asked for load balancing within routing instance, not just BGP
upvoted 3 times

🗳️ 👤 **Mahdi** 1 year, 11 months ago

Should be BC
upvoted 1 times

🗳️ 👤 **piipo** 2 years ago

Selected Answer: BC

BC
<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/load-balancing-bgp-session.html>
upvoted 1 times

🗳️ 👤 **hbstyleboy** 2 years, 3 months ago

Selected Answer: CD

should be c n d
upvoted 2 times

🗳️ 👤 **dragosky** 2 years, 5 months ago

Selected Answer: CD

i was wrong, actually is CD:

To enable these features on a Layer 3 VPN routing instance, include the `vpn-unequal-cost equal-external-internal` statement at the `[edit routing-instances routing-instance-name routing-options multipath]` hierarchy level and the `vrf-table-label` statement at the `[edit routing-instances routing-instance-name]` hierarchy level.

upvoted 3 times

  **dragosky** 2 years, 5 months ago

B,C

`[edit protocols bgp group external]`

`user@R1# set multipath`

`[edit policy-options policy-statement loadbal]`

`user@R1# set from route-filter 10.0.0.0/16 orlonger`

`user@R1# set then load-balance per-packet`

Apply the load-balancing policy.

`content_copy zoom_out_map`

`[edit routing-options]`

`user@R1# set forwarding-table export loadbal`

upvoted 1 times

  **DarkSpirit** 2 years, 6 months ago

Selected Answer: CD

I think should be C and D, because of using multipath inside L3VPN

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/ref/statement/multipath-edit-routing-options.html>

upvoted 1 times

  **penguin02007** 2 years, 6 months ago

Should be B and C.

upvoted 1 times

  **mohdema** 2 years, 6 months ago

Selected Answer: BC

Fortunately, the Juniper Networks BGP implementation supports the notion of a bandwidth community. This extended community encodes the bandwidth of a given next hop, and when combined with multipath, the load-balancing algorithm distributes flows across the set of next hops proportional to their relative bandwidths. Put another way, if you have a 10-Mbps and a 1-Mbps next hop, on average nine flows will map to the high-speed next hop for every one that uses the low speed.

Use of BGP bandwidth community is supported only with per-packet load balancing.

The configuration task has two parts:

Configure the external BGP (EBGP) peering sessions, enable multipath, and define an import policy to tag routes with a bandwidth community that reflects link speed.

Enable per-packet (really per-flow) load balancing for optimal distribution of traffic.

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/load-balancing-bgp-session.html>

upvoted 4 times

You are asked to enforce user authentication using a captive portal before users access the corporate network. Which statement is correct in this scenario?

- A. HTTPS is the default protocol for a captive portal.
- B. A captive portal can be bypassed using an allowlist command containing a device's IP address.
- C. When enabled, a captive portal must be applied to each individual interface.
- D. All Web browser requests are redirected to the captive portal until authentication is successful.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **kim1276** 7 months, 2 weeks ago

correct answer: C. When enabled, a captive portal must be applied to each individual interface
upvoted 1 times

🗳️ 👤 **Theker** 1 year, 10 months ago

D is correct
upvoted 1 times

🗳️ 👤 **mohdema** 2 years, 6 months ago

Selected Answer: D

You can set up captive portal authentication on your switch to redirect all Web browser requests to a login page that requires users to input a username and password before they are allowed access. Upon successful authentication, users are allowed access to the network and redirected to the original page requested.

Junos OS provides a customizable template for the captive portal window that allows you to easily design and modify the look of the captive portal login page. You can modify the design elements of the template to change the look of your captive portal login page and to add instructions or information to the page. You can also modify any of the design elements of a captive portal login page.

The first screen displayed before the captive login page requires the user to read the terms and conditions of use. By clicking the Agree button, the user can access the captive portal login page.

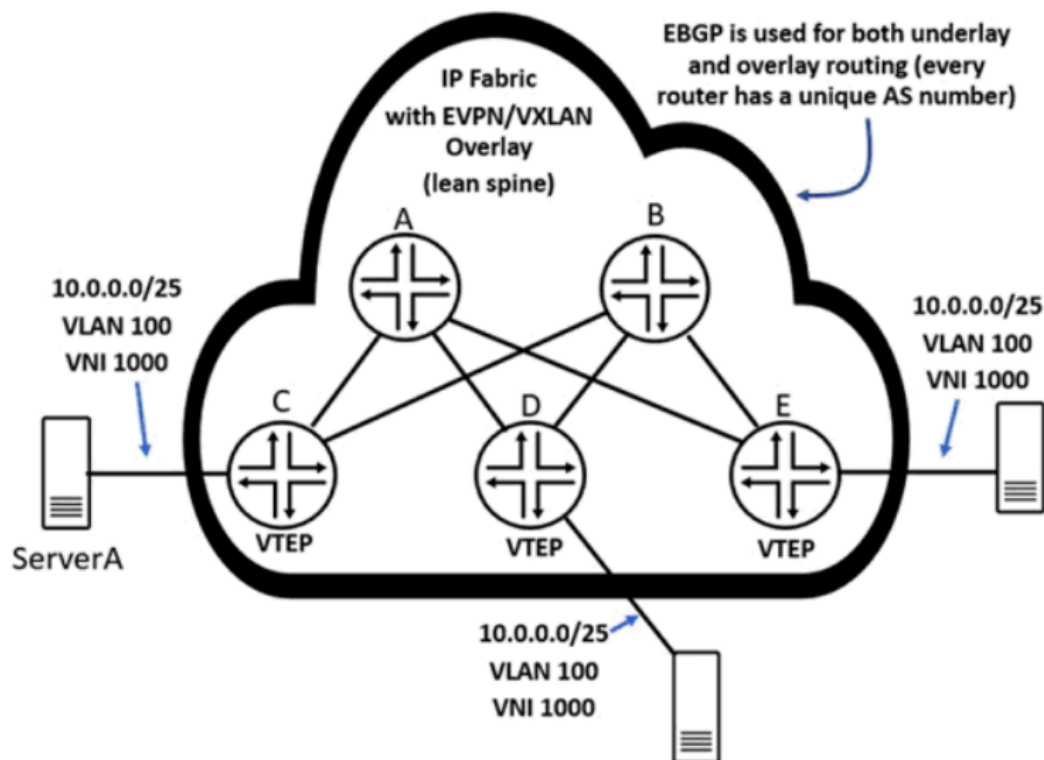
<https://www.juniper.net/documentation/us/en/software/junos/user-access/topics/topic-map/user-authentication-captive-portal.html>

upvoted 2 times

Click the Exhibition button.

