



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

```
user@router> show ospf route 30.0.0.0/24
Topology default Route Table:
```

Prefix	Path Type	Route Type	NH Type	Metric	NextHop Interface	Nexthop Address/LSP
30.0.0.0/24	Ext2	Network	IP	0	ge-0/0/1.0	5.0.0.1

```
user@router> show route protocol ospf 30.0.0.0/24
```

```
inet.0: 21 destinations, 23 routes, (21 active, 0 holddown, 0 hidden)
```

You notice an inconsistency between the routing table and the OSPF database, as shown in the exhibit.

What are two reasons for this behavior? (Choose two.)

- A. The LSA is a Type 4 LSA.
- B. An OSPF export policy is being applied to the route.
- C. An OSPF import policy is being applied to the route.
- D. The LSA is a Type 5 LSA.

Suggested Answer: BC

Community vote distribution

CD (100%)

 **lam3n** Highly Voted 3 years, 2 months ago

OSPF import policy allows you to prevent external routes from being added to the routing tables of OSPF neighbors. The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements. The filtering is done only on external routes in OSPF.

C&D

upvoted 7 times

 **manafo** Most Recent 2 years ago

Selected Answer: CD

C and D are correct

upvoted 2 times

 **Dlbyam** 2 years, 5 months ago

C and D are correct

upvoted 1 times

 **EBNPW9** 2 years, 6 months ago

Selected Answer: CD

see answers by lam3n and minmon_6789

upvoted 3 times

 **oceans1908** 2 years, 6 months ago

C and D

upvoted 1 times

 **jmbecerra** 2 years, 8 months ago

Selected Answer: CD

C and D

upvoted 2 times

 **Dlbyam** 2 years, 9 months ago

C and D are correct. Import policy in RT and OSPF External type 2 route.

upvoted 1 times

 **Anniesyed1234** 2 years, 10 months ago

C & D are correct

upvoted 1 times

🗨️ 👤 **minmon_6789** 2 years, 11 months ago

C and D are correct.

upvoted 1 times

🗨️ 👤 **TT98** 2 years, 11 months ago

Should be C&D. Export policy lists the name of the routing policy to be evaluated when routes are being exported from the routing table into OSPF. This rules out B because with prefix absent in route table, OSPF LSDB could not be populated.

upvoted 1 times

🗨️ 👤 **ak_sisko** 3 years ago

C and D

upvoted 1 times

🗨️ 👤 **Juniperguy** 3 years, 1 month ago

C and D.

As this is an external route, LSA 5 is used for this. LSA 4 is used for discovering the ASBR.

Also, as the route is in the RIB in, but not in the routing table, this has to be done by a import policy.

upvoted 4 times

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

you can have import policy in ospf, it doesn't impact the LSDB, but the routing table and it works only for external routes, as in this case. It's definitely C and D

https://www.juniper.net/documentation/en_US/junos/topics/topic-map/ospf-routing-policy.html

upvoted 2 times

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

I think B and D. In OSPF only export policy

upvoted 1 times

🗨️ 👤 **minmon_6789** 2 years, 11 months ago

You can apply an import policy in OSPF. OSPF import policies apply to external routes only.

So C and D are correct.

upvoted 2 times

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

I'd say it's C and D

upvoted 2 times

```
(65001)R1-----R2-----R3(65001)

[edit]
user@R2# run show route 11.11.11.0/24

inet.0 : 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

11.11.11.0/24      *[BGP/170] 00:04:55, localpref 100
                  AS path: 65001 I, validation-state: unverified
                  > to 172.16.1.1 via ge-0/0/0.0
                  [BGP/170] 00:10:33, localpref 100
                  AS path: 65001 65001 I, validation-state: unverified

[edit]
user@R2# show protocols bgp
group R1 {
  neighbor 172.16.1.1 {
    peer-as 65001;
  }
}
group R3 {
  neighbor 172.16.2.1 {
    peer-as 65001;
  }
}
local-as 65002;

[edit]
user@R2# show policy-options
policy-statement lb {
  then {
    load-balance per-packet;
  }
}
policy-statement prepend {
  term 1 {
    then as-path-prepend 65001;
  }
}

[edit]
user@R2# show routing-options
forwarding-table {
  export lb;
}
```

R2 is receiving the same route from R1 and R3. You must ensure that you can load balance traffic for that route. Referring to the exhibit, which two configuration changes will allow load balancing? (Choose two.)

- A. Apply the prepend policy as an import policy under group R1.
- B. Configure multipath under the global BGP configuration.
- C. Configure multipath under group R1.
- D. Apply the prepend policy as an import policy under group R3.

Suggested Answer: AB

Community vote distribution

AB (100%)

 **mohdema** 1 year, 9 months ago

Selected Answer: AB

AS path must be equal

multipath should be applied globally so that it spans all the configurations

upvoted 1 times

 **BobbyAxelrod** 2 years, 4 months ago

Selected Answer: AB

A & B based on the comments

upvoted 2 times

🗨️ 👤 **amart6936** 2 years, 10 months ago

Starting in Junos OS Release 18.1R1 BGP multipath is supported globally at [edit protocols bgp] hierarchy level.

AS-Path must be equal so prepending R1 is needed to match the AS-Path values.

upvoted 2 times

🗨️ 👤 **minmon_6789** 2 years, 11 months ago

I can not config option multipath in BGP global hierarchy.

upvoted 3 times

🗨️ 👤 **hcccc** 3 years, 3 months ago

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/policy-prepend-as-numbers-to-bgp-as-paths.html>

upvoted 3 times

🗨️ 👤 **dix** 3 years, 4 months ago

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

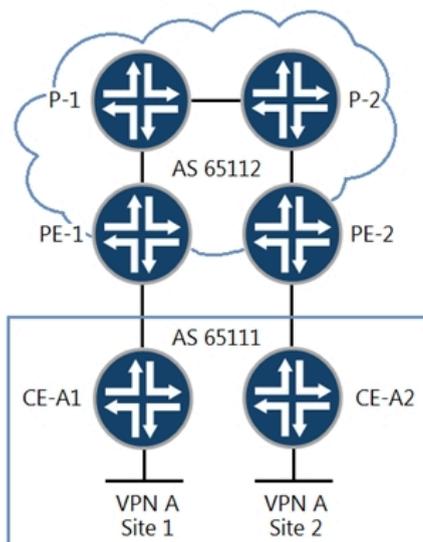
upvoted 1 times

🗨️ 👤 **yemicontrol** 3 years, 4 months ago

<https://www.juniper.net/documentation/us/en/software/junos/sampling-forwarding-monitoring/bgp/topics/concept/bgp-multipath-unequal-understanding.html>

Answer A&B is correct

upvoted 4 times



```
[edit routing-instances CE-A1]
user@PE-1# show
instance-type vrf;
interface ge-0/0/9.0;
route-distinguisher 10.222.222.3:2;
vrf-target target:65511:101;
protocols {
  bgp {
    group CE-A1 {
      type external;
      peer-as 65111;
      neighbor 192.168.0.2;
    }
  }
}

[edit routing-instances CE-A2]
user@PE-2# show
instance-type vrf;
interface ge-0/0/9.0;
route-distinguisher 10.222.222.3:2;
vrf-target target:65511:101;
protocols {
  bgp {
    group CE-A2 {
      type external;
      peer-as 65111;
      neighbor 192.168.6.2;
    }
  }
}
```

Referring to the exhibit, hosts in Site 1 and Site 2 are unable to communicate with each other through the Layer 3 VPN. What is the problem?

- A. The two sites are in the same AS.
- B. The two sites are using the same instance type.
- C. The two sites are using the same route target.
- D. The two sites are using the same route distinguisher.

Suggested Answer: A

Juniperguy Highly Voted 3 years, 1 month ago

As both routers are in the same AS: 65111, the PE wont export the route from the other CE, because it sees itself in the path. To overcome this as-override, has to be used, or put the customer in different AS.

So answer A

upvoted 14 times

NinjaCloud Most Recent 2 years, 4 months ago

In case the PE see same ASN in the ASPATH the route is rejected to avoid any loop.

The loop prevention mechanisms that PE can use:

- as-override
- advertise-peer-as
- remove-private

upvoted 3 times

kasqureshi 2 years, 4 months ago

A. The AS is same at both ends so, CE will reject the route seeing its own AS in the BGP update.

upvoted 1 times

DirkBaert 2 years, 11 months ago

Correct = D, RD must be different on each PE

upvoted 2 times

kasqureshi 2 years, 4 months ago

RDs can be same for a vpn instance. A is the answer.

upvoted 1 times

  **dodds** 2 years, 8 months ago

I think the answer is A. RD can be the same on each PE. It must be different for each vrf.

upvoted 2 times

```
[edit class-of-service]
user@router# show
classifiers {
  dscp classifierX {
    forwarding class low-priority {
      loss-priority low code-points 000000;
      loss-priority high code points 000001;
    }
    forwarding class medium-priority {
      loss-priority low code-points 000010;
      loss-priority high code points 000011;
    }
    forwarding class high-priority {
      loss-priority low code-points 000100;
      loss-priority high code points 000101;
    }
  }
}

forwarding-classes {
  class low-priority queue-num 0;
  class medium-priority queue-num 1;
  class high-priority queue-num 2;
  class network_control queue-num 3;
}
```

You manage an MX Series device which includes the configuration shown in the exhibit. Traffic marked with DSCP 000011 is entering the ge-1/0/4 interface at

102 Mbps. The traffic exits the device on the ge-1/0/5 interface. No other traffic is transiting the router.

In this scenario, what happens to traffic exceeding 100 Mbps?

- A. Traffic exceeding 100 Mbps is redirected to a rate limiter.
- B. Traffic exceeding 100 Mbps is buffered.
- C. Traffic exceeding 100 Mbps is dropped.
- D. Traffic exceeding 100 Mbps is forwarded.

Suggested Answer: C

Community vote distribution

D (100%)

 **dix**  3 years, 4 months ago

After you create a custom classifier you need to apply it to an interface within the class-of-service stanza. Other hand, the interface ge-1/0/5 is Gbps then the traffic is not exceeded normally; I think the information is uncompleted, however if we suppose there is a 100Mbps rate limit then the packets are classified based on the least significant bit of an incoming packet's CoS field, (1 -> high, 0 -> low), then the traffic exceeding 100Mbps should be dropped.

upvoted 8 times

 **sandpit**  2 years, 9 months ago

The information in question is not full, Interfaces & Scheduler-maps configurations are missing which is available on few other websites. From the schedulers config it shows "transmit-rate percent 10" for "med-pri-scheduler" which applies to ge-1/0/5 on exit interface.

