



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



CERTIFICATION TEST

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

Which statement is true about an OSPF broadcast link?

- A. All routers form an adjacency only with the BDR
- B. All routers form an adjacency only with the DR
- C. All routers form an adjacency with all other routers
- D. All routers form an adjacency with both the DR and BDR

Suggested Answer: D

Reference:

<https://sites.google.com/site/amitsciscozone/home/juniper-junos/junos-ospf-designated-router>

Community vote distribution

D (100%)

🗳️ 👤 **achon** 2 years, 3 months ago

D of course

upvoted 1 times

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

Selected Answer: D

D is correct

upvoted 1 times

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

Selected Answer: D

Correct is D


upvoted 1 times

🗳️ 👤 **t567** 3 years, 12 months ago

Answer:D

upvoted 4 times

Click the Exhibit button.



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

- A. The BGP peering between R1 and R4 should use loopback interface addresses
- B. The BGP peering between R1 and R4 should use physical interface addresses
- C. The BGP peerings between R1, R2, and R3 should use loopback interface addresses
- D. The BGP peerings between R1, R2, and R3 should use physical interface addresses

Suggested Answer: BC

Community vote distribution

BC (100%)

bartahr 3 years, 2 months ago

Selected Answer: BC

B, C are correct

upvoted 3 times

bartahr 3 years, 2 months ago

Selected Answer: BC

B and C are correct

upvoted 1 times

villa 3 years, 5 months ago

B and C correct

upvoted 2 times

Click the Exhibit button.

```
[edit protocols]
user@router# show
isis {
    interface ge-0/0/0.0;
}
```

Referring to the exhibit, which statement about the IS-IS interface is true?

- A. The ge-0/0/0.0 interface will act as an L1/L2 interface
- B. The ge-0/0/0.0 interface will act as an L2 interface only
- C. The ge-0/0/0.0 interface will act as an L1 interface only
- D. The ge-0/0/0.0 interface will not be assigned to a level

Suggested Answer: A

Community vote distribution

A (100%)

🗲️ 👤 **bartahr** 3 years, 2 months ago

Selected Answer: A

By default, interface in ISIS is in both levels
upvoted 1 times

🗲️ 👤 **FR99** 3 years, 4 months ago

The IS-IS level number, which can be 1 or 2. The default is for the device to operate as both a Level 1 and 2 device.
upvoted 1 times

🗲️ 👤 **majorcompugeek** 3 years, 7 months ago

Selected Answer: A

https://www.juniper.net/documentation/us/en/software/junos/isis/topics/ref/statement/level-edit-protocols-isis-interface.html#id-11257325__d30766e133
upvoted 3 times

Click the Exhibit button.

```
user@router> show route 10.100.110.1 hidden detail
```

```
inet.0: 33 destinations, 33 routes (22 active, 0 holddown, 11 hidden)
10.100.110.0/24 (1 entry, 0 announced)
  BGP Preference: 170/-101
    Next hop type: Unusable, Next hop index: 0
    Address: 0xc3ca334
    Next-hop reference count: 11
    State: <Hidden Int Ext>
    Local AS: 65514 Peer AS: 65514
    Age: 13
    Validation State: unverified
    Task: BGP_65514.192.168.0.2
    AS path: 65511 I
    Accepted
    Localpref: 100
    Router ID: 192.168.0.2
```


Referring to the exhibit, why is the route hidden?

- A. The wrong BGP address family is enabled for the BGP session
- B. The route has yet to be verified
- C. The protocol next hop is not reachable
- D. The MPLS LSP to the 192.168.0.2 peer is down

Suggested Answer: C

Community vote distribution

C (100%)

 **bartahr** 3 years, 2 months ago

Selected Answer: C

Next hope is not reachable

upvoted 1 times

Click the Exhibit button.

```
[edit]
user@r1# show protocols mpls
no-cspf;
label-switched-path r1-to-r3 {
    to 192.168.100.1;
    bandwidth 500m;
}
```

Referring to the exhibit, which statement is true?

- A. The router will attempt to signal the LSP along the IGP shortest path to 192.168.100.1
- B. The router will prune links with insufficient bandwidth from the path before beginning the signaling process
- C. The router will analyze the traffic engineering database to determine the best path through the network
- D. The router will precalculate a valid path through the network for LSP r1-to-r3

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ 👤 **Juniperguy** Highly Voted 👍 4 years, 2 months ago

As constrained shortest path first is disabled, the pruning of links bandwidth won't happen right? The answer is A? The router will attempt to signal the LSP along the IGP shortest path to 192.168.100.1

upvoted 10 times

🗳️ 👤 **TT98** Highly Voted 👍 4 years, 1 month ago

Because CSPF is disabled, visibility of overall topology is no longer available. LSP signaling will happen along IGP SP. So, correct answer is A.

upvoted 7 times

🗳️ 👤 **FR99** Most Recent 🕒 3 years, 1 month ago

Selected Answer: A

Answer is A. With CSPF disabled the Router will use just the IGP path to establish the LSP using RSVP signaling.

upvoted 1 times

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

Selected Answer: A

Since no-cspf, router will use IGP

upvoted 1 times

🗳️ 👤 **majorcompugeek** 3 years, 7 months ago

Selected Answer: A

Due to `no-cspf`, the route will follow IGP to 192.168.100.1

upvoted 2 times

🗳️ 👤 **GJATANKL** 3 years, 9 months ago

A. CSPF is disabled hence no reference to the TE DB.

upvoted 2 times

🗳️ 👤 **Dibyam** 3 years, 12 months ago

A. The router will attempt to signal the LSP along the IGP shortest path to 192.168.100.1

is correct, because of no-cspf command. CSFP is disabled meaning no mpls visibility to prune insufficient bandwidth.

upvoted 4 times

Which two statements are true for GRE tunneling? (Choose two.)

- A. GRE tunnel endpoints must have a valid route to the remote endpoint
- B. GRE tunnels support multiple logical units per interface
- C. GRE tunnels are stateful by default
- D. GRE tunnels support only one logical unit per interface

Suggested Answer: AB

Reference:

https://www.juniper.net/documentation/en_US/junos/topics/topic-map/filtering-unicast-packets-multicast-tunnel-interfaces.html#id-configuring-unicast-tunnels

Community vote distribution

AB (100%)

FR99 3 years ago

Selected Answer: AB

AB makes sense

upvoted 2 times

nushadu 3 years, 8 months ago

```
ed@vMX_R1# run show interfaces gre terse
```

```
Interface Admin Link Proto Local Remote
```

```
gre up up
```

```
gre.0 up up inet6 2001:db8:1::1/126
```

```
fe80::200:ff:fc00:0/64
```

```
gre.3 up up inet 192.168.13.1/30
```

```
inet6 fc00:3333:3333::1/126
```


```
fe80::200:ff:fc00:0/64
```

[edit]

```
ed@vMX_R1#
```

upvoted 3 times

Click the Exhibit button.