Referring to the exhibit, ServerA sends a single IP packet destined to 10.0.0.127.

Which two statements correctly describe the behavior of the resulting outbound VXLAN packets that contain the original packet destined to 10.0.0.127? (Choose two.)



- A. Router E will replicate and send a copy of the received VXLAN packet to router D.
- B. Router C will send a VXLAN packet destined only to router D and router E.
- C. Router D will not replicate and send a copy of the received VXLAN packet to router E.
- D. Router C will send a single VXLAN packet to one remote VTEP.

Suggested Answer: AD

Community vote distribution

BC (100%)

JoeSun Highly Voted 2 years, 6 months ago

Selected Answer: BC

It is about how EVPN handle BUM traffic. The source leaf will send the unknown unicast traffic to all leaf that with same VLAN/VTEP.
upvoted 8 times

MichelT Most Recent 3 months ago

Selected Answer: BC

10.0.0.127 is the broadcast address for 10.0.0.0/25. Ingress replication leads to answer B and C.
upvoted 1 times

kim1276 7 months, 2 weeks ago

correct answer: B. Router C will send a VXLAN packet destined only to router D and router E.
C. Router D will not replicate and send a copy of the received VXLAN packet to router E.
upvoted 2 times

somanyquestions 1 year, 8 months ago

Selected Answer: BC

The source leaf (C) will send the broadcast traffic to all VTEPs of leafs that with same VLAN (D and E). those routers will not forward it again...
upvoted 2 times

🗨️ 👤 **harrypogi** 1 year, 9 months ago

Selected Answer: BC

bc is correct

upvoted 1 times

🗨️ 👤 **Yunas** 1 year, 10 months ago

BC is correct,

ingress BUM traffic would be sent by local PE, other PEs will not replicate BUM traffic to remote PEs

upvoted 1 times

🗨️ 👤 **Anarky19** 2 years ago

Router E will replicate and send a copy of the received VXLAN packet to router D

Ingress Replication

Ingress replication involves an ingress device replicating every incoming BUM packet and sending them as a separate unicast to the remote egress devices. Ingress replication happens through EVPN route type 3, also called as inclusive multicast ethernet tag (IMET) route

upvoted 2 times

🗨️ 👤 **Anarky19** 2 years ago

Router C will send a VXLAN packet destined only to router D and router E

Underlay Multicast:


In underlay multicast, the underlay network replicates the traffic through a multicast group. Forwarding BUM traffic using underlay multicast requires the configuration of IP multicast in the underlay network. A single copy of the BUM traffic moves from the ingress or source VTEP towards the underlay transport network. The network forwards this copy along the multicast tree so that it reaches all egress or destination VTEPs participating in the given multicast group.

upvoted 1 times

Click the Exhibition button.

A BGP network has been designed to provide resiliency and redundancy to a multihomed customer network.

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



```
user@R1# show protocols bgp group ext-peers
type external;
local-address 192.168.0.1;
peer-as 64522;
neighbor 172.16.255.2 {
  multihop {
    ttl 1;
  }
}
```

```
user@R1# show routing-options
autonomous-system 65411;
static {
  route 172.16.255.2/32 next-hop [ 10.10.1.2 10.10.2.2 ];
}
```

```
user@R1> show route 172.16.255.2/32 terse
```

```
inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	V	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	?	172.16.255.2/32	S	5			>10.10.1.2 10.10.2.2	

```
user@R1> show route forwarding-table matching 172.16.255.2/32
```

```
Routing table: default.inet
```

```
Internet:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
172.16.255.2/32	user	1	10.10.1.2	ucst	590	5	ge-0/0/0.0

- A. Both the next hops will be used to forward traffic to R2.
- B. A routing policy will be required to forward traffic to both next hops.
- C. The TTL value of 1 is set to limit the scope of the EBGP session.
- D. The ttl statement must be configured to accommodate peering to a loopback address of a directly connected peer.

Suggested Answer: BD

Community vote distribution

BC (75%)

BD (25%)

kim1276 7 months, 2 weeks ago

correct answer

- B. A routing policy will be required to forward traffic to both next hops.
 - C. The TTL value of 1 is set to limit the scope of the EBGP session
- upvoted 2 times

Anarky19 1 year, 6 months ago

B and C are correct.

D can't be right because of this: Each router must establish the peering session with the loopback address of the remote router. You can configure this session with the local-address statement, which alters the peer address header information in the BGP packets.

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/multihop-sessions.html#:~:text=Each%20router%20must%20establish%20the%20peering%20session%20with%20the%20loopback%20address%20of%20the%20remote>
upvoted 1 times

🗳️ 👤 **somanyquestions** 1 year, 8 months ago

Selected Answer: BC

- A. nope! clearly there is only one next-hop
 - B. Yes! A routing policy will be required to forward traffic to both next hops.
 - C. Yes! multihop is needed to extent the scope! but ttl is set to limit the scope of an EBGp sessions loopback;
<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/multihop-sessions.html>
 - D. nope! Without the multihop statement, EBGp will only work on interface addresses of directly connected peer.
- upvoted 2 times

🗳️ 👤 **harrypogi** 1 year, 9 months ago

Selected Answer: BC

bc correct
upvoted 1 times

🗳️ 👤 **bendarkel** 1 year, 11 months ago

Selected Answer: BD

How can C be correct when peering eBGp using loopbacks? Shouldn't a TTL of 2 be required to bring up the peering? The packet arrives with a TTL of one. The TTL is decremented to 0 and the packet discarded. Setting the TTL to 2 allows the peering to form using the loopbacks.
upvoted 3 times