Percentage shows 10% of 1GIG interfaces which drops any traffic above 100m on Ge-1/0/5. So I stick with the answer "Dropped"

upvoted 5 times

 **CptBlack** 10 months ago

Nope. "transmit-rate" by itself only ****guarantees**** a certain amount of bandwidth under ****congestion****. The question clearly states that there is NO congestion by the way. Unless you add "rate-limit" parameter, the "transmit-rate" DOES NOT by itself police or discard traffic. So if

no "rate-limit" is specified, and given that there is no congestion, all traffic will be forwarded out ge-1/0/5 (and "transmit-rate" in this case does nothing).

upvoted 1 times

  **ztw3587t** Most Recent 1 year, 3 months ago

Selected Answer: D

D is correct. It considering the available information we can assume that the default scheduler is in action and the traffic will be forwarded.

upvoted 1 times

  **Nikitas** 2 years, 10 months ago

One can't assume a rate limit of 100M, there's no policer specified. And the question specifically states that no other traffic is transiting the router. So, all traffic should be forwarded (Answer D). See here: <https://www.juniper.net/documentation/us/en/software/junos/cos-security-devices/topics/concept/cos-scheduler-default-security-setting.html>

«By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues.»

upvoted 3 times

  **Juniperguy** 3 years, 1 month ago

Yeah, also agree there is missing some information here.

upvoted 2 times

What is the purpose of the cluster-list attribute within a BGP route reflector group?

- A. to disable internal cluster re-advertisements
- B. to facilitate loop detection within the route reflector network
- C. to define the router that first advertised the route to the route reflector
- D. to override the router ID value within the cluster

Suggested Answer: *B*

  **yemicontrol** 3 years, 4 months ago

However, when two RRs are clients of each other and the routes are being reflected from one cluster to another, only one of the cluster IDs is included in the cluster list. This is because having one cluster ID in the cluster list is adequate for loop prevention in this case.

Answer B

upvoted 2 times

  **yemicontrol** 3 years, 4 months ago

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

upvoted 1 times

```

user@host# show protocols ospf
area 0.0.0.6 {
  nssa {
    default-lsa {
      default-metric 10;
      metric-type 1;
      type-7;
    }
  }
  no-summaries;
  area-range 192.168.16.0/20;
}

```

Referring to the ABR configuration shown in the exhibit, which three statements are correct? (Choose three.)

- A. The ABR advertises a default route to the NSSA using a Type 7 LSA.
- B. The ABR advertises a single Type 3 summary LSA to the backbone area for all Type 1 and Type 2 LSAs in the 192.168.16.0/20 range.
- C. The ABR advertises a Type 5 external LSA to the backbone area for each Type 7 LSA in the NSSA.
- D. The ABR does not summarize any routes within the 192.168.16.0/20 range.
- E. The ABR advertises a single Type 5 external LSA to the backbone area for all Type 7 LSAs in the NSSA.

Suggested Answer: ABC

Community vote distribution

ABC (100%)

🗳️ **ztw3587t** 1 year, 3 months ago

A, B and E are correct and D and E are incorrect. E, in special, is applied only on NSSA scenario
upvoted 1 times

🗳️ **johnmerchan** 1 year, 6 months ago

Selected Answer: ABC

B = CORRECT

With the no-summary keyword, the NSSA ABR will not advertise the inter-area routes (Type 3 and Type 4 summary routes) inside the NSSA, instead will advertise a default route. This default route will be propagated inside the NSSA as Type 3 LSA.
upvoted 1 times

🗳️ **much2furious** 1 year, 7 months ago

ABC is correct.

E is not correct: the router will send in area 0 the smaller routes, because area range command specified under area area-number stanza only summarizes intra-area prefixes.

Area range command under NSSA Stanza summarizes Type 7 LSA's before getting translated as Type 5 LSA's

<https://community.juniper.net/communities/community-home/digestviewer/viewthread?MID=71198>

upvoted 1 times

🗳️ **Rophy** 1 year, 10 months ago

I don't understand why so many guys think B is right, 'area-range' is under nssalevel rather than area level, So it wouldn't summary the type 1 and type 2 lsas in the area 6.
upvoted 1 times

🗳️ **Panadol** 2 years, 1 month ago

sorry ABE not ACE

upvoted 1 times

🗳️ **Panadol** 2 years, 1 month ago

ACE is the correct answer

upvoted 1 times

🗨️ **networkingcontrol** 2 years, 2 months ago

Correct answer: A,B,C

rea-range x.x.x.x/x --> Affect to outbound.

o-summaries --->affects inbound routing information. In this case, we won't accept LSA-3 into area 6.

upvoted 3 times

🗨️ **dexjov** 2 years, 5 months ago

B is controversial, as it's under the nssa, so it summarizes only type 7 routes, if it would be outside of nssa stanza, then B would be correct. On the other hand, answers C and E are mutually exclusive... Guess there are only 2 correct answers, A & C

upvoted 3 times

🗨️ **Dlbyam** 2 years, 5 months ago

ACE are correct

upvoted 1 times

🗨️ **Dlbyam** 1 year, 11 months ago

Sorry, ABC are correct

upvoted 1 times

🗨️ **oceans1908** 2 years, 6 months ago

Abc are correct

upvoted 2 times

🗨️ **oceans1908** 2 years, 7 months ago

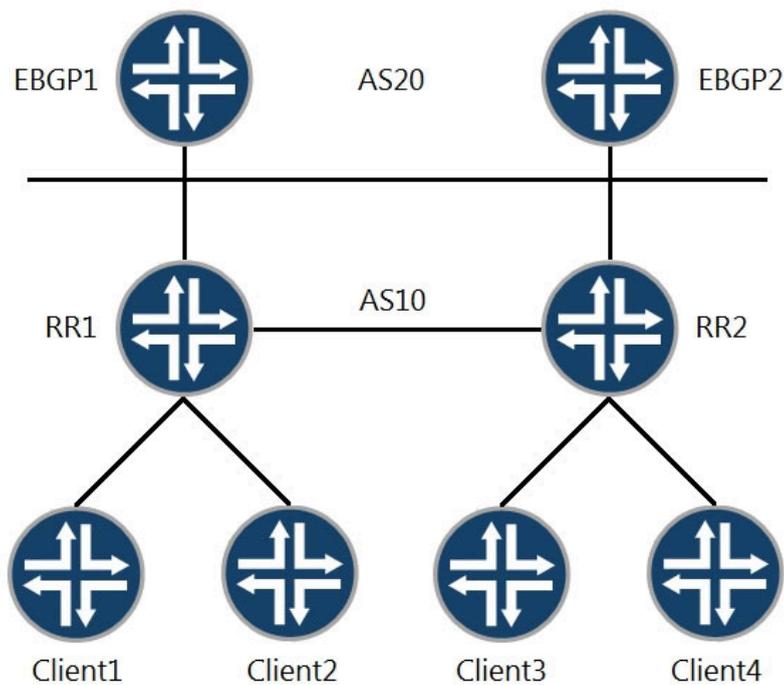
ACE the summary type 7 will be send to the backbona as a type 5

upvoted 1 times

🗨️ **Juniperguy** 3 years, 1 month ago

Agree with ABC.

upvoted 3 times



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

- A. RR1 advertises routes learned from Client1 to RR2 with itself as the next hop.
- B. RR2 advertises routes learned from Client3 to EBGP2 with itself as the next hop.
- C. RR1 and RR2 need the same cluster ID to exchange routes learned from their clients.
- D. RR2 adds its cluster ID when advertising routes from Client4 to Client3.

Suggested Answer: *BD*

PsamK 1 year, 9 months ago

Agree with BD, nexthopself is only relevant from clients. RR take advantage of iGP for nexthop reachability within AS. By default ebgp next hop is self on directly connected interface
upvoted 1 times

networkingcontrol 2 years, 2 months ago

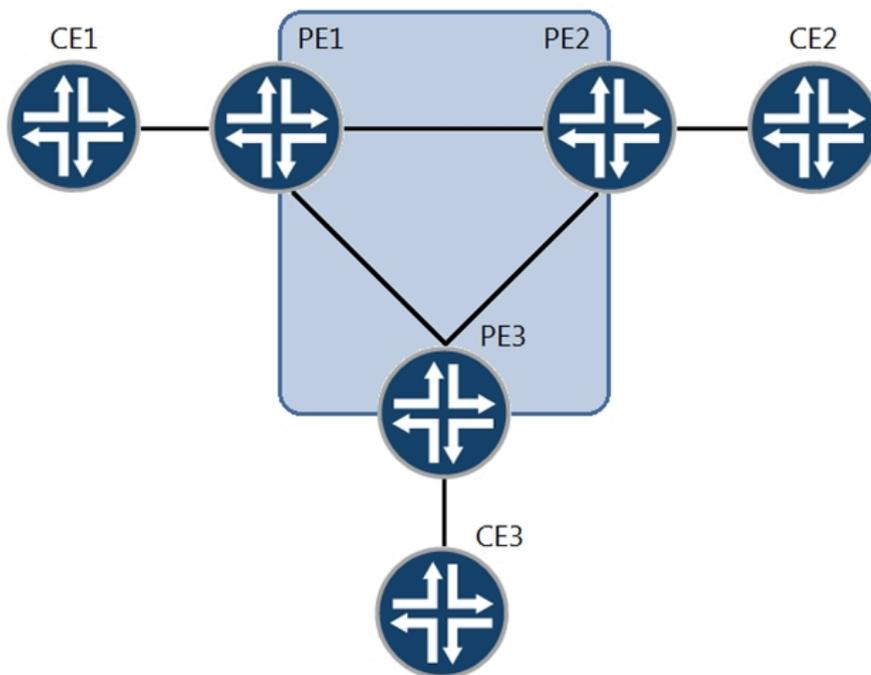
RRs that belong to same cluster ID shouldn't exchange routes.
Hence, B,D
upvoted 1 times

oceans1908 2 years, 7 months ago

Agree with BD
upvoted 1 times

Juniperguy 3 years, 1 month ago

Agree with BD
upvoted 1 times



You are provisioning Layer 2 circuits between sites CE1, CE2, and CE3.
Referring to the exhibit, which statement is true?

- A. A point-to-multipoint LSP must be created between sites.
- B. Each site must have only one VLAN configured to the PE.
- C. Site PE1 must have a point-to-multipoint link configured towards the core.
- D. Two VLANs must be configured from PE 1 to CE 1.

Suggested Answer: D

Community vote distribution

D (100%)

Juniperguy Highly Voted 3 years, 1 month ago
Agree with D.

As its two layer2 circuits, we need a vlan for each of the circuits. so 2x
upvoted 7 times

Sparks026 Most Recent 1 year, 11 months ago

Option B is correct according to this link -

https://www.juniper.net/documentation/en_US/release-independent/nce/topics/example/layer-two-circuits-ethernet-configuring-detailed-solutions.html

upvoted 1 times

BobbyAxelrod 2 years, 4 months ago

Selected Answer: D

we need a vlan for each of the circuits.

upvoted 1 times

oceans1908 2 years, 7 months ago

Agree with B

upvoted 1 times

oscarsd 2 years, 6 months ago

Not Agree, statement stand for a l2c for each site, l2c is p2p, not multipoint so you favor to configure 2 circuit for CE

upvoted 3 times

  **venn_phan** 2 years, 10 months ago

Agree with B

upvoted 2 times

What information is stored in a VRF table for a BGP Layer 2 VPN? (Choose three.)