Referring to the exhibit, the GRE tunnel between R-1 and R-2 allows connectivity between User-1 and Server-1. When User-1 communicates with Server-2 with packets that are 1472 bytes in size, no packet fragmentation occurs. User-1 can communicate with Server-1 with packets that are up to 1448 bytes in size with no packet fragmentation. However, if the packet size is larger than 1448 bytes, packet fragmentation occurs.

Why is the packet fragmentation occurring between User-1 and Server-1 in this scenario?

- A. The GRE header adds 20 bytes to the packet
- B. The GRE header adds 24 bytes to the packet
- C. The IP header adds 20 bytes to the packet
- D. The IP header adds 24 bytes to the packet

Suggested Answer: B

Community vote distribution

B (100%)

FR99 3 years ago

Selected Answer: B

B is correct

upvoted 1 times

bartahr 3 years, 2 months ago

Selected Answer: B

GRE header is 24 bytes

upvoted 1 times

What happens when a packet matches a static route with the next hop parameter set to reject?

- A. The system silently drops the packet
- B. An ICMP message is sent to the source and the packet is forwarded
- C. An ICMP message is sent to the source and the packet is dropped
- D. The packet is forwarded and the packet is marked as rejected in the header

Suggested Answer: C

Reference:

<https://www.informit.com/articles/article.aspx?p=30666&seqNum=5>

Community vote distribution

C (100%)

🗨️ 👤 **FR99** 3 years, 1 month ago

Selected Answer: C

Answer is C.

'Reject' sends the ICMP destination unreachable msg.

'Discard' silently drops the packet.

Discard—Discard a packet silently, without sending an ICMP message.

Reject—Discard a packet, and send an ICMP destination unreachable message.

upvoted 1 times

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

Selected Answer: C

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

upvoted 1 times

Which two statements are correct about the BGP MED attribute? (Choose two.)

- A. BGP uses the MED value when peering to two or more connections to the same upstream AS
- B. BGP routes require the MED attribute be defined
- C. BGP uses the MED value when peering to two different upstream ASs
- D. BGP assumes the MED value to be 0, if not already defined



Suggested Answer: AD

Reference:

https://www.juniper.net/documentation/en_US/junos/topics/topic-map/med-attribute.html

Community vote distribution

AD (100%)

  **bartahr** 3 years, 2 months ago

Selected Answer: AD

- MED is a non-transitive attribute and optional.
- When you do not explicitly configure a metric value, the MED value is equivalent to 0
- When the routing table contains two routes to the same host in a neighboring AS, a MED metric assigned to each route can determine which to include in the forwarding table.

upvoted 1 times

What is the Junos default router priority advertisement value for IS-IS?

- A. 64
- B. 32
- C. 0
- D. 127

Suggested Answer: A

Reference:

https://www.juniper.net/documentation/en_US/junos/topics/concept/routing-protocol-is-is-security-designated-router-understanding.html#:~:text=If%20routers%20in%20the%20network,a%20priority%20value%20of%2064

Community vote distribution

A (100%)

FR99 3 years, 1 month ago

Selected Answer: A

Correct answer is A

upvoted 1 times

bartahr 3 years, 2 months ago

Selected Answer: A

A router's priority is from 0 through 127. By default, routers have a priority value of 64.

upvoted 1 times

nushadu 3 years, 8 months ago

ed@vMX_R2# show protocols isis

interface em0.2;

interface lo0.0;

[edit]

ed@vMX_R2# run show isis interface extensive em0.2

IS-IS interface database:

em0.2

Index: 64, State: 0x6, Circuit id: 0x2, Circuit type: 3

LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding, IIH max size: 1492

Adjacency advertisement: Advertise, Layer2-map: Disabled

Interface Group Holddown Delay: 20 s, remaining: 0 s

Level 1

Adjacencies: 1, Priority: 64, Metric: 10

Hello Interval: 3.000 s, Hold Time: 9 s

Designated Router: vMX_R2.02 (us)

Level 2

Adjacencies: 1, Priority: 64, Metric: 10

Hello Interval: 3.000 s, Hold Time: 9 s

Designated Router: vMX_R2.02 (us)

[edit]

ed@vMX_R2#

upvoted 2 times

Which three attributes are well-known mandatory BGP attributes? (Choose three.)

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

Suggested Answer: ABE

Reference:

<https://www.kwtrain.com/blog/bgp-pt2>

Community vote distribution

ABE (100%)

🗲️ 👤 **FR99** 3 years, 1 month ago

Selected Answer: ABE

ABE correct answer

upvoted 1 times

🗲️ 👤 **bartahr** 3 years, 2 months ago

Selected Answer: ABE

- Local preference is well-known discretionary

- MED is optional nontransitive

upvoted 1 times

Click the Exhibit button.

```
[edit protocols ospf]
user@router# show
reference-bandwidth 100m
area 0.0.0.0 {
    interface ge-1/0/0.0 {
        interface-type p2p;
    }
    interface ge-3/0/0.0 {
        priority 128;
    }
    interface xe-0/0/0.0 {
        interface-type nbma;
    }
}
```

Referring to the exhibit, which statement is correct?

- A. Interface ge-3/0/0.0 has a default metric of 10
- B. Interface xe-0/0/0.0 can only form a single adjacency
- C. Interface xe-0/0/0.0 has a default metric of 10
- D. Interface ge-1/0/0.0 can only form a single adjacency

Suggested Answer: D

Community vote distribution

D (100%)

FR99 3 years, 1 month ago

Selected Answer: D

D is correct

upvoted 1 times

bartahr 3 years, 2 months ago

Selected Answer: D


P2P links can have only one neighbor

upvoted 1 times

Click the Exhibit button.

```
[edit]
user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet6 {
      address 2001:0:0:1::2/64;
    }
  }
}
gr-0/0/0 {
  unit 0 {
    tunnel {
      source 192.168.1.1;
      destination 192.168.1.2;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.100.1/24;
    }
  }
}
fxp0 {
  unit 0 {
    family inet {
      address 10.0.1.12/24;
    }
  }
}
```

```
[edit]
user@R1# show routing-options
rib inet6.0 {
  static {
    route 2001:0:0:2::0/64 next-hop gr-0/0/0.0;
  }
}
static {
  route 0.0.0.0/0 next-hop 10.0.1.1;
  route 192.168.1.2/32 next-hop 192.168.100.2;
}
```



You have configured IPv6 over IPv4 tunneling, as shown in the exhibit. However, hosts connected to network 2001:0:0:1::0/64 cannot communicate with hosts on network 2001:0:0:2::0/64. The router R2 has a similar configuration as the R1 router.