🗳️ 👤 **somanyquestions** 1 year, 8 months ago

see; <https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/multihop-sessions.html> "We use a TTL value of 1 to ensure that the session cannot be established across any other backdoor links in the network. "
upvoted 1 times

🗳️ 👤 **65briang** 1 year, 6 months ago

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/multihop-sessions.html>
NOTE: When multihop is configured, the Junos OS sets the TTL value of 64, by default.
A TTL value of 1 is sufficient to enable an EBGp session to the loopback address of a directly connected neighbor.
upvoted 1 times

🗳️ 👤 **Takigama** 2 years ago

Selected Answer: BC

- BC
- A. we can see the forwarding table only has one next hop (cant be this)
 - B. technically true (but not the only requirement)
 - C. true enough
 - D. Technically the routers are directly connected, just not the peers cause its loopback to loopback, so in some ways its a semantics/wording of the question that needs to be noted.
- upvoted 2 times

🗳️ 👤 **Anarky19** 2 years ago

- B. You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:
Per-packet load balancing
Class of service (CoS)
- upvoted 1 times

🗳️ 👤 **DarkSpirit** 2 years, 6 months ago

Selected Answer: BC

- B - because only one next hop in forwarding table
 - C- becuse TTL=1 limiting bgp connection to be established only over loopbacks interface
<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/multihop-sessions.html>
Use the multihop statement to alter the default use of the neighbor's physical address. In addition, you can also specify a time-to-live (TTL) value in the BGP packets to control how far they propagate. We use a TTL value of 1 to ensure that the session cannot be established across any other backdoor links in the network.
- upvoted 4 times

Click the Exhibition button.

You are asked to troubleshoot voice quality issues on your newly implement VoIP network. You notice that the voice packets are being dropped.

You have verified that the packets are correctly marked for expedited forwarding queue.

Referring to the exhibit, what must you configure to solve the problem?

```

[edit]
user@R1# show class-of-service
classifiers {
    dscp voip {
        import default;
    }
}
interfaces {
    ge-1/0/0 {
        unit 0 {
            classifiers {
                dscp voip;
            }
        }
    }
}
user@R1> show interfaces ge-1/0/0 extensive
Physical interface: ge-1/0/0, Enabled, Physical link is Up
Interface index: 154, SNMP ifIndex: 527, Generation: 157
Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error: None, Loop Detect PDU Error:
None,
Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control:
Enabled,
Auto-negotiation: Enabled, Remote fault: Online
Pad to minimum frame size: Disabled
Media type: Copper
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Auto-negotiation: Enabled, Remote fault: Online
Pad to minimum frame size: Disabled
Media type: Copper
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Schedulers : 0
Hold-times : Up 0 ms, Down 0 ms
Damping : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state: unsuppressed
Current address: 4c:96:14:93:9a:95, Hardware address: 4c:96:14:93:9a:95
Last flapped : 2022-05-16 11:44:33 PDT (21:23:22 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 894761 0 bps
Output bytes : 681004 240 bps
Input packets: 13083 0 pps
Output packets: 11321 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Dropped traffic statistics due to STP State:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Total errors 0 0
Filter statistics:
Input packet count 13083
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 11320
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 0 (0x00)
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority Limit
% bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none
Interface transmit statistics: Disabled



```

- A. You must configure a multifeild classifier to put the VoIP traffic in the correct queue.
- B. You must configure a rewrite rule to ensure that the traffic is scheduled properly in the device.
- C. You must configure a scheduler to allocate bandwidth to the expedited forwarding queue.
- D. You must configure a policer to ensure that the queue is not being starved.

Suggested Answer: *C*

Community vote distribution



  **harrypogi** 9 months, 1 week ago

Selected Answer: C

C


by default only BE and NC are allocated with bandwidth

upvoted 3 times

Click the Exhibit button.

Remote branches connect to the corporate WAN through access switches. The access switches connect to access ports on the WAN distribution switch, as shown in the exhibit. L2PT has previously been configured on the tunnel Layer 2 traffic across the WAN. You decide to move the L2PT tunnel endpoints to the access switches. When you apply the L2PT configuration to the access switches, the ports that connect the access switches to the distribution switch shut down.

Which action would solve this problem?



- A. Configure the links between the access switches and the distribution switch as a trunk port.
- B. Disable the BPDU block function on the access switches.
- C. Disable the BPDU block function on the distribution switch.
- D. Configure a GRE tunnel to encapsulate the L2PT traffic across the WAN.

Suggested Answer: A

Community vote distribution

A (100%)

Alex_NYC 3 months, 3 weeks ago

Selected Answer: B

Answer B

Note: Access interfaces in an L2PT-enabled VLAN should not receive L2PT-tunneled PDUs. If an access interface does receive L2PT-tunneled PDUs, there might be a loop in the network, and the device will shut down the interface.

upvoted 1 times

harrypogi 9 months ago

Selected Answer: A

a correct

upvoted 1 times

mohdema 1 year, 6 months ago

Selected Answer: A

Access interfaces in an L2PT-enabled VLAN should not receive L2PT-tunneled PDUs. If an access interface does receive L2PT-tunneled PDUs, there might be a loop in the network, and the device will shut down the interface.

<https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/topic-map/layer-2-protocol-tunneling.html>

upvoted 2 times

  **penguin02007** 1 year, 6 months ago

A is correct.


<https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/topic-map/layer-2-protocol-tunneling.html#d116e87>

upvoted 1 times

Click the Exhibit button.

You want to provide Layer 2 connectivity between campus sites using Ethernet switches through a metro Ethernet service provider who is using Q-in-Q tagging on their network.


Referring to the exhibit, what are two design considerations in this environment? (Choose two.)