- A. Layer 2 encapsulation
- B. local site ID
- C. remote interface of local CE device
- D. logical interlace provisioned to local CE device
- E. label-switched path

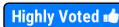
Suggested Answer: ADE

Community vote distribution

ABD (100%)

 **dre8051**  2 years, 11 months ago

L2 BGP VPN (Kempella) is ABD
upvoted 8 times

 **carroyoc**  2 years, 6 months ago

Selected Answer: ABD
Based on the book, Junos Layer 2 VPNs

The VRF table is populated with information provisioned for the local CE device and contains:

The local site ID
The site's layer 2 encapsulation
The logical interfaces provisioned to the local CE device
upvoted 6 times

 **Panadol**  2 years, 2 months ago

Agree with ABD
upvoted 2 times

 **EBNPW9** 2 years, 6 months ago

Selected Answer: ABD
See carroyoc's and venn_phan's comments.
upvoted 2 times

 **oceans1908** 2 years, 6 months ago

ADE local interface - label - encapsulation
upvoted 1 times

 **hrum** 2 years, 7 months ago

BGP L2VPN - ABD
upvoted 2 times

 **DevinK** 2 years, 7 months ago

Agree with Venn_Phan
ABD
upvoted 2 times

 **sandpit** 2 years, 10 months ago

(layer 2 encap + local site ID + logical interface provision to local CE)
upvoted 2 times

 **venn_phan** 2 years, 10 months ago

he VRF table is populated with information provisioned for the local CE device and contains:

- The local site ID;
- The site's Layer 2 encapsulation;

- The logical interfaces provisioned to the local CE device; and
 - A label base used to associated received traffic with one of the logical interfaces
- upvoted 3 times

  **Juniperguy** 3 years, 1 month ago

Agree with A, D,E

YOu need the logial interface for the routing instance, label-switched path and layer 2 encapsulation, ie: VLAN

upvoted 1 times

Which two statements about wide and narrow metrics used in IS-IS are correct? (Choose two.)

- A. Wide metrics are enabled with the wide-metrics-only parameter under protocols isis hierarchy.
- B. Narrow metrics are enabled by default and use 8 bits in TLVs to send information.
- C. Wide metrics are sent by default and use 24 bits in TLVs to send information.
- D. Disabling narrow metrics results in external routes being leaked from L1 to L2 areas automatically.

Suggested Answer: AB

Community vote distribution

CD (100%)

🗳️ **vasiliy19** Highly Voted 3 years, 2 months ago

C, D

Juniper support Narrow and Wide metrics by default.

upvoted 8 times

🗳️ **anonymonkey** 3 years, 1 month ago

Agree C&D

upvoted 3 times

🗳️ **ztw3587t** Most Recent 1 year, 3 months ago

Selected Answer: CD

CD are correct.

upvoted 1 times

🗳️ **jncie_examer** 1 year, 11 months ago

Selected Answer: CD

Agree with others

upvoted 1 times

🗳️ **Panadol** 2 years, 2 months ago

C&D for me is the correct answer

upvoted 1 times

🗳️ **networkingcontrol** 2 years, 2 months ago

Note that narrow metric use 6bits rather than 8bits.

upvoted 1 times

🗳️ **networkingcontrol** 2 years, 2 months ago

Correct answer: A, D

upvoted 1 times

🗳️ **NinjaCloud** 2 years, 4 months ago

B is fals because narrow metric uses 6 bits.

Answer: A,D

upvoted 1 times

🗳️ **carroyoc** 2 years, 7 months ago

Selected Answer: CD

With these TLVs, IS-IS metrics can have values up to 16,777,215 (224 - 1).

By default, Junos OS supports the sending and receiving of wide metrics.

<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/concept/isis-wide-metrics.html>

upvoted 2 times

🗳️ **oceans1908** 2 years, 7 months ago

CD both are enabled by default, narrow uses 6bits; when only wide auto l1 to l2

upvoted 1 times

🗨️ 👤 **Anniesyed1234** 2 years, 10 months ago

c & d are correct
upvoted 1 times

🗨️ 👤 **Sct38** 3 years ago

C+D. Wide and Narrow advertised by default. Export policy not required to export from L1 to L2 when only wide metrics enabled.
upvoted 3 times

🗨️ 👤 **Juniperguy** 3 years, 1 month ago

I think the question might be wrong.

narrow metrics use 6 bits, and wide metrics use 24 bit.

Both narrow and wide are enabled by default, but narrow is preferred.

So i would say answer C is correct, but ABD is not correct, unless the answer is abit wrong?

upvoted 1 times

🗨️ 👤 **Juniperguy** 3 years, 1 month ago

nvm, its AD
upvoted 2 times

🗨️ 👤 **anonymonkey** 3 years, 1 month ago

Not A

"Don't confuse wide-metrics with wide-metrics-only, as I stated before wide-metrics are enabled on Juniper by default,"

<https://networkengineering.stackexchange.com/questions/64469/understanding-the-is-is-wide-metric-style-for-traffic-engineering-in-juniper-mx>

upvoted 1 times

🗨️ 👤 **anonymonkey** 3 years, 1 month ago

Also explained here by juniper where the command is explained to change how the already enabled by default "wide metric" changes the parameters of this metric's algorithm

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/isis-wide-metrics.html>

upvoted 1 times

🗨️ 👤 **Sheet** 3 years, 1 month ago

A,D

With wide-metric, there is no distinction between L1 routes and L1 external routes

upvoted 3 times

```
[edit]
user@R4# run show pim rps
Instance: PIM.master

address-family   INET
RR address      Type      Mode      Holdtime  Timeout  Groups  Group prefixes
22.22.22.22     bootstrap sparse    150       108       0       224.0.0.0/4
33.33.33.33     bootstrap sparse    150       108       2       224.1.0.0/16

[edit]
user@R4# run show route 22.22.22.22

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

22.22.22.22/32    *[IS-IS/18] 00:32:27, metric 10
                  > to 10.1.1.2 via ge-0/0/0.0

inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0        *[Static/5] 00:13:55
                  > to 10.1.1.6 via ge-0/0/1.0

[edit]
user@R4# run show route 33.33.33.33

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

33.33.33.33/32   *[IS-IS/18] 00:32:43, metric 10
                  > to 10.1.1.6 via ge-0/0/1.0

inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0        *[Static/5] 00:14:25
                  > to 10.1.1.6 via ge-0/0/1.0

[edit]
user@R2# run show protocols pim
rp {
  bootstrap {
    family inet {
      priority 200;
    }
  }
  local {
    address 22.22.22.22;
    group-ranges {
      224.0.0.0/4;
    }
  }
}
interface all;

[edit]
user@R3# show protocols pim
rp {
  bootstrap {
    family inet {
      priority 210;
    }
  }
  local {
    address 33.33.33.33;
    group-ranges {
      224.1.0.0/16;
    }
  }
}
interface all;
```

R4 is directly connected to both RPs (R2 and R3). R4 is currently sending all joins upstream to R3 but you want to load balance the joins between both RPs.

Referring to the exhibit, which configuration change will solve this issue?

- A. Configure the join-load-balance parameter under PIM on R4.
- B. Configure the default route in inet.2 on R4 from R3 as the next hop to both R3 and R2.
- C. Configure the group-range parameter to be the same on R2 and R3.
- D. Configure the bootstrap priority on R2 to be the same as R3.

Suggested Answer: A

Community vote distribution

A (67%)

D (33%)

 **magmartin**  2 years, 1 month ago

I am alone but I think the answer is C. :)

The multicast network is using a bootstrap router to announce the available RPs to the rest of PIM neighbors. The priority value shown in the exhibit is used to choose which router will act as bootstrap, but it does not indicate which one will be the RP, so option D is discarded.

The definition of join-load-balance says "you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source".

I think this refers to balancing traffic from a specific source to the destination using several equal cost paths, but it does not mean that we will share the load between different RPs as the question asks, so I would discard also option A.

As we are using a bootstrap router, we can check that the group to RP mapping is the following (<https://www.rfc-editor.org/rfc/rfc7761.html#section-4.7.1>)

If we choose option C, we will have a tie in the first 2 steps and we will let the output of the hash function decide, which I think is some kind of "random" number which helps to balance the traffic to different RPs.

upvoted 7 times

 **amnesiac1** 2 years, 1 month ago

first part of your answer is correct. However on the second part, they won't tie on the first step, R3 wins.

upvoted 1 times

 **amnesiac1** 2 years, 1 month ago

agree, answer is C

upvoted 3 times

 **ztw3587t**  1 year, 3 months ago

Selected Answer: A

The correct answer is A.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors.

There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/task/mcast-pim-join-load-balance.html>

upvoted 1 times

 **Tomerd** 2 years, 5 months ago

D for me,

A: is incorrect -

You configure PIM join load balancing on the non-RP routers in the PIM domain.

upvoted 2 times

 **albert1687** 2 years, 2 months ago

Exactly and R4 is a non-RP. R2 and R3 are RPs. So 'A' is correct.

upvoted 1 times

 **Tomerd** 2 years, 5 months ago

D for me,

A: is incorrect -

You configure PIM join load balancing on the non-RP routers in the PIM domain.

upvoted 1 times

 **ptt** 2 years, 6 months ago

100% not A, because the multicast section i got zero percent correct after i chose A for this question.

upvoted 3 times

 **EBNPW9** 2 years, 6 months ago

Selected Answer: A

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/task/mcast-pim-join-load-balance.html>

upvoted 1 times

🗨️ 👤 **carroyoc** 2 years, 7 months ago

Selected Answer: A

The domain bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routers use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. A priority of 0 disables the function for IPv4 and does not cause the routing device to send bootstrap router packets with a 0 in the priority field. The routing device with the highest priority value is elected to be the bootstrap router.

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

This means that the priority is "like the priority of OSPF election DR route" and we require the load balance configuration only

upvoted 2 times

🗨️ 👤 **TungNguyenTruc** 2 years, 7 months ago

Selected Answer: A

* PIM join load balancing:

Normally, if multiple equal-cost paths toward source exist, only 1 will pass the RPF check

For PIM sparse mode, load sharing can be enabled so that all equal-cost links can be used on an (S,G) by (S,G) basis

upvoted 2 times

🗨️ 👤 **TungNguyenTruc** 2 years, 7 months ago

[edit protocols pim]

user@router# show

join-load-balance;

upvoted 2 times

🗨️ 👤 **boyseven777** 2 years, 7 months ago

Selected Answer: D

R3 has a higher priority at the moment

upvoted 3 times

🗨️ 👤 **Dbyam** 2 years, 8 months ago

A. Configure the join-load-balance parameter under PIM on R4.

is the correct answer.