How would you solve this problem?

- A. Configure an IGP across the tunnel interfaces
- B. Configure an IPv6 address on the tunnel interfaces
- C. Configure the next hop of the inet6.0 static route to point to the physical interface between the routers
- D. Configure the next hop of the inet6.0 static route to point to the IPv4 address of the remote router

Suggested Answer: B

Community vote distribution

B (100%)

Joohn2022 2 years, 4 months ago

The ipv6 addresses are in different subnets, shouldnt it be A?

upvoted 1 times

FR99 3 years, 1 month ago

Selected Answer: B

I think B is the correct answer

upvoted 1 times

Ragnahardt 3 years, 6 months ago

B is the correct answer.

Everything else is correctly configured.

IMPORTANT NOTE: In real scenarios, it is not strictly needed to configure an IPv6. What would be totally needed is to define IPv6 support under the 'gr' interface, as you would do for IPv4 support.

Example:

```
gr-0/0/0 {
  unit 0 {
    tunnel {
```

```
source 192.168.1.2;
destination 192.168.1.1;
}
family inet; ← ← ← ← ←
family inet6; ← ← ← ← ←
}
```

upvoted 4 times

🗲️ 👤 **[Removed]** 3 years, 9 months ago

first the ip addressing is wrong , for ipv6 over ipv4 gre the tunnels do not need ipv6 address as the packet is encapsulated , only thing needed is a route to put the packet into the tunnel , as the tunnel is a point-to-point connection is needed a route to the remote router

upvoted 1 times

🗲️ 👤 **mixeren** 3 years, 10 months ago

Seems like coorrect answer is B:

<https://www.juniper.net/documentation/us/en/software/junos/interfaces-encryption/topics/topic-map/configuring-ip-tunnel-interfaces.html>

upvoted 3 times

🗲️ 👤 **kamalelsherif** 3 years, 10 months ago

IPv6 must be configured on the tunnel interface - correct answer is B

upvoted 3 times

🗲️ 👤 **Chaimaaaaaaaa** 4 years ago

I guess D is the correct answer

upvoted 3 times

🗲️ 👤 **dodds** 4 years ago

Not B?

upvoted 1 times

In a stateless IPv6 auto-configuration scenario, what is the host's IPv6 address if the interface's MAC address is 12:34:ab:cd:ef:56?

- A. fe80::1234:abff:fe:cd:ef56/64
- B. fe8::1234:abff:fe:cd:ef56/64
- C. fec0::1234:abff:fe:cd:ef56/64
- D. fe80::1234:abcd:ef56/64

Suggested Answer: A

🗨️ **nushadu** 3 years, 8 months ago
 ed@vMX_R1# run show interfaces terse em2
 Interface Admin Link Proto Local Remote
 em2 up down

[edit]
 ed@vMX_R1# run show interfaces em2 | match HARD
 Interface flags: Hardware-Down SNMP-Traps
 Current address: 0c:9e:24:d6:00:02, Hardware address: 0c:9e:24:d6:00:02

[edit]
 ed@vMX_R1# set interfaces em2.0 family inet6

[edit]
 ed@vMX_R1# commit
 commit complete

[edit]
 ed@vMX_R1# run show interfaces terse em2
 Interface Admin Link Proto Local Remote
 em2 up down
 em2.0 up down inet6 fe80::e9e:24ff:fed6:2/64

[edit]
 ed@vMX_R1#
 upvoted 1 times

🗨️ **kkched** 3 years, 10 months ago
 Using SLAC scenario for configuring 64 bit Interface ID part of IP address
 Subnet ID = 12:34:ab:fffe:fe:cd:ef:56 (inserting FFFE between first 24 bit Company ID and last 24 bit Manufacturer ID)
 then add subnet ID part of Link local scope which is FE80::/10
 We then find : fe80::1234:abff:fe:cd:ef56
 upvoted 2 times

🗨️ **LittleMing** 3 years, 11 months ago
 I think correct answer is A (the question is host's IPv6 EUI-64)
 Ref : <https://blog.marquis.co/ipv6-and-junos-stateless-address-autoconfiguration-slaac/>
 upvoted 4 times

🗨️ **FrankZane** 3 years, 11 months ago
 Not D?
 upvoted 1 times

🗨️ **majorcompugeek** 3 years, 7 months ago
 No because D just appends the MAC address to fe80::<mac>/64. The right answer is derived using the follow EUI-64 conversion process:

1. Identify the 24 bit company ID and manufacture ID.

MAC: 12-34-ab-cd-ef-56

Company ID (First 24 bits/First half of MAC): 12-34-ab

Manufacture ID (Last 24 bits/Last half of MAC): cd-ef-56

2. Rewrite/Concatenate the new 64 bit address as follows <company id>-ff-fe-<manufacture id>: 12-34-ab-ff-fe-cd-ef-56

3. Drop the dashes and group into 4 byte groups separated by a colon (xxxx:xxxx:xxxx:xxxx). You should get 4 groups: 64 bit Address:

1234:abff:fe56

4. Now new address into the following to generate the new IPv6 SLAAC address: fe80::<new 64 bit address>/64 -> fe80::1234:abff:fe56/64
upvoted 3 times

Click the Exhibit button.

```
[edit routing-options]
user@router# show
generate {
    defaults {
        preference 5;
    }
    route 0.0.0.0/0 policy ISP-NET;
}

[edit]
user@router# show policy-options
policy-statement ISP-NET {
    term 1 {
        from protocol bgp;
        then accept;
    }
    term 2 {
        then reject;
    }
}
```

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

- A. The router will install the 0.0.0.0/0 route into the routing table when no BGP prefixes are present
- B. The router will remove the 0.0.0.0/0 route from the routing table when no BGP prefixes are present
- C. The router will remove the 0.0.0.0/0 route from the routing table when any BGP prefixes are present
- D. The router will install the 0.0.0.0/0 route into the routing table when any BGP prefixes are present

Suggested Answer: BD

Community vote distribution

BD (100%)

 **bartahr** 3 years, 2 months ago

Selected Answer: BD

B and D are correct

upvoted 1 times

Click the Exhibit button.

```
[edit protocols mpls]
user@router# show
label-switched-path R1-to-R6 {
    to 172.17.20.6;
    install 10.3.0.0/24 active;
}

[edit routing-options]
user@router# show
static {
    route 10.3.0.0/24 {
        lsp-next-hop R1-to-R6;
    }
}
```

Both configuration hierarchies shown in the exhibit have been committed to your MX Series device.