- A. VXLAN could be implemented on your network across this service provider network.
- B. Each campus switch shown must have a C-Tag 300 configured.
- C. L2PT is required on the SP network to support the spanning tree protocol.
- D. Each campus switch shown must have S-Tag 300 configured.

Suggested Answer: DC

Community vote distribution



somanyquestions Highly Voted 1 year, 8 months ago

Selected Answer: AC

A. VXLAN could be implemented on your network across this service provider network.
could be a valid option.

B. Each campus switch shown must have a C-Tag 300 configured.

no, since the provider is using S-Tag 300 and QinQ, you can use C-Tag 300 but also other, it's not a must.

C. L2PT is required on the SP network to support the spanning tree protocol.

yes, is needed!

D. Each campus switch shown must have S-Tag 300 configured.

no, the S-Tag is taken care of the providers QinQ

upvoted 7 times

oz1i Most Recent 11 months, 3 weeks ago

Selected Answer: AC

AC is correct.

upvoted 1 times

rfg8366 1 year ago

Is fair to say, that considering that this is not an EVPN related question, CD do seem to be correct

upvoted 1 times

Yunas 1 year, 10 months ago

A can be correct as we have layer 3 access to the all Camnpus's switche loopback ip address.

upvoted 1 times

sangki 2 years, 5 months ago

i think bc

s-tag is handled by SP.

upvoted 3 times

🗨️ 👤 **JoeSun** 2 years, 6 months ago

Selected Answer: BC

VXLAN , L2PT and S-Tag(service VLAN) are handled by ISP.

C-Tag (customer VLAN) is handled by customer. In fact, in customer view, 2 campus switches are direct connected to each others so that they have to configure same vlan info whatever it is trunk or access port.

upvoted 1 times

🗨️ 👤 **mohdema** 2 years, 6 months ago

Selected Answer: CD

<https://www.juniper.net/documentation/us/en/software/junos/multicast-l2/topics/ref/statement/layer2-protocol-tunneling-edit-vlans-l2pt-ex-series.html>

upvoted 4 times

🗨️ 👤 **hbstyleboy** 2 years, 3 months ago

What is this B.S ?

The link doesn't even justify your answers at all. Don't troll plz.

Administer should ban MF like u

upvoted 2 times

You must ensure that all routes in the 10.0.0/8 address range are not advertised outside of your AS.
Which well-known BGP community should be assigned to these addresses to accomplish this task?

- A. no-export
- B. no-peer
- C. internet
- D. no-advertise

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ **harrypogi** 9 months ago

Selected Answer: A

a no export

upvoted 1 times

🗨️ **mohdema** 1 year, 6 months ago

Selected Answer: A

For specifying the BGP community attribute only, you also can specify community-ids as one of the following well-known community names defined in RFC 1997:

no-advertise—Routes containing this community name are not advertised to other BGP peers.

no-export—Routes containing this community name are not advertised outside a BGP confederation boundary.

no-export-subconfed—Routes containing this community are advertised to IBGP peers with the same AS number, but not to members of other confederations.

llgr-stale—Adds a community to a long-lived stale route when it is readvertised.

no-llgr—Marks routes which a BGP speaker does not want to be retained by LLGR. The Notification message feature does not have any associated configuration parameters.

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/statement/community-edit-routing-options.html>

upvoted 4 times

Which statement is correct about IS-IS?

- A. IS-IS uses areas and an autonomous system.
- B. Level 1/2 routers automatically inject a default route to the nearest Level 1 router.
- C. Level 2 routers must share the same area address.
- D. Level 1 routers route traffic between autonomous systems.

Suggested Answer: A

Community vote distribution

A (73%)

B (27%)

 **mohdema** Highly Voted 2 years, 6 months ago

Selected Answer: A

Level 1/2 routers automatically inject a default route to the nearest Level 1 router.

It's the other way around

upvoted 9 times

 **FR99** Most Recent 1 month, 1 week ago

Selected Answer: A

In IS-IS, a single AS can be divided into smaller groups called areas.

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>

upvoted 1 times

 **Balanar** 1 year ago

Selected Answer: A

When a routing device that operates as both a Level 1 and Level 2 router determines that it can reach at least one area other than its own, it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router.

First: L1/2 router is set attached bit to its L1 PDU, if only if there is another area other than its own area. if there is only one are, it will not do it.

Second: the default route is introduced by the L1 only router.

upvoted 1 times

 **jprat** 1 year, 5 months ago

Selected Answer: A

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>

upvoted 1 times

 **somanyquestions** 1 year, 8 months ago

Selected Answer: A

for B to be correct, it is not the L1/L2 router who adds the default route, but the L1 router connected to based on the ATTACHED bit.

"When a routing device that operates as both a Level 1 and Level 2 router (Router B) determines that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router (Router A) introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router (Router B)."

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>

A is then the most accurate. areas. the s as multiple = 2. an AS. so single AS..

upvoted 3 times

 **harrypogi** 1 year, 9 months ago

Selected Answer: A

a is correct

upvoted 1 times

 **piipo** 2 years ago

Selected Answer: B

B

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html>

upvoted 1 times

🗨️ 👤 **hbstyleboy** 2 years, 3 months ago

Selected Answer: A

No, ISIS (Intermediate System to Intermediate System) Level 1/2 routers do not automatically inject a default route to the nearest Level 1 router.

upvoted 2 times

🗨️ 👤 **hbstyleboy** 2 years, 3 months ago

Ans is A

No, ISIS (Intermediate System to Intermediate System) Level 1/2 routers do not automatically inject a default route to the nearest Level 1 router.

In ISIS, Level 1 routers only advertise their connected networks to other Level 1 routers within the same area. Level 2 routers advertise their connected networks to other Level 2 routers throughout the entire ISIS domain. When a Level 2 router receives a route from another Level 2 router, it can forward traffic to that destination directly or through a Level 1 router.

If a router needs a default route, it must be manually configured or learned through some other means, such as a dynamic routing protocol like BGP (Border Gateway Protocol) or static routing. Default routes are used to forward traffic to a destination when no specific match is found in the routing table.

upvoted 2 times

🗨️ 👤 **sanalainen** 2 years, 3 months ago

Selected Answer: B

When a routing device that operates as both a Level 1 and Level 2 router (Router B) determines that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router (Router A) introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router (Router B).