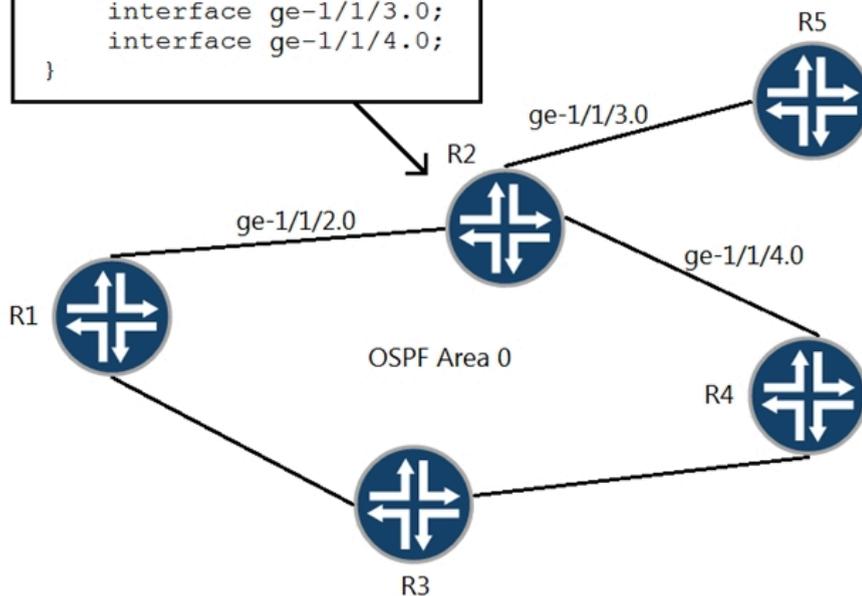
upvoted 1 times

🗨️ 👤 **Anniesyed1234** 2 years, 10 months ago

Agree with A

upvoted 1 times

```
[edit protocols ospf]
user@R2# show
overload;
area 0.0.0.0 {
  interface ge-1/1/2.0;
  interface ge-1/1/3.0;
  interface ge-1/1/4.0;
}
```



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

- A. R2 stops sending LSAs into the network.
- B. The OSPF interface metrics on R2 are all set to 65535.
- C. R1 will never forward transit traffic through R2.
- D. Transit traffic from R1 to R4 will traverse R3.

Suggested Answer: *BD*

Community vote distribution

BD (67%)

BC (33%)

ztw3587t 1 year, 3 months ago

Selected Answer: BD

BD is correct

upvoted 2 times

wingnut 1 year, 4 months ago

A- Incorrect

B- Yes correct

10.100.255.2/32 *[OSPF/10] 00:06:46, metric 65538

> to 10.100.254.1 via ge-0/0/1.0

C- Incorrect it will take the path is there it not a better path

D- Correct if metric is lower which is usually will be

so answer is B and D

upvoted 1 times

Jen_663 1 year, 8 months ago

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

BC

upvoted 1 times

mohdema 1 year, 9 months ago

Selected Answer: BC

Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.

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

upvoted 1 times

  **Dbyam** 2 years, 5 months ago

B and D are correct

upvoted 2 times

  **oceans1908** 2 years, 7 months ago

B and D

upvoted 2 times

  **Anniesyed1234** 2 years, 10 months ago

B & D correct

upvoted 2 times

  **Anniesyed1234** 2 years, 10 months ago

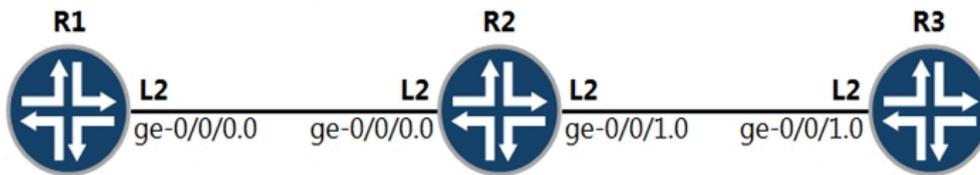
D is the right answer

upvoted 2 times

```

user@R2# show protocols isis
level 1 disable;
interface ge-0/0/0.0;
interface ge-0/0/1.0 {
    level 2 metric 300;
}

```



```

user@R1# show protocols isis
level 1 disable;
interface ge-0/0/0.0;

user@R3# show protocols isis
level 1 disable;
interface ge-0/0/1.0;

```

AREA 49.0001

Referring to the exhibit, what will the IS-IS cost be for R1 to reach R3?

- A. 73
- B. 20
- C. 301
- D. 310

Suggested Answer: D

Community vote distribution

A (100%)

Juniperguy Highly Voted 3 years, 1 month ago

Remember metric is only "counted" on outbound interface. So when the path is R1 to R3, its 10+63(as 63 is max when 300 is set on the outbound interface towards R3 when wide metrics is not enabled).

If the traffic came the other way from R3 to R1, it would actually be 10+10 = 20
upvoted 12 times

anonymonkey 3 years, 1 month ago

"A" agreed with the explanation. The answer would be 310 if wide-metric-only were configured
upvoted 2 times

ztw3587t Most Recent 1 year, 3 months ago

Selected Answer: A

A is correct: 10 + 63 -> 73
upvoted 2 times

azou25 2 years, 4 months ago

A is the correct

3.3.3.3/32 *[IS-IS/18] 00:00:20, metric 73

> to 10.1.0.2 via ge-0/0/0.0

Config R1 :

```

root@vMX-1> show configuration protocols isis
level 1 disable;
interface ge-0/0/0.0;
interface lo0.0 {
level 1 passive;
}

```

Config R2 :

```
root@vMX-2> show configuration protocols isis
level 1 disable;
interface ge-0/0/0.0;
interface ge-0/0/1.0 {
level 2 metric 300;
}
interface lo0.0 {
passive;
}
```

config R3 :

```
root@vMX-3> show configuration protocols isis
level 1 disable;
interface ge-0/0/1.0;
interface lo0.0 {
passive;
}
```

upvoted 1 times

🗨️ **EBNPW9** 2 years, 6 months ago

Selected Answer: A

"By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs. To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, you must include the wide-metrics-only statement in the IS-IS configuration."

<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/concept/isis-wide-metrics.html>

upvoted 1 times

🗨️ **oceans1908** 2 years, 6 months ago

Selected Answer: A

I labbed it up and answer is 73, 63 (which is max port cost under narrow metrics) plus 10

upvoted 2 times

🗨️ **oceans1908** 2 years, 6 months ago

Cost for R1 to reach R3, so the cost of the path will be the cost of ge-0/0/1.0 of R3 which is 10 and the cost of ge-0/0/0.0 of R2 which is also 10. So it must be 20.

upvoted 1 times

🗨️ **chris_2_a** 2 years, 6 months ago

Selected Answer: A

Test it in lab environment and the answer is A.
below

```
R1> show isis route
```

```
172.16.56.0/24 2 44 73 int ge-0/0/1.0 IPV4 P3
```

upvoted 3 times

🗨️ **hrum** 2 years, 7 months ago

10+63 - A

upvoted 1 times

🗨️ **oceans1908** 2 years, 7 months ago

A 63 plus 10 as narrow is still active

upvoted 1 times

🗨️ **TT98** 2 years, 11 months ago

I think answer should be 73 (10+63). Without wide metric enabled, TLV from R2 will advertise metric of 63 to R1.

upvoted 2 times

🗨️ **Sct38** 3 years ago

Answer is 73 as both wide and narrow metrics enabled; router will only use narrow metric; 300 will become 63 when advertised in narrow.

upvoted 3 times

🗨️ **Juniperguy** 3 years, 1 month ago

Agree. A is correct

upvoted 2 times

🗨️ 👤 **Sheet** 3 years, 1 month ago

A is correct.

Wide metric is not enabled, thus, the max cost of for metric 300 will be 63.

upvoted 3 times

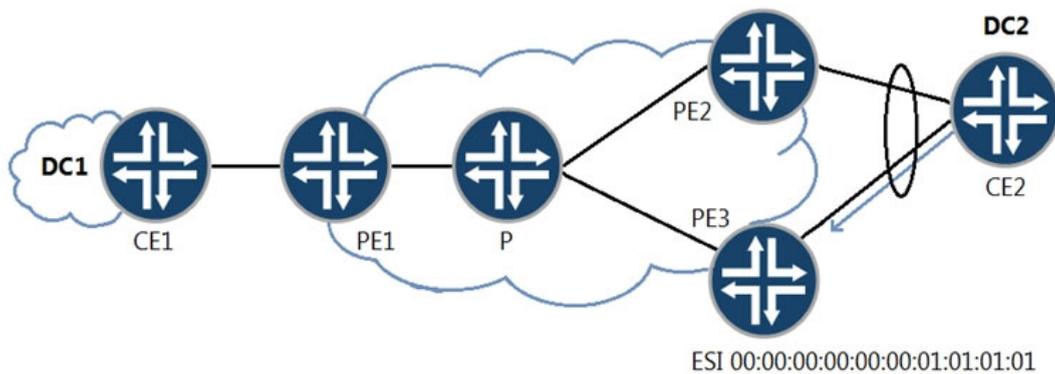
🗨️ 👤 **hbstyleboy** 3 years, 3 months ago

D is correct:

Default IS-IS metric is 10, so $300 + 10 = 310$

<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/concept/isis-wide-metrics.html>

upvoted 2 times



Referring to the exhibit, traffic sent from CE-A2 to PE3 does not loop back to CE-A2 through PE2.

Which two EVPN functions accomplish this task? (Choose two.)

- A. multicast ingress replication
- B. aliasing
- C. split horizon
- D. designated forwarder election

Suggested Answer: CD

Community vote distribution

CD (100%)

zwt3587t 1 year, 3 months ago

Selected Answer: CD

CD are correct.

- Split Horizon avoid the BUM (Broadcast, Unknown, Multicas) traffic to be forward to CE for more one PE to avoid Loop.
- Designated Forwarder election process choose one PE to forward BUM traffic to avoid Loop.

upvoted 1 times

somanyquestions 2 years, 10 months ago

CD;

https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-bgp-multihoming-overview.html#evpn-multihoming-overview_d3873e262

upvoted 2 times

Which two types of LSAs have an area scope? (Choose two.)

- A. Type 2
- B. Type 5
- C. Type 11
- D. Type 7

Suggested Answer: AC

Community vote distribution

AD (75%)

AC (25%)

TT98 Highly Voted 3 years, 2 months ago

Answer should be AD. Type 2's flooding scope is obvious. Type 7 is never flooded outside of the NSSA, so it has an area flooding scope. Type 11 has AS flooding scope, similar to type 5.

upvoted 11 times

ztw3587t Most Recent 1 year, 3 months ago

Selected Answer: AD

AD is correct

upvoted 1 times

ztw3587t 1 year, 3 months ago

AD is correct

upvoted 1 times

johnmerchan 1 year, 6 months ago

Selected Answer: AC

<https://www.firewall.cx/networking-topics/routing/ospf-routing-protocol/1178-ospf-lsa-types-explained.html>

type 2 y type 11 el proceso de inundación se realiza en la propia área por esto la respuesta mientras que 5 y 7 lo realizan debido a que permite la distribución externa

upvoted 1 times

TgDgNU 1 year, 11 months ago

Selected Answer: AD

It's obviously A + D

Type 2 is changed to Type5\Type7 when transferred to another area

Type 7 is never advertised outside nssa

upvoted 1 times

Panadol 2 years, 2 months ago

A and D

upvoted 1 times

networkingcontrol 2 years, 2 months ago

Answer:A, D

LSA1-4,7 -> Area scope

LDSA5 --> Domain scope

upvoted 1 times

Dlbyam 2 years, 5 months ago

Type 1 2 and 7 have area scope. They do not cross their own area

upvoted 1 times

EBNPW9 2 years, 6 months ago

Selected Answer: AD

https://techhub.hp.com/eginfolib/networking/docs/switches/5710/5200-4992_l3-ip-rtng_cg/content/517702239.htm

upvoted 1 times

🗨️ 👤 **oceans1908** 2 years, 7 months ago

A and D

upvoted 1 times

🗨️ 👤 **DiByam** 2 years, 8 months ago

Type 2 and Type 7 is the correct answer

upvoted 1 times

🗨️ 👤 **nguyen03t2** 2 years, 11 months ago

It should be A, D

upvoted 1 times

🗨️ 👤 **M_A_T** 2 years, 11 months ago

A and D I would say. 1 2 3 4 area scope. 5 domain scope. 6 no idea. 7 area 8 no idea. 9 link. 10 area 11 domain.

upvoted 3 times

🗨️ 👤 **Juniperguy** 3 years, 1 month ago

Agree with TT98. AD.