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

- A. Traffic destined to 10.3.0.1 will use the R1-to-R6 LSP as a next hop
- B. Traffic destined to 10.3.0.1 will not use the R1-to-R6 LSP as a next hop
- C. The active 10.3.0.0/24 prefix installed in the route table will have a route preference of 5
- D. The active 10.3.0.0/24 prefix installed in the route table will have a route preference of 7

Suggested Answer: AD

Community vote distribution

AC (100%)

 **bartahr** 3 years, 2 months ago

Selected Answer: AC

A -> route will use because static route is using it as next next hop

C -> RSVP route preference is 7, comparing to static route preference 5 -> it will use static route
upvoted 3 times

 **Panadol** 3 years, 4 months ago

Answer is A and C

upvoted 1 times

 **mahmutaydin** 3 years, 7 months ago

Answer A C

root@AT# run show route 10.3.0.0

inet.0: 30 destinations, 36 routes (30 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both


10.3.0.0/24 *[Static/5] 09:21:27

> to 198.51.100.1 via ge-0/0/1.0, label-switched-path AT-to-DA

[RSVP/7/1] 00:00:15, metric 11

> to 198.51.100.1 via ge-0/0/1.0, label-switched-path AT-to-DA

upvoted 2 times

 **minmon_6789** 3 years, 9 months ago

A & C are correct.

[edit]

root@R1# show routing-options

static {

```
route 10.0.0.3/32 {  
  lsp-next-hop TO_R3;  
}  
}  
router-id 10.0.0.1;
```

[edit]

root@R1# run show route 10.0.0.3

inet.0: 9 destinations, 10 routes (9 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

10.0.0.3/32 *[Static/5] 00:05:12

> to 172.16.0.2 via ge-0/0/1.0, label-switched-path TO_R3

Static route with next-hop is an LSP has the preference of 5.

upvoted 3 times



  **fakrulalam** 3 years, 9 months ago

The answer should be A & C:

MPLS Default preference: 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs

<https://www.juniper.net/documentation/us/en/software/junos/mpls/topics/ref/statement/preference-edit-protocols-mpls.html>



upvoted 1 times

  **Dibyam** 3 years, 12 months ago

StaticLSP preference is 6. So the best answer should be (The active 10.3.0.0/24 prefix installed in the route table will have a route preference of 6).

Hence, A and C is correct answer.

upvoted 1 times

  **TT98** 4 years, 1 month ago

Answer should be AC. install <dst> active will be discovered via RSVP with pref 7 while static route <dst> lsp-next-hop R1-to-R6 will be installed as a static route with pref 5.

upvoted 4 times

  **kent2612** 3 years, 9 months ago

since it is 'install 10.3.0.0/24 active' shouldn't the answer be D?

upvoted 1 times

  **cosminaas** 3 years, 6 months ago

Active knob just tell us that the route learned from RSVP should be installed in inet.0 rather than inet.3, but the preference remain 7 and therefore the route learned via static which has preference of 5 will be selected for forwarding.

upvoted 2 times

You are asked to change the default TTL handling behavior on your Junos device to ensure that the RSVP-signaled LSPs in your MPLS network cannot be mapped.

Which configuration should be performed to accomplish this task?

- A. Configure the no-decrement-ttl parameter for each LSP on the ingress device
- B. Configure the no-propagate-ttl parameter for each LSP on the egress device
- C. Configure the no-propagate-ttl parameter for each LSP on the ingress device
- D. Configure the no-decrement-ttl parameter for each LSP on the egress device

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ 👤 **val4y_** 2 years, 2 months ago

Selected Answer: A

On the ingress of the LSP, if you include the no-decrement-ttl statement, the ingress router negotiates with all downstream routers using a proprietary RSVP object, to ensure all routers are in agreement. If negotiation succeeds, the whole LSP behaves as one hop to transit IP traffic.

upvoted 2 times

🗳️ 👤 **val4y_** 2 years, 2 months ago

Correct answer is A

On the ingress of the LSP, if you include the no-decrement-ttl statement, the ingress router negotiates with all downstream routers using a proprietary RSVP object, to ensure all routers are in agreement. If negotiation succeeds, the whole LSP behaves as one hop to transit IP traffic.

no-decrement-ttl;

upvoted 1 times

🗳️ 👤 **achon** 2 years, 3 months ago

Selected Answer: A

it is ingress LSR, so A or C

"the default TTL handling behavior on your Junos device", so, A

upvoted 2 times

🗳️ 👤 **bartahr** 3 years ago

Selected Answer: A

It should be A

- no-decrement-ttl - ingress LSR use special Label Request object to signal all the routers in LSP (configured only on ingress LSR)

- no-propagate-ttl - each router in the path has to be configured with no-propagate-ttl

upvoted 2 times

🗳️ 👤 **FR99** 3 years ago

Selected Answer: A

After reading all comments here, the answer should be A

upvoted 2 times

🗳️ 👤 **Ragnahardt** 3 years, 6 months ago

A is the correct answer.

'no-decrement-ttl':

It is Juniper proprietary. Third party devices may not recognize this feature.

It is used at the ingress LSR only.

It can be globally configured at MPLS protocol definition or per LSP.

'no-propagate-ttl':

It is not Juniper proprietary. It exists as an alternative for third party compatibility.

It must be used on all LSR in the LSP.

It can only be configured at global MPLS configuration.

Due above characteristics:

Options B and D are immediately discarded. Either 'no-decrement-ttl' or 'no-propagate-ttl' can be configured at the ingress LSR but not at egress LSR ('no-propagate-ttl' can but must be configured on all the path, ingress LSR included).

Option C is discarded since 'no-propagate-ttl' cannot be configured per LSP. Is globally configured at the MPLS protocol.

upvoted 3 times

🗨️ 👤 **mahmutaydin** 3 years, 7 months ago

Answer is C

On the ingress router, you can include the no-propagate-ttl statement. The no-propagate-ttl statement applies to all LSPs, regardless of whether they are RSVP-signaled or LDP-signaled. Once set, all future LSPs traversing through this router behave as a single hop to IP packets. LSPs established before you configure this statement are not affected.

upvoted 2 times

🗨️ 👤 **minmon_6789** 3 years, 9 months ago

"The operation of the no-propagate-ttl statement is interoperable with other vendors' equipment.