[https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html#routing-is-is-overview__d12254e173]

upvoted 1 times

🗨️ 👤 **mistadave** 2 years, 5 months ago

Selected Answer: B

B is your answer here

upvoted 2 times

🗨️ 👤 **dragosky** 2 years, 5 months ago

Is the other way around:

level 1 routers generate a default route to the closest level 1-2 router to reach prefixes outside of their own area.

level 1-2 routers can redistribute prefixes from level 2 to level 1 so that level 1 routes can choose the most optimal path.

upvoted 2 times

🗨️ 👤 **dragosky** 2 years, 5 months ago

A

An IS-IS network is a single autonomous system (AS), also called a routing domain, that consists of end systems and intermediate systems. End systems are network entities that send and receive packets. Intermediate systems send and receive packets and relay (forward) packets. (Intermediate system is the Open System Interconnection [OSI] term for a router.) ISO packets are called network PDUs.

In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area; when the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs. No IS-IS area functions strictly as a backbone.

upvoted 1 times

🗨️ 👤 **davidrsr** 2 years, 6 months ago

Selected Answer: B

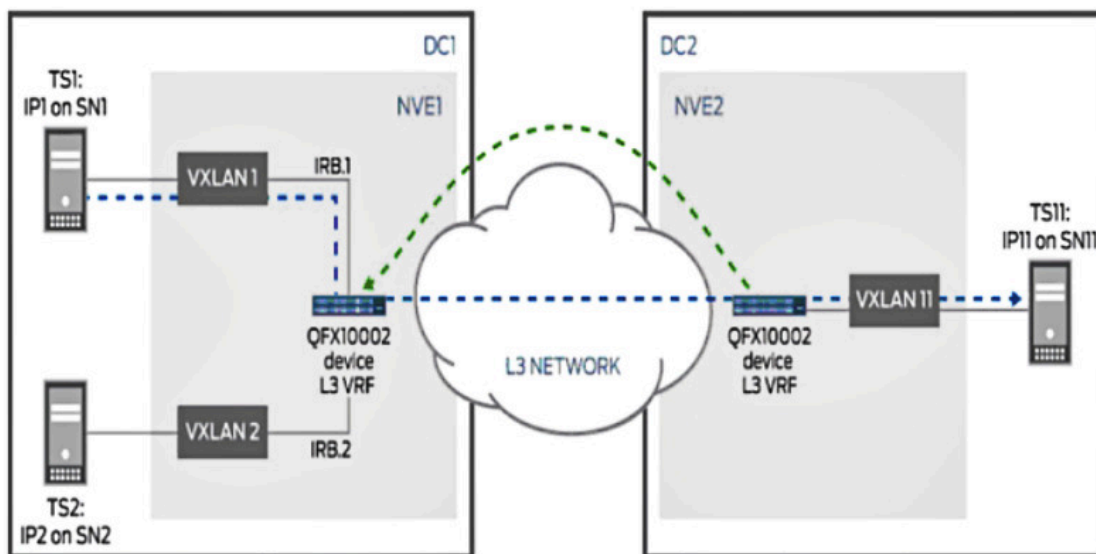
ISIS use area and level for peering, not AS. And L1/2 router generate default route to Level 1 is correct.

upvoted 2 times

Click the Exhibit button.

The connection between DC1 and DC2 is routed as shown in the exhibit.

In this scenario, which statement is correct?



- A. The border devices must be able to perform Layer 3 routing and provide IRB functionality.
- B. L3VPN must be enabled to advertise reachability.
- C. An IP prefix route provides encoding for intra-subnet forwarding.
- D. Type 2 and Type 5 routes will be exchanged between DC1 and DC2.

Suggested Answer: A

Community vote distribution

A (80%)

D (20%)

penguin02007 Highly Voted 2 years ago

A is correct.

upvoted 6 times

random_name Highly Voted 2 years, 1 month ago

May be correct

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-route-type5-understanding.html>

upvoted 5 times

h4lv4r3 Most Recent 11 months ago

A is correct.

D cant be correct cause no where does it state what protocol it is using.

upvoted 2 times

harrypogi 1 year, 3 months ago

Selected Answer: A

a correct

upvoted 1 times

piipo 1 year, 6 months ago

Selected Answer: A

A

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-route-type5-understanding.html>



upvoted 2 times

hbstyleboy 1 year, 9 months ago

Selected Answer: A

Agree. Type 2 is not needed. So D is not correct.

upvoted 1 times

  **chicharito14** 1 year, 10 months ago

I think it does not need Type 2 so A is correct

upvoted 2 times

  **dragossky** 1 year, 11 months ago

Selected Answer: D

type 2 and type 5

upvoted 1 times

BGP multipath or multihop are not configured in your network.

In this scenario, what is the correct sequence for BGP active route selection?

A. higher local preference

shortest AS path

lowest peer address

lowest router ID

lower origin code

B. higher local preference

shortest AS path

lower origin code

lowest router ID

lowest peer address

C. higher local preference

lowest router ID

lowest peer address

lower origin code

shortest AS path

D. higher local preference

shortest AS path

lowest router ID


lowest peer address

lower origin code

Suggested Answer: D

Community vote distribution

B (100%)

 **mohdema** Highly Voted 1 year, 6 months ago

Selected Answer: B

BGP Path Selection process follows this order

1. Weight (Bigger is better)
2. Local preference (Bigger is better)
3. Self originated (Locally injected is better than iBGP/eBGP learned)
4. AS-Path (Smaller is better)
5. Origin
6. MED (Smaller is better)
7. External (Prefer eBGP over iBGP)
8. IGP cost (Smaller is better)
9. EBGP Peering (Older is better)
10. Router- ID

<http://www.next-itsolutions.co.uk/wp-content/uploads/2015/09/CCIE-BGP-Best-Path-Selection1.png>

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/bgp/topics/concept/routing-protocols-address-representation.html>

upvoted 7 times

 **harrypogi** Most Recent 9 months, 1 week ago

Selected Answer: B

b is correct

upvoted 1 times

 **Rudy6969** 1 year, 5 months ago

go with B

upvoted 2 times

🗒️ 👤 **carlitox** 1 year, 7 months ago

seems like B:

Algorithm steps

BGP Next Hop—Protocol

Local Preference—Localpref

AS Path—AS path

Origin—AS

Multiple Exit Discriminator

EBGP vs. IBGP—Local AS

Cost to IGP peer—Metric

Cluster List

Router ID


Peer ID

Multipath ignores router and peer id I cannot find changes by multihop
upvoted 4 times

Click the Exhibit button.