Type 2 is obvious. And type 7 is only used in NSSA. From the site that anonymonkey mentions it says under type 7: "The flooding scope of Type-7 LSA is a single area."

upvoted 4 times

🗨️ 👤 **anonymonkey** 3 years, 1 month ago

100% agree, I put the wrong letters/selection with my validation and never spoke to 11 as they are similar to 5

upvoted 1 times

🗨️ 👤 **anonymonkey** 3 years, 1 month ago

AC, Type 5 are domain scopes, Type 2 are area scopes: <https://ipmazariegos.com/2016/01/06/ospf-lsa-flooding-scope/>

upvoted 1 times

🗨️ 👤 **anonymonkey** 3 years, 1 month ago

A&D, not A&C

upvoted 1 times

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

Shouldn't the answer be type 10 (OSPF Area Scope Opaque LSA) and type 11?

upvoted 2 times

```

user@R1> show ospf3 interface
Interface      State      Area      DR ID      BDR ID      Nbrs
ge-0/0/0.0    DR         0.0.0.0   172.16.1.2 172.16.1.1   1
ge-0/0/0.0    PtToPt    0.0.0.1   0.0.0.0    0.0.0.0     1
ge-0/0/1.0    BDR        0.0.0.1   172.16.1.1 172.16.1.2   1

```

```

user@R1> show ospf3 neighbor
ID            Interface  State  Pri    Dead
172.16.1.1    ge-0/0/0.0 Full   128    39
  Neighbor-address fe80::20c:29ff:fef9:7f7b
  Area 0.0.0.0
172.16.1.1    ge-0/0/0.0 Full   128    37
  Neighbor-address fe80::20c:29ff:fef9:7f7b
  Area 0.0.0.1
172.16.1.1    ge-0/0/1.0 Full   128    37
  Neighbor-address fe80::20c:29ff:fef9:7f85
  Area 0.0.0.1

```

Referring to the exhibit, which OSPFv3 configuration is implemented on router R1?

- A. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0
- B. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 virtual-link neighbor-id 172.16.1.2
- C. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 secondary
- D. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 interface-type p2p

Suggested Answer: D

Community vote distribution

C (86%)

14%

 **johnmerchan** 1 year, 6 months ago

Selected Answer: D

OSPF interface type point to point (p2p)
no designated router/backup designated router election occurs
Used between adjacencies involving exactly two neighbors
saves time getting the adjacency to a full state
Useful on back to back ethernet links
upvoted 1 times

 **mohdema** 1 year, 9 months ago

Selected Answer: C

By default, a single interface can belong to only one OSPF area. You can configure a single interface to belong in multiple OSPF areas. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. When configuring a secondary interface, consider the following:

For OSPFv2, you cannot configure point-to-multipoint and nonbroadcast multiaccess (NBMA) network interfaces as a secondary interface because secondary interfaces are treated as a point-to-point unnumbered link.

Secondary interfaces are supported for LAN interfaces (the primary interface can be a LAN interface, but any secondary interfaces are treated as point-to-point unnumbered links over the LAN). In this scenario, you must ensure that there are only two routing devices on the LAN or that there are only two routing devices on the LAN that have secondary interfaces configured for a specific OSPF area.

<https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/topic-map/configuring-ospf-areas.html#id-example-configuring-multiarea-adjacency-for-ospf>

upvoted 1 times

- 🗨️ **Dlbyam** 1 year, 11 months ago
C is correct because of secondary keyword, which will be DR plus P2p
upvoted 1 times
- 🗨️ **EBNPW9** 2 years, 6 months ago
Selected Answer: C
I've tested in a lab. Also see TT98's and Sheet's answers.
upvoted 2 times
- 🗨️ **oscarsd** 2 years, 6 months ago
Selected Answer: C
C is correct answer
upvoted 3 times
- 🗨️ **somanyquestions** 2 years, 10 months ago
C:
A logical interface can be configured as primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.
https://www.juniper.net/documentation/en_US/junos/topics/reference/configuration-statement/secondary-edit-protocols-ospf.html
upvoted 4 times
- 🗨️ **JNCXX** 2 years, 6 months ago
agree with C
upvoted 1 times
- 🗨️ **ak_sisko** 3 years ago
Correct C
upvoted 2 times
- 🗨️ **Juniperguy** 3 years, 1 month ago
Agree with Sheet. adding secondary to the configuration makes the interface PtToPt

Correct C
upvoted 2 times
- 🗨️ **Sheet** 3 years, 1 month ago
C is correct. Secondary adjacencies are always point-to-point.
upvoted 3 times
- 🗨️ **vasiliy19** 3 years, 2 months ago
It look like there should be 2 answers - C,D
upvoted 1 times
- 🗨️ **mikkamilla** 3 years, 2 months ago
D wouldn't even commit, I guess
upvoted 1 times
- 🗨️ **oscarsd** 2 years, 6 months ago
You're right, it won't commit
upvoted 1 times
- 🗨️ **carroyoc** 2 years, 6 months ago
[edit protocols ospf3 area 0.0.0.1]
'interface ge-0/0/0.0'
Duplicate interface configured under area 0.0.0.0
error: configuration check-out failed
upvoted 2 times
- 🗨️ **TT98** 3 years, 2 months ago
Answer is C. This is an example of a multiarea adjacency. Interface ge-0/0/0.0 has both DR and PTP, which with ordinary OSPF adjacency rules will be impossible.
upvoted 4 times

```
[edit]
user@router# show protocols bgp group ISP
export ISP;
neighbor 172.16.20.20 {
    peer-as 65310;
}
neighbor 192.168.51.200 {
    peer-as 65441;
}
```

```
[edit]
user@router# show policy-options
policy-statement ISP {
    term ROUTES {
        from {
            protocol bgp;
            as-path LOCAL-ROUTES;
        }
        then accept;
    }
}
as-path LOCAL-ROUTES "(.*)";
```

Your network is connected to two different ISPs and you notice that they are using your network for transit traffic. In this scenario, which two configuration statements will solve this problem? (Choose two.)

- A. set policy-options policy-statement ISP term REST then reject
- B. set policy-options policy-statement ISP term ROUTES then reject
- C. set policy-options as-path LOCAL-ROUTES ()€λ€
- D. set policy-options as-path LOCAL-ROUTES λ+(65441|65310)€λ€

Suggested Answer: AC

Community vote distribution

AC (50%)

AD (25%)

BC (25%)

 **somanyquestions** Highly Voted 2 years, 10 months ago

A; default is accept, so leftovers should be rejected.

C; You can use AS path regular expressions to create a null AS path that matches routes (prefixes) that have originated in your AS. These routes have not been advertised to your AS by any external peers. To create a null AS path, use the parentheses operator enclosed in quotation marks with no intervening spaces: "()"

https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/policy-configuring-as-path-regular-expressions-to-use-as-routing-policy-match-conditions.html#understanding-as-path-regular-expressions-for-use-as-routing-policy-match-conditions__id-10256235
upvoted 11 times

 **ztw3587t** Most Recent 1 year, 3 months ago

Selected Answer: AC

A and C are correct.

*A - It's necessary to create the term REST to reject the ISPs routes

*C - You can use AS path regular expressions to create a null AS path that matches routes (prefixes) that have originated in your AS. These routes have not been advertised to your AS by any external peers. To create a null AS path, use the parentheses operator enclosed in quotation marks with no intervening spaces:

```
[edit policy-options]
as-path null-as "()";
```

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/policy-configuring-as-path-regular-expressions-to-use-as-routing-policy-match-conditions.html>

upvoted 1 times

🗨️ **ztw3587t** 1 year, 3 months ago

Selected Answer: AC

A and C are correct.

*A - It's necessary to create the term REST to reject the ISPs routes

*C - You can use AS path regular expressions to create a null AS path that matches routes (prefixes) that have originated in your AS. These routes have not been advertised to your AS by any external peers. To create a null AS path, use the parentheses operator enclosed in quotation marks with no intervening spaces:

[edit policy-options]

```
as-path null-as "()";
```

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/policy-configuring-as-path-regular-expressions-to-use-as-routing-policy-match-conditions.html>

upvoted 1 times

🗨️ **mohdema** 1 year, 9 months ago

Selected Answer: AD

A. set policy-options policy-statement ISP term ROUTES then reject

B. set policy-options as-path LOCAL-ROUTES ")"

C. set policy-options policy-statement ISP term REST then reject

D. set policy-options as-path LOCAL-ROUTES "(65310 | 65441) +"

It's D*

{m,n} Match at least m and at most n repetitions of term

{m} Match exactly m repetitions of term

{m,} Match m or more repetitions of term

* Match 0 or more repetitions of term, same as {0,}

+ Match 1 or more repetitions of term, same as {1,}

? Match 0 or 1 repetitions of term, same as {0,1}

| Match one of the two terms on either side of the pipe

- Used to represent a range

'...>() Used to group terms, or indicate null with no space

upvoted 1 times

🗨️ **mohdema** 1 year, 9 months ago

(...>() Used to group terms, or indicate null with no space

upvoted 1 times

🗨️ **Dlbyam** 1 year, 10 months ago

A and C are correct

upvoted 1 times

🗨️ **bartahr** 1 year, 11 months ago

Selected Answer: BC

C is correct, we need to mark our internal BGP routes - this means AS_PATH "()"

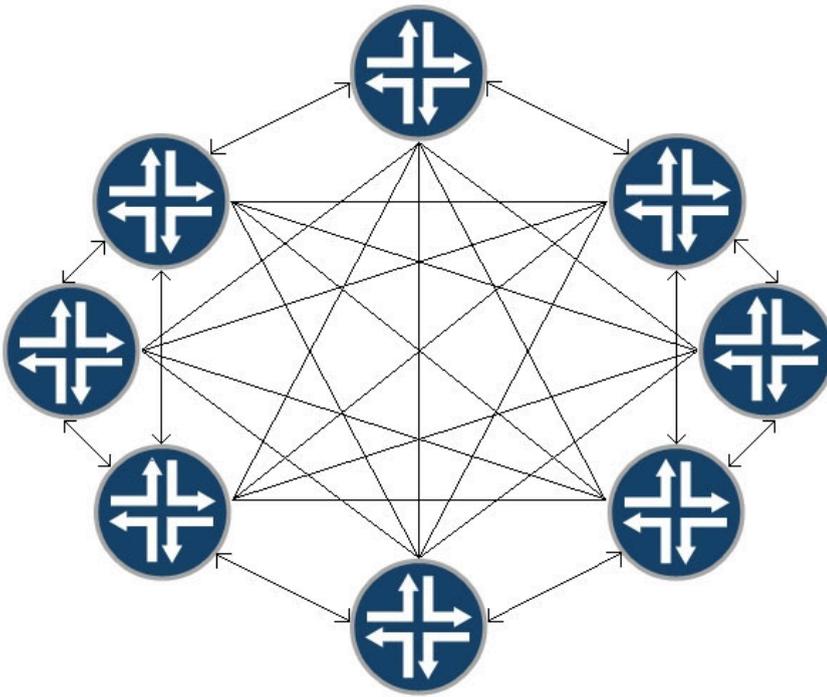
B is correct answers - we need to forbid our routes to be used by our EBGp peers, so we shouldn't export all our local routes, we should export all other BGP routes from other BGP peers which is default BGP behavior

upvoted 1 times

🗨️ **guesswho123** 1 year, 9 months ago

B is incorrect, we have to export our routes.

upvoted 1 times



A customer wants to reduce LSP flooding in their IS-IS network.
Which parameter should you change to accomplish this task?