However, you must ensure that all routers are configured identically."

"no-propagate-ttl" have to be configured on all routers along LSP, not only the ingress.

A is correct.

upvoted 1 times

🗨️ 👤 **GJATANKL** 3 years, 9 months ago

Answer is A. "no-decrement-ttl" on the ingress router and "no-propagate-ttl" on all MPLS routers

upvoted 2 times

🗨️ 👤 **bokko123** 4 years ago

The answer should be A as "Configure the no-decrement-ttl parameter for each LSP on the ingress device"

The configuration of the "no-propagate-ttl" have to be set on the ingress router including all LSPs.

https://www.juniper.net/documentation/en_US/junos/topics/usage-guidelines/mppls-disabling-normal-ttl-decrementing.html

upvoted 2 times

🗨️ 👤 **TT98** 4 years, 1 month ago

Answer should be A. "no-propagate-ttl" parameter needs to be defined on every LSR to have effect, so option B cannot be answer.

upvoted 3 times

Which action would you use to connect two virtual switches that are configured on the same router?

- A. Create a VRF routing instance
- B. Create a forwarding routing instance
- C. Connect the virtual switches with a cable
- D. Configure an irb interface

Suggested Answer: B

Community vote distribution

C (100%)

🗳️ **giorgi2023** 1 year, 9 months ago

Selected Answer: C

Only with Cable
upvoted 1 times

🗳️ **bartahr** 2 years, 11 months ago

Selected Answer: C

Only with Cable, Virtual Switches uses the same MAC
upvoted 1 times

🗳️ **FR99** 3 years ago

Selected Answer: C

Answer is C, with a cable
upvoted 1 times

🗳️ **Ragnahardt** 3 years, 6 months ago

Selected Answer: C

C is the correct answer:

"

For virtual routers, you can accomplish this task using either a logical tunnel interface or by looping two interfaces together with a single cable. For virtual switches, this process works only using the external cable method.

"

upvoted 3 times

🗳️ **mahmutaydin** 3 years, 7 months ago

C.

For virtual routers, you can accomplish this task using either a logical tunnel interface or

by looping two interface with a single cable.

For virtual switches, this works only using external cable. Reason is spanning tree

protocols do not function properly. Because all virtual switches use the same MAC

address as part of their bridge ID in BPDUs. Virtual switch's mac cannot be changed

upvoted 2 times

🗳️ **and00** 3 years, 11 months ago

At some point you might want to interconnect the virtual routers and virtual switches that are local to a single chassis. For virtual routers, you can accomplish this task using either a logical tunnel interface or by looping two interfaces together with a single cable. For virtual switches, this process works only using the external cable method.

upvoted 3 times

🗨️ 👤 **bokko123** 4 years ago

The answer should be "A"

routing instane has instance-type, Virtual switch.

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

upvoted 1 times

🗨️ 👤 **funkeymonkey** 3 years, 12 months ago

The question is on how to connect two virtual switch instances. There are two methods: for layer 3, configure irb interface; for layer 2, connect physically. I'd go C

upvoted 4 times

Which two functions are performed by the OSPF designated router? (Choose two.)

- A. It advertises link-state information to the AS
- B. It designates some routers as inactive when not needed
- C. It forms adjacencies with all the other OSPF routers on the link
- D. It chooses the backup designated router



Suggested Answer: AC

Reference:

<https://sites.google.com/site/amitsciscozone/home/juniper-junos/junos-ospf-designated-router>

Community vote distribution

AC (100%)

  **bartahr** 3 years, 2 months ago

Selected Answer: AC

A, C are correct

upvoted 1 times

The IPv6 Neighbor Discovery Protocol (NDP) performs the same function as which two IPv4 protocols? (Choose two.)

- A. ICMP
- B. ARP
- C. DNS
- D. DHCP



Suggested Answer: AB

Reference:

[https://www.juniper.net/documentation/en_US/junos/topics/topic-map/ipv6-neighbor-discovery.html#:~:text=Neighbor%20discovery%20for%20IPv6%20replaces,Discovery%20protocol%20\(NDP\)%20messages](https://www.juniper.net/documentation/en_US/junos/topics/topic-map/ipv6-neighbor-discovery.html#:~:text=Neighbor%20discovery%20for%20IPv6%20replaces,Discovery%20protocol%20(NDP)%20messages)

Community vote distribution

AB (100%)

  **achon** 2 years, 3 months ago

Selected Answer: AB

One of the key features of NDP is its ability to perform address resolution without the need for the Address Resolution Protocol (ARP) used in IPv4. NDP accomplishes this by using a combination of ICMPv6 messages and multicast address resolution.

upvoted 1 times

  **FR99** 3 years ago

Answer could be D instead of A due to SLAAC being used as a replacement for DHCP in some use cases.

upvoted 1 times

  **donkey_12** 3 years, 2 months ago

AB, As discovery protocol are ICMP and ARP

upvoted 1 times

  **donkey_12** 3 years, 2 months ago

Neighbor discovery for IPv6 replaces the following IPv4 protocols: router discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

[https://www.juniper.net/documentation/us/en/software/junos/neighbor-discovery/topics/topic-map/ipv6-neighbor-discovery.html#:~:text=Neighbor%20discovery%20for%20IPv6%20replaces,Discovery%20protocol%20\(NDP\)%20messages](https://www.juniper.net/documentation/us/en/software/junos/neighbor-discovery/topics/topic-map/ipv6-neighbor-discovery.html#:~:text=Neighbor%20discovery%20for%20IPv6%20replaces,Discovery%20protocol%20(NDP)%20messages)

upvoted 1 times

Which two statements are true about IP and GRE tunnels? (Choose two.)

- A. The protocol field is changed in the inner IP packet header
- B. Tunnel traffic is encrypted
- C. The TTL field is changed in the inner IP packet header
- D. Tunnel endpoints need a valid route to the remote endpoint

Suggested Answer: *CD*

🗨️ 👤 **Raul21** 2 years, 9 months ago

C&D

"The inner packet (also known as the payload packet) is not modified, except for the time-to-live (TTL) field, which is reduced"
upvoted 1 times

🗨️ 👤 **nushadu** 3 years, 8 months ago

C & D

RFC 2784 Generic Routing Encapsulation March 2000

3.1. Forwarding Decapsulated IPv4 Payload Packets

When a tunnel endpoint decapsulates a GRE packet which has an IPv4 packet as the payload, the destination address in the IPv4 payload packet header MUST be used to forward the packet and the TTL of the payload packet MUST be decremented.
upvoted 3 times

🗨️ 👤 **minmon_6789** 3 years, 9 months ago

TTL in the inner IP header is not changed.
upvoted 1 times

🗨️ 👤 **minmon_6789** 3 years, 9 months ago

I think in this question only D is correct.
Do you have any comments?
upvoted 3 times

Which two IP addresses are considered Martian addresses? (Choose two.)