You are troubleshooting an EVPN-VXLAN IP fabric and observe the loop shown in the exhibit.

Which two steps would you take to further troubleshoot this problem? (Choose two.)



- A. Verify that the same ESI is configured on the link from the host and that it matches the source.
- B. Issue the `show route table bgp.evpn.0` command on Leaf2 and verify that Type 4 routes are present.
- C. Issue the `show route table bgp.evpn.0` command on Leaf2 and verify that Type 3 routes are present.
- D. Verify that the same ESI is configured on the two links from the source.

Suggested Answer: BC

Community vote distribution

BD (100%)

harrypogi 9 months, 1 week ago

Selected Answer: BD

bd is correct

upvoted 2 times

davidrsr 1 year, 6 months ago

Selected Answer: BD

EVPN type 4 routes - Ethernet Segment Routes are needed in multi-homing scenario and used for Designated Forwarder Election.

upvoted 4 times

mohdema 1 year, 6 months ago

Type 2 route, MAC with IP advertisement route—Type 2 routes are per-VLAN routes, so only PEs that are part of a VNI need these routes. EVPN allows an end host's IP and MAC addresses to be advertised within the EVPN Network Layer reachability information (NLRI). This allows for control plane learning of ESI MAC addresses. Because there are many Type 2 routes, a separate route-target auto-derived per VNI helps to confine their propagation. This route type is supported by all EVPN switches and routers.

Type 5 route, IP prefix Route—An IP prefix route provides encoding for inter-subnet forwarding. In the control plane, EVPN Type 5 routes are used to advertise IP prefixes for inter-subnet connectivity across data centers. To reach a tenant using connectivity provided by the EVPN Type 5 IP prefix route, data packets are sent as Layer 2 Ethernet frames encapsulated in the VXLAN header over the IP network across the data centers.

upvoted 1 times

Click the Exhibit button.

Referring to the outputs shown in the exhibit, which two statements are correct about the IS-IS adjacency? (Choose two.)

```


user@R1> show isis adjacency extensive
R2
  Interface: ge-1/0/0.0, Level: 2, State: Up, Expires in 7 secs
  Priority: 64, Up/Down transitions: 1, Last transition: 00:02:19 ago
  Circuit type: 2, Speaks: IP, IPv6, MAC address: 4c:96:14:93:9a:96
  Topologies: Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: R2.02, IP addresses: 10.1.1.2
  Transition log:
  When                State      Event      Down reason
  Mon May 16 11:53:33  Up        Seenself
user@R2> show isis adjacency extensive
R1
  Interface: ge-1/0/1.0, Level: 2, State: Up, Expires in 20 secs
  Priority: 64, Up/Down transitions: 1, Last transition: 00:01:55 ago
  Circuit type: 3, Speaks: IP, IPv6, MAC address: 4c:96:14:93:9a:95
  Topologies: Unicast
  Restart capable: No, Adjacency advertisement: Advertise
  LAN id: R2.02, IP addresses: 10.1.1.1
  Transition log:
  When                State      Event      Down reason
  Mon May 16 11:53:33  Up        Seenself

```

- A. R1 is configured to participate in both Level 1 and Level 2.
- B. R2 is configured to participate in both Level 1 and Level 2.
- C. R1 is configured to participate in Level 2 only.
- D. R2 is configured to participate in Level 2 only.

Suggested Answer: BC

Community vote distribution



Takigama Highly Voted 2 years ago

Selected Answer: BC

Circuit type identifies the router you are running the command ON, not the connected router.

In the diagram, R1 says Circuit type 2 - Therefore R1 (not R2) is level 2 only

In the diagram, R2 says Circuit type 3 - Therefore R2 (and not R1) is Level 1 and level 2

Other comments here say similar but say it wrong or have the wrong answer. The circuit type in "show isis adjacency extensive" output is always related to the router you ran the command on!

upvoted 13 times

tester2478 Highly Voted 1 year, 11 months ago

Selected Answer: AD

Circuit type (detail output only): Displays the circuit type, which can be 1—Level 1 only, 2—Level 2 only, or 3—Level 1 and Level 2;

from JNCIS-ENT study guide

upvoted 8 times

e98fa04 Most Recent 3 months, 2 weeks ago

Selected Answer: BC

labbed this and configured my MX-1 to disable level 1 and it says circuit type 2 while MX-2 says circuit 3
upvoted 1 times

🗨️ 👤 **Alex_NYC** 3 months, 3 weeks ago

Selected Answer: AD

The circuit type in the output of show isis adjacency extensive refers to the local router's capability on that interface, not the neighbor's.

It indicates what type of IS-IS adjacency the local router is forming over that circuit (interface). However, the neighbor must support the same or a compatible circuit type for an adjacency to be established.

For example:

If your router is configured as L1, it can only form L1 adjacencies.

If it is L2, it will only form L2 adjacencies.