- A. [edit protocols isis] user@router# set spf-options rapid-runs 5
- B. [edit protocols isis interface <interface-name>] user@router# set csnp-interval 65535
- C. [edit protocols isis interface <interface-name>] user@router# set lsp-interval 1000
- D. [edit protocols isis interface <interface-name>] user@router# set mesh-group <mesh-group-number>

Suggested Answer: B

Community vote distribution

D (75%)

B (25%)

Juniperguy Highly Voted 3 years, 1 month ago

Answer D, for Mesh group.

CSNP will only be sent on topology change and expiration.

upvoted 6 times

ztw3587t Most Recent 1 year, 3 months ago

Selected Answer: D

D is correct, mesh-group

upvoted 1 times

NetworkGuy420 1 year, 4 months ago

Option D:

<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/example/isis-mesh.html>

After the adjacencies have been up for about 38 minutes, the output shows that Device R1 has received 73 link-state PDUs and sent 37 link-state PDUs. In the same topology in the same amount of time without the mesh group enabled, Device R1 would have received roughly 156 link-state PDUs and sent roughly 117 link-state PDUs.

<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/example/isis-mesh.html>

"After the adjacencies have been up for about 38 minutes, the output shows that Device R1 has received 73 link-state PDUs and sent 37 link-state PDUs. In the same topology in the same amount of time without the mesh group enabled, Device R1 would have received roughly 156 link-state PDUs and sent roughly 117 link-state PDUs."

upvoted 1 times

🗨️ 👤 **johnmerchan** 1 year, 6 months ago

Selected Answer: B

<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/concept/isis-csnp-interval-understanding.html>

La recomendación general es utilizar los valores predeterminados o aumentar el intervalo CSNP si hay un gran número de circuitos de radiodifusión que necesitan ser suministrados con CSNP nuevos. Aumentar el intervalo puede ayudar a proteger contra inundaciones de CSNP.

upvoted 1 times

🗨️ 👤 **mohdema** 1 year, 9 months ago

Selected Answer: D

A mesh group is a set of routing devices that are fully connected. That is, they have a fully meshed topology.

Junos OS supports IS-IS mesh groups as documented in RFC 2973, IS-IS Mesh Groups.

When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.

upvoted 1 times

🗨️ 👤 **Panadol** 2 years, 2 months ago

D mesh group

upvoted 1 times

🗨️ 👤 **NinjaCloud** 2 years, 4 months ago

Correct Answer: D

Explanation: When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.

upvoted 3 times

🗨️ 👤 **EBNPW9** 2 years, 6 months ago

Selected Answer: D

I agree with mikkamilla. Also see <https://www.juniper.net/documentation/us/en/software/junos/isis/topics/example/isis-mesh.html>

upvoted 1 times

🗨️ 👤 **Anniesyed1234** 2 years, 10 months ago

Mesh group is used for reducing redundant LSP, I say if redundant LSP is not mentioned B could be right however we do see from the picture that redundant LSPs would be flowing

upvoted 3 times

🗨️ 👤 **M_A_T** 2 years, 11 months ago

I would say mesh groups are invented for a situation like this so D would be my answer.

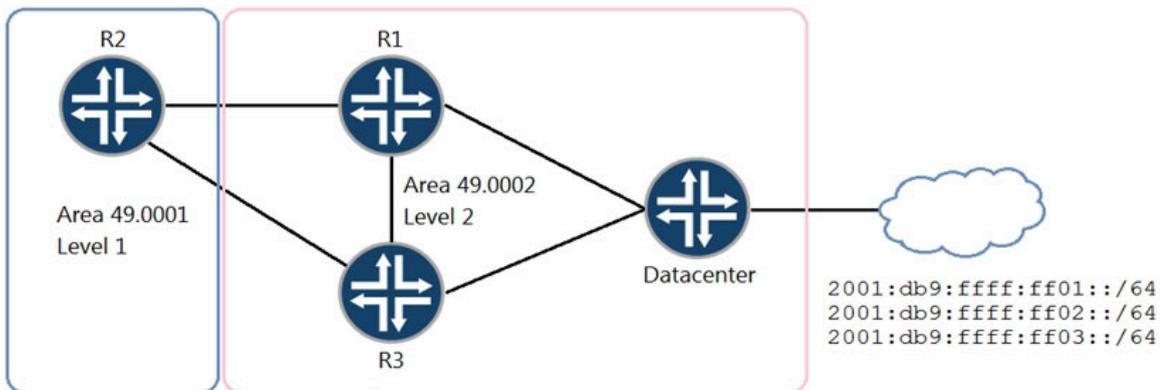
upvoted 4 times

🗨️ 👤 **mikkamilla** 3 years, 3 months ago

what about mesh-groups?

<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/concept/isis-mesh-group-understanding.html>

upvoted 4 times



A network designer wants to ensure that traffic from R2 destined for 2001:db9:ffff:ff00::/62 always traverses the R2-R1 link if that link is available.

Referring to the exhibit, which configuration change will satisfy this requirement?

A.

```
user@R1# show protocols isis
export leak-v6;

user@R1# show policy-options
policy-statement leak-v6 {
  term DC-routes {
    from {
      protocol isis;
      level 2;
      route-filter 2001:db9:ffff:ff00::/62 orlonger;
    }
    to level 1;
    then accept;
  }
}
```

B.

```
user@R2# show protocols isis
export leak-v6;

user@R2# show policy-options
policy-statement leak-v6 {
  term DC-routes {
    from {
      protocol isis;
      level 2;
      route-filter 2001:db9:ffff:ff00::/62 orlonger;
    }
    to level 1;
    then accept;
  }
}
```

C.

```
user@R1# show protocols isis
import leak-v6;

user@R1# show policy-options
policy-statement leak-v6 {
  term DC-routes {
    from {
      protocol isis;
      level 1;
      route-filter 2001:db9:ffff:ff00::/62 orlonger;
    }
    to level 2;
    then accept;
  }
}
```

Suggested Answer: A

A is correct because the configuration must be applied on L1/L2 Router and also it's necessary be applied on L2 route.

upvoted 1 times

🗨️ **ztw3587t** 1 year, 3 months ago

A is correct

upvoted 1 times

🗨️ **mohdema** 1 year, 9 months ago

My bad it's on R1 thought R2 was the level 2

upvoted 1 times

🗨️ **mohdema** 1 year, 9 months ago

B

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/example/example-configuring-is-is-route-leaking-L2-to-L1.html>

upvoted 1 times

🗨️ **bartahr** 1 year, 11 months ago

B is correct. Export should be done on R2 since R2 is L1/L2 router

upvoted 2 times

🗨️ **oceans1908** 2 years, 7 months ago

A on router 1 the l2 l1 router

upvoted 1 times

🗨️ **12_sdeEQ** 2 years, 10 months ago

Agree with Answer A.

upvoted 1 times

🗨️ **Juniperguy** 3 years, 1 month ago

Agree with Answer A.

1# the config has to be done on R1, as R1 is in Level 2 and has to export the route to layer 1.

Answer B is the wrong router to do configuration on and C is routes from area 1 to area 2.

upvoted 3 times

🗨️ **hcccc** 3 years, 3 months ago

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/example/example-configuring-is-is-route-leaking-L2-to-L1.html>

upvoted 2 times

```
[edit]
user@R2# run show isis adjacency
Interface    System  L   State      Hold (secs) SNPA
ge-0/0/0.0   R3      1   Up          6 0:50:56:93:54:4b
ge-0/0/0.0   R3      2   Up          7 0:50:56:93:54:4b
ge-0/0/1.0   R4      2   Up          7 0:50:56:93:54:4b

[edit]
user@R2# show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 172.16.2.2/30;
      }
      family iso;
    }
  }
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.2/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 22.22.22.22/32;
      }
      family iso;
      address 49.0001.0022.2222.0022.00;
    }
  }
}
protocols {
  isis {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
    interface lo0.0 {
      level 1 disable;
    }
  }
}
```

R2 has IS-IS adjacencies with R3 and R4. You want to ensure that R2 has both a level 1 and level 2 adjacency to both R3 and R4, but R2 only has one adjacency with R4.

Referring to the exhibit, which configuration change will solve this issue?

- A. Change the IS-IS area on R4 to match R2.
- B. Remove the level 1 disable configuration from R4.
- C. Remove the level 1 disable configuration from R2.
- D. Change the IS-IS area on R2 to match R4.

Suggested Answer: C

Community vote distribution

A (83%)

B (17%)

 **JuniperGuy** Highly Voted 3 years, 1 month ago

A - Change the IS-IS area on R4 to match R2.

upvoted 10 times

 **dix** Highly Voted 2 years, 10 months ago

If R2 and R4 are in different areas, and if R4 has "level1 disable" under [protocols isis], it generates same effect on R2...only L2 adjacency between R2 and R4. Then Option A and option B are valid...

upvoted 5 times

 **dix** 2 years, 10 months ago

On others pages... option A is selected.

upvoted 1 times

 **ztw3587t** Most Recent 1 year, 3 months ago

Selected Answer: A

A is correct. It's necessary that both routers are on the same area ID to establish L1/L2 adjacency.

upvoted 1 times

🗨️ **ztw3587t** 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ **Dlbyam** 1 year, 10 months ago

A and B, both are correct

upvoted 1 times

🗨️ **OGab** 2 years, 1 month ago

Selected Answer: A

the reason why there is only a L2 session is because the area is not matching on R2

upvoted 1 times

🗨️ **oscarsd** 2 years, 6 months ago

Selected Answer: A

Correct A, for L1 adjacency to form need to be in same Area ID

upvoted 3 times

🗨️ **oscarsd** 2 years, 6 months ago

Correct A, for L1 adjacency to form need to be in same Area ID

upvoted 1 times

🗨️ **exambreaker** 3 years, 2 months ago

I meant only L2 adjacency..

upvoted 1 times

🗨️ **exambreaker** 3 years, 2 months ago

The correct answer here is A as it looks like R4 is with different area id and in this case, R2 will form only L1 adjacency with it as per the rule

upvoted 3 times

🗨️ **Sheet** 3 years, 1 month ago

But, if you change the area ID, it will break the L1 adjacency with R3. The requirement in the question is that R2 should form L1 and L2 adjacencies with both the routers.