- A. 0.0.0.0/8
- B. 192.168.0.0/8
- C. 240.0.0.0/4
- D. 169.254.0.0/16

Suggested Answer: AC

Reference:

https://www.juniper.net/documentation/en_US/junos/topics/topic-map/recognize-martian-addr-routing.html

Community vote distribution

AC (100%)

🗳️ 👤 **FR99** 3 years ago

Selected Answer: AC

AC are correct

upvoted 1 times

🗳️ 👤 **nushadu** 3 years, 8 months ago

ed@vMX_R2# run show route martians

inet.0:

0.0.0.0/0 exact -- allowed

0.0.0.0/8 orlonger -- disallowed

127.0.0.0/8 orlonger -- disallowed

192.0.0.0/24 orlonger -- disallowed

240.0.0.0/4 orlonger -- disallowed

224.0.0.0/4 exact -- disallowed

224.0.0.0/24 exact -- disallowed

inet.1:

upvoted 1 times

🗳️ 👤 **donkey_12** 3 years, 2 months ago

So correct answer is only A, is it not?? 240.0.0.0/4 is disallowed

upvoted 1 times

Click the Exhibit button.

```
[edit policy-options]
user@R1# show
policy-statement direct2ospf {
  term 1 {
    from {
      protocol direct;
      route-filter 172.10.1.0/24 exact;
    }
    then accept;
  }
}

[edit protocols]
user@R1# show
ospf {
  export direct2ospf;
  area 0.0.0.1 {
    interface ge-1/0/0.0;
  }
}

[edit protocols]
user@R2# show
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
    interface lo0.0;
  }
  area 0.0.0.1 {
    interface ge-1/0/0.0;
  }
}
```

Referring to the exhibit, which statement is correct?

- A. R2 is an ASBR
- B. R1 is a backbone router
- C. R2 is an ABR
- D. R1 is an ABR

Suggested Answer: C

Community vote distribution

C (100%)

FR99 3 years, 1 month ago

Selected Answer: C

C is correct as R2 has configured an interface on OSPF area 0 and another interface on OSPF area 1
upvoted 2 times

Click the Exhibit button.

```
[edit protocols ospf]
user@router# show
reference-bandwidth 10g;
area 0.0.0.0 {
    interface ge-1/0/0.0 {
        priority 255;
    }
    interface ge-3/0/0.0 {
        priority 128;
    }
    interface xe-0/0/0.0 {
        interface-type nbma;
    }
}
```

Referring to the exhibit, which statement is correct?

- A. Interface xe-0/0/0.0 has a default metric of 10
- B. Interface ge-3/0/0.0 has a default metric of 10
- C. Interface ge-1/0/0.0 can only form a single adjacency
- D. Interface xe-0/0/0.0 can only form a single adjacency

Suggested Answer: B

Community vote distribution

B (100%)

FR99 3 years, 1 month ago

Selected Answer: B

With 10Gig of OSPF reference bandwidth configured, a 1Gig interface would have an OSPF cost of 10
upvoted 1 times

nushadu 3 years, 8 months ago

ed@vMX_R1# run show interfaces em0 extensive | match speed
Speed: 1000mbps

[edit]

ed@vMX_R1# run show ospf interface em0.2 extensive | match cost
Type: P2P, Address: 10.0.0.1, Mask: 255.255.255.0, MTU: 1500, Cost: 10
Topology default (ID 0) -> Cost: 10

[edit]

ed@vMX_R1# show protocols ospf | display set
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface lo0.0 metric 500
set protocols ospf area 0.0.0.0 interface em0.2 interface-type p2p
set protocols ospf area 0.0.0.0 interface em0.2 bfd-liveness-detection minimum-interval 1000
set protocols ospf reference-bandwidth 10g

[edit]

ed@vMX_R1#
upvoted 2 times

mixeren 3 years, 10 months ago

Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:
cost = ref-bandwidth/bandwidth
upvoted 3 times

🗨️ 👤 **shaokan1** 3 years, 10 months ago

reference-bandwidth of 10g means a 1g interface will have a metric of 10.

upvoted 4 times

🗨️ 👤 **Dlbyam** 3 years, 12 months ago

For NBMA OSPF network-type, we can form multiple adjacency with static neighbor relationship, so not true. Also, reference-BW/interface-BW provides ospf cost. So, B is true answer.

upvoted 4 times

🗨️ 👤 **DirkBaert** 4 years ago

nmda = non broadcast, so I think D

upvoted 3 times

🗨️ 👤 **fakrulalam** 3 years, 9 months ago

<https://ipccisco.com/lesson/ospf-network-types-ccnp/>

OSPF DR and BDR selection is done for Nonbroadcast Multi Access(NBMA) Networks.

upvoted 1 times

Click the Exhibit button.

```
[edit protocols ospf area 0.0.0.0]
user@router# show
interface ge-0/0/0.0 {
    bfd-liveness-detection {
        minimum-interval 500;
    }
}
```

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

- A. The OSPF neighbor will be declared down if BFD hello packets are not received for 1.5 seconds
- B. The OSPF neighbor will be declared down if BFD hello packets are not received for 5 seconds
- C. The OSPF neighbor will be declared down if 500 BFD hello packets are missed
- D. The OSPF neighbor will be declared down if three BFD hello packets are missed

Suggested Answer: AD

Community vote distribution

AD (100%)