If it is L1/L2, it can form both types depending on the neighbor's capability.

upvoted 1 times

🗨️ 👤 **bob1296** 8 months ago

100% B and C. I labbed this out, only takes a minute or so to spin up a simple ISIS direct topology. When modifying the Level type on an interface via the "set protocols isis interface <name> level {1 or 2} disable" command, the circuit type changes according to how the LOCAL router is configured when doing a "show isis adjacency extensive" command.

upvoted 1 times

🗨️ 👤 **Totalstranger** 1 year, 1 month ago

Selected Answer: AD

Circuit 1 - L1

Circuit 2 - L2

Circuit 3 - L1/L2

upvoted 3 times

🗨️ 👤 **scorpfe** 1 year, 5 months ago

Selected Answer: BC

1—Level 1 only

2—Level 2 only

3—Level 1 and Level 2

upvoted 1 times

🗨️ 👤 **jprat** 1 year, 5 months ago

Selected Answer: BC

@R1 Circuit Type:2, and Level: 2 (mean only in level 2)

@R2 Circuit Type:3, and Level: 2 (mean on level 1 & level 2)

upvoted 1 times

🗨️ 👤 **scorpfe** 1 year, 5 months ago

Selected Answer: BC

Circuit type:

1—Level 1 only

2—Level 2 only

3—Level 1 and Level 2

upvoted 1 times

🗨️ 👤 **harrypogi** 1 year, 9 months ago

Selected Answer: BC

bc is correct

takigama is right

upvoted 1 times

🗨️ 👤 **Jammyx5** 1 year, 11 months ago

AD

Circuit type:

- 1 –Level 1 only.
 - 2 –Level 2 only.
 - 3 –Level 1 and Level 2
- upvoted 3 times

🗲️ 👤 **piipo** 2 years ago

Selected Answer: AD

AD

Circuit type

upvoted 5 times

🗲️ 👤 **Rudy6969** 2 years, 5 months ago

A and D

upvoted 2 times

🗲️ 👤 **dragosky** 2 years, 6 months ago

Selected Answer: AD

CC types 2 - L2 only 3- L1/L2

upvoted 6 times

🗲️ 👤 **mohdema** 2 years, 6 months ago

Selected Answer: BC

R1 (and not R1) has circuit type 2 which means that R2 is configured to participate in Level 2 only, answer D is correct.

R2 (and not R1) has circuit type 3 which means that R1 is configured to participate in both Level 1 and Level 2, answer A is correct.

upvoted 6 times

🗲️ 👤 **penguin02007** 2 years, 6 months ago

A and D is correct.

upvoted 2 times

🗲️ 👤 **Obodio** 2 years, 7 months ago

R2 has circuit type 2 which means that R2 is configured to participate in Level 2 only, answer D is correct.

R1 has circuit type 3 which means that R1 is configured to participate in both Level 1 and Level 2, answer A is correct.

upvoted 3 times

Which two multicast listener registration protocols are supported in the Junos operating system? (Choose two.)

- A. MLD
- B. DVMRP
- C. IGMP
- D. PIM

Suggested Answer: AC

Community vote distribution

AC (100%)

🗨️ **DinoVercotti** 7 months ago

Selected Answer: AC

AC is correct:

https://www.cisco.com/assets/sol/sb/Switches_Emulators_v2_3_5_xx/help/250/index.html#page/tesla_250_olh/multicast_snooping.html

upvoted 1 times

🗨️ **harrypogi** 9 months, 1 week ago

Selected Answer: AC

ac is correct

upvoted 1 times

🗨️ **mohdema** 1 year, 6 months ago

Selected Answer: AC

Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) are the Multicast Group Membership Discovery (MGMD) protocols

upvoted 3 times


Which three statements are correct about EVPN route types? (Choose three.)

- A. Type 3 routes carry replication information.
- B. Type 2 routes carry endpoint MAC address information.
- C. Type 2 routes carry endpoint IP address information.
- D. Type 5 routes carry replication information.
- E. Type 1 routes carry endpoint MAC address information.

Suggested Answer: BCE

Community vote distribution

ABC (100%)

 **harrypogi** 9 months, 1 week ago

Selected Answer: ABC

abc is correct

upvoted 2 times

 **Mahdi** 11 months, 2 weeks ago

should be ABC

upvoted 2 times

 **mohdema** 1 year, 6 months ago

Selected Answer: ABC

Cisco explains it better: The EVPN control plane advertises the following types of information:

Route type 1 – This is an Ethernet Auto-Discovery (EAD) route type used to advertise Ethernet segment identifier, Ethernet Tag ID, and EVPN instance information. EAD route advertisements may be sent for each EVPN instance or for each Ethernet segment.

Route type 2 – This advertises endpoint reachability information, including MAC and IP addresses of the endpoints or VTEPs.

Route type 3 – This performs multicast router advertisement, announcing the capability and intention to use ingress replication for specific VNIs.

Route type 4 – This is an Ethernet Segment route used to advertise the Ethernet segment identifier, IP address length, and the originating router's IP address.

Route type 5 – This is an IP prefix route used to advertise internal IP subnet and externally learned routes to a VXLAN network.

upvoted 4 times

 **mohdema** 1 year, 6 months ago

Selected Answer: ABC

Type 1 route, Ethernet autodiscovery route

Type 2 route, MAC with IP advertisement route

Type 3 route, inclusive multicast Ethernet tag route

The tunnel types supported in an EVPN route type 3 when BGP-EVPN MPLS is enabled are ingress replication, P2MP MLDP, and composite tunnels

Type4 route, Ethernet segment Route

Type 5 route, IP prefix Route

upvoted 2 times

Click the Exhibit button.