I think R4 needs to change its area, thus B

upvoted 3 times

🗨️ **mikkamilla** 3 years, 2 months ago

i think it's B, there's no offending configuration on R2

upvoted 3 times

You are asked to configure a series of interface policers and firewall filters, which include policers, on the same device. You must ensure that the two configuration methods do not conflict.

What are two considerations when performing this task? (Choose two.)

- A. On inbound traffic, firewall filters are applied before interface policers.
- B. On outbound traffic, interface policers are applied before firewall filters.
- C. On outbound traffic, firewall filters are applied before interface policers.
- D. On inbound traffic, interface policers are applied before firewall filters.

Suggested Answer: CD

Community vote distribution

CD (100%)

🗨️ **NetworkGuy420** 1 year, 4 months ago

C and D

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/policer-order-of-operations-with-firewall-filters.html>

upvoted 1 times

🗨️ **mohdema** 1 year, 9 months ago

Selected Answer: CD

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/policer-order-of-operations-with-firewall-filters.html>

See figure 1

upvoted 1 times

🗨️ **jprat** 1 year, 12 months ago

Selected Answer: CD

C & D for the answer

upvoted 1 times

🗨️ **Juniperguy** 3 years, 1 month ago

C and D is correct.

"If an input firewall filter is configured on the same logical interface as a policer, the policer is executed first."

If an output firewall filter is configured on the same logical interface as a policer, the firewall filter is executed first."

upvoted 4 times

🗨️ **definity** 3 years, 1 month ago

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/policer-order-of-operations-with-firewall-filters.html>

upvoted 4 times

```

user@router> show bgp neighbor 192.168.100.2
Peer: 192.168.100.2+50862 AS 65512 Local: 192.168.100.1+179 AS 65512
  Group: INT                               Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Refresh>
  Options: <GracefulShutdownRcv>
  Local Address: 192.168.100.1 Holdtime: 90 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Number of flaps: 0
  Peer ID: 192.168.100.2   Local ID: 192.168.100.1   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0   SNMP index: 3
  I/O Session Thread: bgpio-0   State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet-vpn-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65512)
  Peer does not support Addpath
  NLRI(s) enabled for color nexthop resolution: inet-unicast
...

```

Referring to the exhibit, the local BGP router is receiving IPv4 routes from the BGP neighbor, but it is not receiving L3 VPN routes from the BGP neighbor.

Which two actions should you take to solve this problem? (Choose two.)

- A. Configure the family inet-vpn unicast statement on the BGP neighbor.
- B. Configure the family inet unicast statement on the local BGP router.
- C. Configure the family inet-vpn unicast statement on the local BGP router.
- D. Configure the family inet unicast statement on the BGP neighbor.

Suggested Answer: AC

Community vote distribution

BC (100%)

 **Juniperguy** Highly Voted 3 years, 1 month ago

Agree. BC
upvoted 7 times

 **TriviumGG** Most Recent 1 year, 1 month ago

Selected Answer: BC

A and D is incorrect, as the output of NLRI clearly shows that the neighbor has the relevant commands configured.
upvoted 1 times

 **NetworkGuy420** 1 year, 4 months ago

Selected Answer: BC

B and C

The output clearly shows how the remote peer is advertising inet-unicast and inet-vpn-unicast. So this basically discards answers A and D for changes in the remote end.

upvoted 1 times

🗨️ 👤 **NetworkGuy420** 1 year, 4 months ago

B and C.

The output shows how the peer is advertising inet-vpn-unicas and inet-unicast. Just that discards changes in the remote end, so A and D are discarded.

upvoted 1 times

🗨️ 👤 **mohdema** 1 year, 9 months ago

Selected Answer: BC

It is clearly shown in the exhibit that inet-unicast is the only NLRI enabled locally, so inet-vpn-unicast needs to be configured too. However, when you do this, it overrides the default NLRI for a peer which is inet-unicast, resulting in only inet-vpn-unicast being advertised to the peer. To continue using inet-unicast this must now also be configured explicitly.

upvoted 1 times

🗨️ 👤 **Panadol** 2 years, 2 months ago

The combination AD or BC both will solve the problem. The point is you need to enable both address families either at BGP(global) , Group or Neighbor level.

upvoted 1 times

🗨️ 👤 **EBNPW9** 2 years, 6 months ago

Selected Answer: BC

abcxyz12 is right. Totally agree.

upvoted 2 times

🗨️ 👤 **oceans1908** 2 years, 7 months ago

B,C peer router send both families

upvoted 4 times

🗨️ 👤 **abcxyz12** 3 years, 2 months ago

This should be A,B.

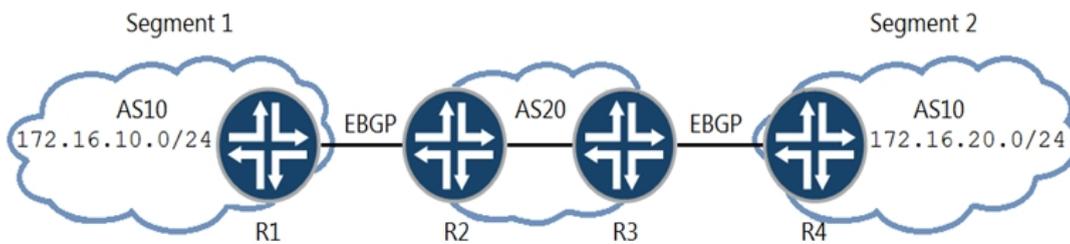
It is clearly shown in the exhibit that inet-unicast is the only NLRI enabled locally, so inet-vpn-unicast needs to be configured too. However, when you do this, it overrides the default NLRI for a peer which is inet-unicast, resulting in only inet-vpn-unicast being advertised to the peer. To continue using inet-unicast this must now also be configured explicitly.

upvoted 2 times

🗨️ 👤 **abcxyz12** 3 years, 2 months ago

Sorry, B,C.

upvoted 13 times



Your network connects two segments of your customer's network as shown in the exhibit. They need to exchange routes between Segment 1 and Segment 2 but both segments use the same AS number.

Which two steps will accomplish this task? (Choose two.)

- A. Configure the routing-options autonomous-system loops 1 parameter on routers R1 and R4.
- B. Configure the BGP group with the advertise-peer-as parameter on routers R2 and R3.
- C. Configure the routing-options autonomous-system loops 1 parameter on routers R2 and R3.
- D. Configure the BGP group with the as-override parameter on routers R1 and R4.

Suggested Answer: AB

somanyquestions Highly Voted 2 years, 10 months ago

AB

on A however I think it should be "routing-options autonomous-system loops 2" on routers R1 and R4.

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/autonomous-systems.html#id-example-enabling-bgp-route-advertisements>

upvoted 5 times

magmartin 2 years, 1 month ago

from the explanation in the juniper web you are right, it suggests that you have to configure "loops (n+1)" if you want to allow "n" occurrences of your AS number in the received AS path. But I have checked this in the lab and it works if you configure "loops n" to allow "n" occurrences of your AS number. So, I think is a mistake on the wording in the juniper web.

upvoted 1 times

oceans1908 Most Recent 2 years, 7 months ago

Correct

upvoted 2 times

minmon_6789 2 years, 11 months ago

A and B are correct. We need to do both of them.

upvoted 2 times

definity 3 years, 1 month ago

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/autonomous-systems.html#id-example-enabling-bgp-route-advertisements>

upvoted 2 times

Totalstranger 2 years ago

this works

Routing-option as loop - not working

BGP family inet loop - working

upvoted 1 times

You have a mixed vendor EVPN environment and you need to ensure VXLAN interoperability between all devices. In this scenario, which statement is correct?

- A. You should only use pure Type 2 routes.
- B. You should only use pure Type 5 routes.
- C. You should only use Type 2 and Type 5 routes.
- D. You should only use Type 6 and Type 2 routes.

Suggested Answer: C

🗨️ 👤 **Sparks026** 1 year, 10 months ago

Option C is correct.

<https://blogs.juniper.net/en-us/service-provider-transformation/juniper-delivers-evpn-vxlan-data-center-fabric-overlays-with-multi-vendor-interoperability>

upvoted 1 times

🗨️ 👤 **somanyquestions** 2 years, 10 months ago

should be B; <https://blogs.juniper.net/en-us/service-provider-transformation/juniper-delivers-evpn-vxlan-data-center-fabric-overlays-with-multi-vendor-interoperability> "The recommended approach of using EVPN Type 5 routes for multi-vendor fabric environments is possible due to the fact that it has fewer dependencies on other EVPN route types and because most vendors have adopted precisely the same Type 5 route IP-VRF-to-IP-VRF interface-less service model."

upvoted 2 times

🗨️ 👤 **Rdpk** 2 years, 8 months ago

I think C is correct,

1)"With IETF standards-based EVPN-VXLAN on Junos software, this can be achieved even when remote locations are located behind multiple IP domains or connected to different vendors supporting EVPN Type 2 (VLAN-aware bundle service) model or EVPN Type 5 (prefix advertisement) model"

2)".....in order to benefit from more advanced features such as EVPN-VXLAN stitching or pure overlay multicast. Alternatively, POD2 could be from a different vendor supporting EVPN Type 5 routes or EVPN Type 2 VLAN-aware bundle service type."

upvoted 1 times

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

It doesn't say all vendors however, so the answer is C.

upvoted 1 times

The link between CE1 and PE1 has a history of flapping. To avoid the impact that flapping causes to the network, you decide to use route damping.

Which statement is correct in this scenario?

- A. Dampening is enabled on interfaces.
- B. Dampened routes decay at a sliding rate known as half-life.
- C. Routes become dampened when the configured max-suppress value is reached.
- D. Dampened routes become active when their figure of merit drops below the reuse value.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **BobbyAxelrod** 2 years, 4 months ago

Selected Answer: D

Answer is D

<https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/ref/command/show-policy-damping.html>

Reuse merit

Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.

upvoted 2 times

🗳️ **cartelluminoso** 2 years, 5 months ago

B is NOT correct because the figure of merit do NOT decay at sliding rate, decays exponentially. Thus the correct answer is D.

upvoted 2 times

🗳️ **austin_powers** 2 years, 7 months ago

B is correct, however D is also correct in my opinion. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.

upvoted 2 times

🗳️ **Ad_Latjes** 2 years, 7 months ago

D : Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.

upvoted 1 times

🗳️ **oceans1908** 2 years, 7 months ago

B is correct

upvoted 1 times

🗳️ **Anniesyed1234** 2 years, 10 months ago

B is the correct answer

upvoted 1 times

🗳️ **Anniesyed1234** 2 years, 10 months ago

Sorry Correct Answer is B

upvoted 1 times

🗳️ **Anniesyed1234** 2 years, 10 months ago

Correct Answer : A

upvoted 1 times

🗳️ **definity** 3 years, 1 month ago

https://www.juniper.net/documentation/en_US/junos/topics/concept/routing-policy-security-damping-parameter-understanding.html
upvoted 1 times

You are deploying a new EVPN service for your customers.

You must build the service based on the following requirements:

- both Layer 2 and Layer 3 functionality must be supported;
- your customers must be able to support multiple VLANs in the same EVPN instance (EVI).

In this scenario, which two types of routing instances should be configured? (Choose two.)