 **DirkBaert** Highly Voted 4 years ago

AD: correct

upvoted 11 times

 **FR99** Most Recent 3 years, 1 month ago

Selected Answer: AD

500ms and by default 3 BFD keepalives should be missed for the session to go down

upvoted 1 times

 **nushadu** 3 years, 8 months ago

ed@vMX_R1# run show bfd session extensive

Detect Transmit

Address State Interface Time Interval Multiplier

10.0.0.2 Up em0.2 3.000 1.000 3

Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 1.000, RX interval 1.000

Session up time 03:04:09, previous down time 00:00:02

Local diagnostic None, remote diagnostic None

Remote state Up, version 1

Session type: Single hop BFD

Min async interval 1.000, min slow interval 1.000

Adaptive async TX interval 1.000, RX interval 1.000

Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3

Remote min TX interval 1.000, min RX interval 1.000, multiplier 3

Local discriminator 16, remote discriminator 16

Echo TX interval 0.000, echo detection interval 0.000

Echo mode disabled/inactive Session ID: 0x0

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

[edit]



ed@vMX_R1#

upvoted 1 times

 **LittleMing** 3 years, 11 months ago

$500\text{ms} * 3 = 1.5 \text{ sec}$

upvoted 3 times

  **virtualwalker** 3 years, 12 months ago

Default Multiplier: 3

upvoted 3 times

Click the Exhibit button.

```
[edit interfaces]
user@router# show
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.1.5/31;
        }
        family mpls;
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.1.21/31;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.2/32;
        }
    }
}

[edit protocols bgp group BGP]
user@router# show
multihop;
local-address 192.168.0.2;
hold-time 30;
family inet {
    unicast;
}
family inet-vpn {
    unicast;
}
family inet6 {
    unicast;
}
```

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

- A. The configuration is for an external BGP session
- B. The local-address statement is required for the BGP session to establish correctly
- C. The multi-hop statement is required for the BGP session to establish correctly
- D. The configuration is for an internal BGP session

Suggested Answer: BD

Community vote distribution

AC (50%)

BD (50%)

 **java70** Highly Voted 3 years, 6 months ago

The parts you can't see in the picture will be as follows.

```
peer-as 65514;
local-as 65514;
neighbor 192.168.0.1;
upvoted 7 times
```

 **val4y_** Most Recent 2 years, 2 months ago

Selected Answer: BD

See what's missing from "java70", B and D are correct.

In most cases, the local-address statement is not required for the BGP session to establish correctly, as the router will select the appropriate source IP address based on its routing table and configured interfaces. However, there are some scenarios where specifying a specific local IP address may

be necessary, such as when there are multiple interfaces with the same IP address or when using a loopback interface as the source address for BGP sessions.

Therefore, while it is not required, configuring the local-address statement can provide additional control over the BGP session's source IP address and may be necessary in some specific scenarios.

upvoted 1 times

🗳️ 👤 **bartahr** 2 years, 11 months ago

Selected Answer: AC

A - If there is no type default is external

C - if you are using loopback you need to use multi-hop in external BGP

upvoted 1 times

🗳️ 👤 **minmon_6789** 3 years, 9 months ago

"i.e. because it has "multihop" defined it must be an external bgp and there is other config outside the screenshot that shows a peer-as"

Hi Takigama,

I see that multihop can be config along with type internal, why you said: "because it has "multihop" defined it must be an external bgp"?

upvoted 1 times

🗳️ 👤 **Takigama** 3 years, 9 months ago

it feels like part of the question is missing... the configuration as it is written wouldnt commit as the DEFAULT is type external, if its type external peer-as must be configured... junos wont let you commit that config unless you either specify a peer as or a type... unless you dont define a neighbor, in which case, this config would commit and achieve nothing.. but without either "type internal" or a "peer as" its almost impossible to answer the question without actual assumptioms..

i.e. because it has "multihop" defined it must be an external bgp and there is other config outside the screenshot that shows a peer-as

upvoted 3 times

🗳️ 👤 **fakrulalam** 3 years, 9 months ago

The answer would be A & C

To commit the config either we need "peer-as" for eBGP session or "type internal" for iBGP session. As none of the config there; I assume it's iBGP session due to multihop command as the local-address is loopback. Question 28 will have some clue.

upvoted 1 times

🗳️ 👤 **[Removed]** 3 years, 9 months ago

The type external command is kinda optional , a bgp config can be set without specifiing type external , anyway the configuration shown cannot be committed because no peer-as was issued , unless no type internal is shown i can assume that an eBGP session is in place to be build , also the multihop statement is a good indicator

upvoted 1 times

🗳️ 👤 **shaokan1** 3 years, 10 months ago

the major difference between the iBGP and eBGP is the AS Number. in iBGP you don't need to add the AS number while in eBGP you must add. so B and D are correct.

upvoted 1 times

🗳️ 👤 **Dlbyam** 3 years, 10 months ago

If it is iBGP connection then "type internal" must be define, unless it is ebgp

upvoted 2 times

🗳️ 👤 **kamalelsherif** 3 years, 10 months ago

- By default the EBGp connection is established between immediately connected devices located in two different ASs because the time-to-live (TTL) value of the EBGp packets is equal to 1. If an EBGp peer is more than one hop away, a multihop BGP session must be established. Correct answers are A & C

upvoted 3 times

🗳️ 👤 **Dlbyam** 3 years, 10 months ago

Yes, A and C are correct. Because if it is ibgp then "type internal" must be defined, but isnot required in ebgp. Also, we can peer ebgp neighbor using lo0 interface as source-interface but then TTL is 1, so we need to define multihop. Thus, A and C are correct answer.

upvoted 3 times

🗳️ 👤 **Dlbyam** 3 years, 12 months ago

I think, B and D are correct because peering IP is loopback_address, which basically used in iBGP, and multihop statement isnot required because underlying IGP will provide reachability between the peering IP.

upvoted 2 times

🗨️ 👤 **Deb123** 3 years, 12 months ago

I believe A & B is correct because of the multihop statement in the configuration

And multihop does require local address to be specified

upvoted 4 times

🗨️ 👤 **DIbyam** 3 years, 12 months ago

Think question is good one, you need to think a little-bit more.

upvoted 2 times

What is the correct description of an Area Border Router (ABR)?

- A. An ABR is an OSPF router with links in two areas, connecting OSPF areas to the backbone
- B. An ABR is an OSPF router that injects routing information from outside the OSPF AS
- C. An ABR is an OSPF router with at least one link in a Layer 2 area
- D. An ABR is an OSPF router with all of its links within an area

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **FR99** 3 years, 1 month ago

Selected Answer: A

Definitely A is the correct answer as an OSPF ABR is by definition a Router with one of its links in Area 0 and at least another in a non-zero Area
upvoted 1 times

🗨️ 👤 **donkey_12** 3 years, 2 months ago

A.
Autonomous system boundary router (ASBR) or an area border router (ABR).
upvoted 1 times

Click the Exhibit button.

```
[edit protocols]
'bgp'
Error in neighbor 192.168.1.2 of group my-int-group:
peer AS number must be configured for an external peer
error: configuration check-out failed
```

You are configuring an IBGP group. When you commit your configuration, you receive the error shown in the exhibit. Which additional configuration parameter must you add to your configuration?