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

```

user@leaf> show route table default-switch.evpn.0 detail
...
2:192.168.100.13:1::5010::00:0c:29:08:04:a0/304 MAC/IP (2 entries, 1 announced)
  *BGP      Preference: 170/-101
            Route Distinguisher: 192.168.100.13:1
            Next hop type: Indirect, Next hop index: 0
            Address: 0xcd690bc
            Next-hop reference count: 12
            Source: 192.168.100.1
            Protocol next hop: 192.168.100.13
            Indirect next hop: 0x2 no-forward INH Session ID: 0x0
            State: <Secondary Active Int Ext>
            Local AS: 65000 Peer AS: 65000
            Age: 8:17      Metric2: 0
            Validation State: unverified
            Task: BGP_65000.192.168.100.1
            Announcement bits (1): 0-default-switch-evpn
            AS path: I (Originator)
            Cluster list: 1.1.1.1
            Originator ID: 192.168.100.13
            Communities: target:65000:5010 encapsulation:vxlan(0x8)
            Import Accepted
            Route Label: 5010
            ESI: 00:00:00:00:00:00:00:00:00
            Localpref: 100
            Router ID: 192.168.100.1
            Primary Routing Table: bgp.evpn.0
            Thread: junos-main

```

- A. The host that the route is associated with is multihomed to two leaf nodes.
- B. The route is a Type 1 EVPN route.
- C. The route is a Type 2 EVPN route.
- D. The host that the route is associated with is single-homed to one leaf node.

Suggested Answer: BD

Community vote distribution

CD (100%)

🗳️ 👤 **TooManyQuestions** 8 months, 2 weeks ago

Selected Answer: CD

C start with 2: and is MAC/IP

D single homed since ESI is all zero's

upvoted 4 times

🗳️ 👤 **harrypogi** 9 months, 1 week ago

Selected Answer: CD

cd is correct

upvoted 2 times

🗳️ 👤 **dragosky** 1 year, 6 months ago

type 2, it starts with 2, so no brainer here....



upvoted 1 times

🗳️ 👤 **mohdema** 1 year, 6 months ago

Selected Answer: CD

Route type is 2 if not mistaken

upvoted 3 times

  **Obodio** 1 year, 7 months ago

Answers C, D - the route is Type 2 EVPN route...

upvoted 3 times

You must provide network connectivity to hosts that fail authentication.


In this scenario, what would be used in a network secured with 802.1X to satisfy this requirement?

- A. Configure the native-vlan-id parameter on the port.
- B. Use the server-reject-vlan command to specify a guest VLAN.
- C. Configure a secondary IP address on the port for unauthenticated hosts.
- D. Configure the port as a spanning tree edge port.

Suggested Answer: B

Community vote distribution

B (100%)

 **mohdema** Highly Voted 1 year, 6 months ago

Selected Answer: B

For a device configured for 802.1X authentication, specify that when the device receives an Extensible Authentication Protocol Over LAN (EAPoL) Access-Reject message during the authentication process between the device and the RADIUS authentication server, supplicants attempting to access the LAN are granted access and moved to a specific bridge domain or VLAN. Any bridge domain, VLAN name or VLAN ID sent by a RADIUS server as part of the EAPoL Access-Reject message is ignored.

upvoted 6 times

 **somanyquestions** 8 months, 2 weeks ago

<https://www.juniper.net/documentation/us/en/software/junos/user-access/topics/ref/statement/server-reject-vlan-edit-protocols-dot1x-authenticator-interface-802-1x.html>

upvoted 1 times

 **harrypogi** Most Recent 9 months, 1 week ago

Selected Answer: B

b is correct

upvoted 1 times

A Layer 2 connection does not extend across data centers. The IP subnet in a Layer 2 domain is confined within a single data center. Which EVPN route type is used to communicate prefixes between the data centers?

- A. Type 1
- B. Type 2
- C. Type 4
- D. Type 5

Suggested Answer: D

Community vote distribution

D (100%)

  **mohdema** Highly Voted 1 year, 6 months ago

Selected Answer: D

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-route-type5-understanding.html#:~:text=In%20the%20control%20plane%2C%20EVPN,subnet%20connectivity%20across%20data%20centers.>
upvoted 5 times

  **harrypogi** Most Recent 9 months, 1 week ago

Selected Answer: D

D is correct
upvoted 2 times

  **hbstyleboy** 1 year, 3 months ago

Type 5 is correct
upvoted 1 times

You are asked to implement fault tolerant RPs in your multicast network.
Which two solutions would accomplish this behavior? (Choose two.)

- A. Use BFD with statically defined RPs.
- B. Use MSDP with statically defined RPs.
- C. Use anycast PIM with statically defined RPs.
- D. Use IGMPv3 with statically defined RPs.

Suggested Answer: *CD*

Community vote distribution

BC (100%)

🗲️ 👤 **harrypogi** 9 months, 1 week ago

Selected Answer: BC

bc is correct

upvoted 1 times

🗲️ 👤 **davidrsr** 1 year, 6 months ago

Selected Answer: BC

IGMP is for join.

upvoted 4 times

🗲️ 👤 **Obodio** 1 year, 7 months ago


Answers: B, C

upvoted 4 times

Click the Exhibit button.

Your network has an unmanaged switch between the hosts and your EX Series switch. After the traffic enters the EX Series switch, each host must be on a separate VLAN.

How would you accomplish this task?



- A. Configure an input firewall filter on interface ge-0/0/3 to match the source MAC or IP address of the hosts to assign the VLANs.
- B. Configure an output firewall filter on interface ge-0/0/1 to match the destination MAC or IP address of the hosts to assign the VLANs.
- C. Configure interface ge-0/0/3 to a mode trunk to assign the VLANs.
- D. Configure VSTP on interface ge-0/0/1 to assign the VLANs.

Suggested Answer: C

Community vote distribution

A (100%)

DarkSpirit Highly Voted 2 years ago

Selected Answer: A

EX switches has VLAN action on ingress interfaces to forward traffic to specific VLAN.

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/topic-map/firewall-filter-match-conditions-and-actions-qfx.html>
vlan VLAN-name

Forward matched packets to a specific VLAN.

NOTE: The vlan action is supported on ingress interfaces only.

upvoted 7 times

anonymonkey Most Recent 10 months, 2 weeks ago

The key is contained within answer C itself as the diagram leads one to think the devices already reside within their respective VLANs. "Configure interface ge-0/0/3 to a mode trunk to assign the VLANs. VLANs within a trunk interface can be allowed not assigned. This function is accomplished via answer A. This is a very tricky question especially the way the diagram is represented.

A is the correct answer.

upvoted 2 times

harrypogi 1 year, 3 months ago

Selected Answer: A

a is correct