- A. multipath
- B. type external
- C. type internal
- D. export <policy name>

Suggested Answer: C

Community vote distribution

C (100%)

FR99 3 years, 1 month ago

Selected Answer: C

Answer is C, the command "type internal" is needed.

upvoted 1 times

donkey_12 3 years, 2 months ago

Question is "You are configuring an IBGP group." - so answer is C - type internal

upvoted 1 times

nushadu 3 years, 8 months ago

ed@vMX_R1# delete protocols bgp group eBGP peer-as

[edit]

ed@vMX_R1# show | compare

[edit protocols bgp group eBGP]

- peer-as 3;

[edit]

ed@vMX_R1# commit check

[edit protocols]

'bgp'

Error in neighbor 10.0.13.3 of group eBGP:

peer AS number must be configured for an external peer

error: configuration check-out failed

[edit]

ed@vMX_R1# rollback

load complete

[edit]

ed@vMX_R1# commit check

configuration check succeeds

[edit]

ed@vMX_R1# show | compare



[edit]

```
ed@vMX_R1# show protocols bgp group eBGP | display set
set protocols bgp group eBGP type external
set protocols bgp group eBGP family inet unicast rib-group inet0-to-test
set protocols bgp group eBGP export EXPORT_TO_eBGP
set protocols bgp group eBGP peer-as 3
set protocols bgp group eBGP neighbor 10.0.13.3
```

[edit]

ed@vMX_R1#

upvoted 2 times

  **bokko123** 4 years ago

A error messages tells us configuration is running as eBGP, so the Answer should be "C".

It needs to specify the config as iBGP.

upvoted 3 times

  **[Removed]** 3 years, 9 months ago

if no type command is issued junos ask for a peer-as statement , but here if you issue type internal the session will be treated as iBGP session , also group is called int-bgp

upvoted 1 times

Click the Exhibit button.

```
user@router> show bgp neighbor 10.1.254.1
Peer: 10.1.254.1 AS 100          Local: 10.1.254.2 AS 65000
  Type: External      State: Active      Flags: <>
  Last State: Idle     Last Event: Start
  Last Error: Open Message Error
  Export: [ ebgp-export ]
  Options: <Preference AddressFamily PeerAs Refresh>
  Address families configured: inet-unicast inet6-unicast
  Holdtime: 90 Preference: 170
  Number of flaps: 15
  Last flap event: RecvNotify
  Error: 'Open Message Error' Sent: 6 Recv: 0
  Error: 'Cease' Sent: 13 Recv: 2

user@router> show log messages | match "open message"
Sep 19 00:07:31 R1 rpd[1325]: bgp_pp_rcv:3124: NOTIFICATION sent to 10.1.254.1+52788 (proto):
code 2 (Open Message Error) subcode 2 (bad peer AS number), Reason: no group for
10.1.254.1+52788 (proto) from AS 1000 found (peer as mismatch), dropping him
...
```

You are troubleshooting a new BGP peering session which is not establishing.
Referring to the exhibit, which statement is true?

- A. The update messages contain an unsupported option
- B. The neighbor does not support IPv6
- C. The peer's AS number is misconfigured
- D. The TCP session is not establishing

Suggested Answer: C

Community vote distribution

C (100%)

FR99 3 years, 1 month ago


Selected Answer: C

C is the correct answer
upvoted 1 times

donkey_12 3 years, 2 months ago

show log said AS1000 is found, while show BGP AS is 100, therefore Peering AS is misconfig or mismatch. - Correct answer is C.
upvoted 1 times

Click the Exhibit button.



You must create a static route on Router-A to the 10.10.5.0/24 network using 10.10.3.2 as the next hop. Referring to the exhibit, which configuration accomplishes this task?

A.

```
{master:0}[edit]
user@Router-A# show routing-options
static {
    route 10.10.5.0/24 next-hop 10.10.3.2;
}
```

B.

```
user@Router-A# show routing-options
static {
    route 10.10.5.0/24 next-hop 10.10.1.2;
}
```

C.

```
{master:0}[edit]
user@Router-A# show routing-options
static {
    route 10.10.5.0/24 {
        next-hop 10.10.3.2;
        qualified-next-hop 10.10.1.2;
    }
}
```

D.

```
{master:0}[edit]
user@Router-A# show routing-options
static {
    route 10.10.5.0/24 {
        next-hop 10.10.3.2;
        resolve;
    }
}
```

Suggested Answer: D

FR99 3 years, 1 month ago

Correct answer is D

upvoted 1 times

donkey_12 3 years, 2 months ago

Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the inet.0 and inet.3 routing tables.

Correct Answer is D.

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

upvoted 1 times

Which two statements describe operations performed by the encapsulating tunnel endpoint in an IP-IP tunnel? (Choose two.)

- A. It decrements the time-to-live (TTL) counter by one in the inner IP header
- B. It modifies the source and destination addresses in the inner IP header
- C. It adds an outer IP header with the destination address of the remote tunnel endpoint
- D. It creates and adds a new inner IP header with the remote destination device's IP address

Suggested Answer: AC

Community vote distribution

AC (100%)

🗲️ 👤 **Ragnahardt** 3 years, 6 months ago

Selected Answer: AC

When the packet is encapsulated with GRE the original TTL is decremented to make sure that the packet does not live forever.
upvoted 3 times

🗲️ 👤 **crramos** 3 years, 8 months ago

Answer is C and D, ttl is not decremented by encapsulator/decapsulator
upvoted 1 times

🗲️ 👤 **Ragnahardt** 3 years, 6 months ago

This is not true.

When the packet is encapsulated with GRE the original TTL is decremented to make sure that the packet does not live forever.
upvoted 1 times

🗲️ 👤 **nushadu** 3 years, 8 months ago

+

IP in IP Encapsulation

To encapsulate an IP datagram using IP in IP encapsulation, an outer IP header [10] is inserted before the datagram's existing IP header, as follows:
upvoted 1 times

🗲️ 👤 **nushadu** 3 years, 8 months ago

rfc2003

When encapsulating a datagram, the TTL in the inner IP header is decremented by one if the tunneling is being done as part of forwarding the datagram; otherwise, the inner header TTL is not changed during encapsulation. If the resulting TTL in the inner IP header is 0, the datagram is discarded and an ICMP Time Exceeded message SHOULD be returned to the sender. An encapsulator MUST NOT encapsulate a datagram with TTL = 0.

The TTL in the inner IP header is not changed when decapsulating. If, after decapsulation, the inner datagram has TTL = 0, the decapsulator MUST discard the datagram. If, after decapsulation, the decapsulator forwards the datagram to one of its network interfaces, it will decrement the TTL as a result of doing normal IP forwarding. See also Section 4.4.
upvoted 2 times