- A. virtual switch
- B. virtual router
- C. VRF
- D. EVPN

Suggested Answer: CD

Community vote distribution

AC (100%)

 **oceans1908** Highly Voted 2 years, 7 months ago

Two instances 1 type virtualswitch and 1 type vrf
upvoted 5 times

 **oceans1908** 2 years, 7 months ago

On later inside VRF and EVPN instance type
upvoted 2 times

 **jncie_examer** Most Recent 1 year, 11 months ago

Selected Answer: AC

Virtual switch instance for multi vlan, vrf instance for l3 vpn
upvoted 1 times

 **OGab** 2 years, 1 month ago

Selected Answer: AC

the reason why AC is because EVPN only supports a single VLAN, therefore since you need to support multiple vlans you must use a virtual switch.
upvoted 3 times

 **Hiwanku** 2 years ago

Starting with Junos OS Release 17.1, VLAN bundle service allows multiple broadcast domains to map to a single bridge domain.
<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-bundled-services.html>
upvoted 2 times

 **EBNPW9** 2 years, 6 months ago

CD

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/task/evpn-irb-configuring.html>
upvoted 3 times

 **sught** 2 years, 10 months ago

CD are correct

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/concept/evpn-bundled-services.html>
upvoted 3 times

 **hrum** 2 years, 7 months ago

NOTE: Integrated routing and bridging (IRB) is not supported for VLAN bundle service.
upvoted 3 times

 **somanyquestions** 2 years, 10 months ago

AC

set routing-instances <name> instance-type virtual-switch
set routing-instances <name> instance-type vrf

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/topic-map/evpn-virtual-switch-configuring.html>

upvoted 4 times

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

- A. Level 1 intermediate systems exchange routing information with Level 1 intermediate systems on other IS-IS areas.
- B. An IS-IS router sets the attached bit in the PDUs it sends to a Level 1 area to indicate that it is a backbone router.
- C. A Level 1 router can only form adjacencies with other Level 1 routers.
- D. Level 2 routers can form adjacencies with either Level 1 or Level 2 routers.

Suggested Answer: BC

Community vote distribution

BC (67%)

CD (33%)

🗨️ **ztw3587t** 1 year, 3 months ago

Selected Answer: BC

Just correating, B and C are correct:

A Level 1/Level 2 system sets the attached bit in the Level 1 PDUs that it generates into a Level 1 area to indicate that it is a Level 2-attached backbone router and that it can be used to reach prefixes outside the Level 1 area.

Only L1/L2 routers can be form adjacency with L1 and L2 routers.

upvoted 2 times

🗨️ **ztw3587t** 1 year, 3 months ago

Selected Answer: CD

CD ia correct

upvoted 1 times

🗨️ **Anniesyed1234** 2 years, 10 months ago

B& C correct

upvoted 1 times

🗨️ **Ramy_Halim** 2 years, 11 months ago

correct

upvoted 1 times

Which two statements are correct about Opaque LSAs in OSPF? (Choose two.)

- A. Type 10 LSAs are used for MPLS traffic-engineering and have area scope.
- B. Type 11 LSAs are used for MPLS traffic-engineering and have area scope.
- C. Type 11 LSAs are used for MPLS label exchange and have link-local scope
- D. Type 9 LSAs are used for graceful-restart and have link-local scope

Suggested Answer: AD

Community vote distribution

AD (100%)

EBNPW9 2 years, 6 months ago

9 - link-local scope, 10 - area scope, 11 - domain (AS) scope. so, A&D
upvoted 2 times

Brabus 2 years, 6 months ago

Selected Answer: AD

A&D as junos doesn't support 11 lsa.
upvoted 3 times

oceans1908 2 years, 7 months ago

A,D as type 11 is not supported by junos
upvoted 2 times

ak_sisko 3 years ago

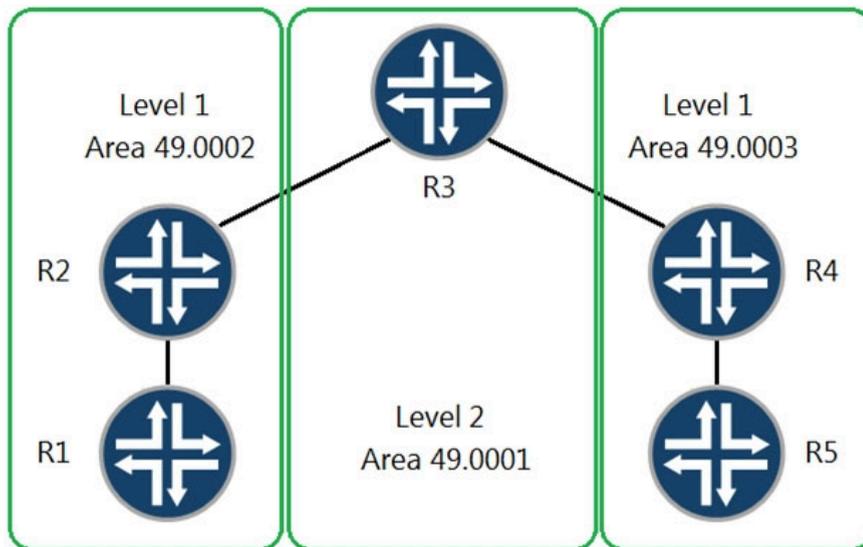
Correct A,D
upvoted 1 times

Juniperguy 3 years, 1 month ago

Type 11 lsa is not used on junos
upvoted 4 times

juanangel128 3 years ago

Agree with u
upvoted 1 times



All adjacencies have been formed, no extra options have been configured, and no policies have been written. Referring to the exhibit, which two statements are correct? (Choose two.)

- A. R2 will create a default route and send it as a TLV to R1
- B. R1 cannot reach R5
- C. R1 can reach R5
- D. R1 will create its own default route that points to R2

Suggested Answer: AC

Community vote distribution

CD (100%)

- 🗨️ **Juniperguy** Highly Voted 3 years, 1 month ago
 CD. R1 sees attach bit in Level1 LSP and it installs default route in routing table.
 Also i dont see any reason why R1 cant reach R5
 upvoted 12 times
- 🗨️ **somanyquestions** 2 years, 10 months ago
https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html#routing-is-is-overview_d8374e173 L1 router installs it is correct
 upvoted 3 times
- 🗨️ **ztw3587t** Most Recent 1 year, 3 months ago
Selected Answer: CD
 C and D are correct:
 The R2 is a L1/L2 router and because of that it sets the attached bit in the Level 1 PDUs and the R1 Router (L1) generates its own default route that points to R2. As a result R1 reaches R5 router and another routers out of L1 area.
 upvoted 1 times
- 🗨️ **ztw3587t** 1 year, 3 months ago
Selected Answer: CD
 CD ia correct
 upvoted 1 times
- 🗨️ **NetworkGuy420** 1 year, 4 months ago
Selected Answer: CD
https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/concept/is-is-routing-overview.html#routing-is-is-overview_d8374e173
 upvoted 1 times
- 🗨️ **DPK1001** 1 year, 10 months ago
 C and D are correct

upvoted 1 times

🗨️ 👤 **Panadol** 2 years, 2 months ago

CD is correct

upvoted 1 times

🗨️ 👤 **EBNPW9** 2 years, 6 months ago

Selected Answer: CD

Juniperguy and somanyquestions are totally correct.

upvoted 1 times

🗨️ 👤 **oceans1908** 2 years, 7 months ago

CD, as a L1 chooses the best TLV with attached bit set as the next hop for the default

upvoted 1 times

🗨️ 👤 **vegass** 2 years, 8 months ago

Selected Answer: CD

cd is correct

upvoted 1 times

🗨️ 👤 **nguyen03t2** 2 years, 11 months ago

CD is correct

upvoted 1 times

🗨️ 👤 **anonymonkey** 3 years, 1 month ago

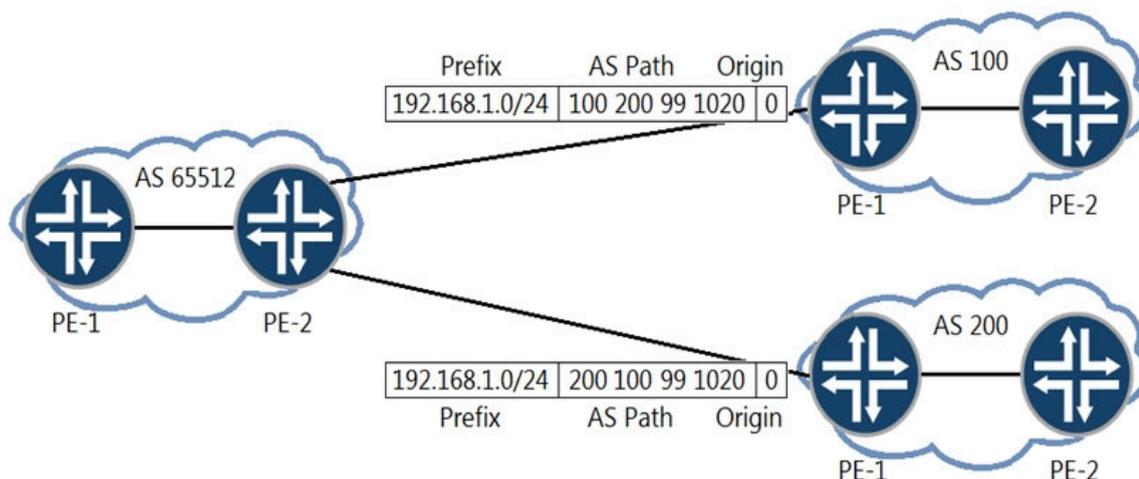
<https://www.juniper.net/documentation/us/en/software/junos/isis/topics/example/logical-systems-isis-policy-default-route.html>

upvoted 1 times

🗨️ 👤 **Sheet** 3 years, 1 month ago

R1 will create it's own default route

upvoted 2 times



You are the administrator of AS 65512. You are learning the 192.168.1.0/24 prefix from both AS 100 and AS 200. You want traffic destined to the 192.168.1.0.0/24 prefix to exit your AS towards AS 200.

How would you accomplish this task?

- A. Configure an import routing policy on PE-2 to set a higher MED on the path learned from AS 100.
- B. Configure an import routing policy on PE-2 to modify the origin attribute on the path learned from AS 100.
- C. Configure an import routing policy on PE-2 to set a higher local preference value on the path learned from AS 200.
- D. Configure an import routing policy on PE-2 to append the AS path attribute on the path learned from AS 100.

Suggested Answer: C

ak_sisko Highly Voted 3 years ago

Correct answer is C
upvoted 6 times

[Removed] Most Recent 1 year, 4 months ago

A - probably should word but in practice in lab does not seem to
B - Yes will work proven in lab
C- Will work
D- As Path Prepending
So A-Maybe , and B,C,D are all valid answers
upvoted 1 times

TriviumGG 1 year, 1 month ago

First I was thinking of MED, but it is used to tell the other AS's how to enter your network.
Since the remote AS are different they won't see the difference in metrics between the two routes with different MED.
upvoted 1 times

NinjaCloud 2 years, 3 months ago

Local preference works only for iBGP sessions. Correct answer is D
upvoted 1 times

Ahmed2021 2 years, 3 months ago

Local preference is working for internal route only not for the route learned between ASs so the correct answer is D
upvoted 2 times

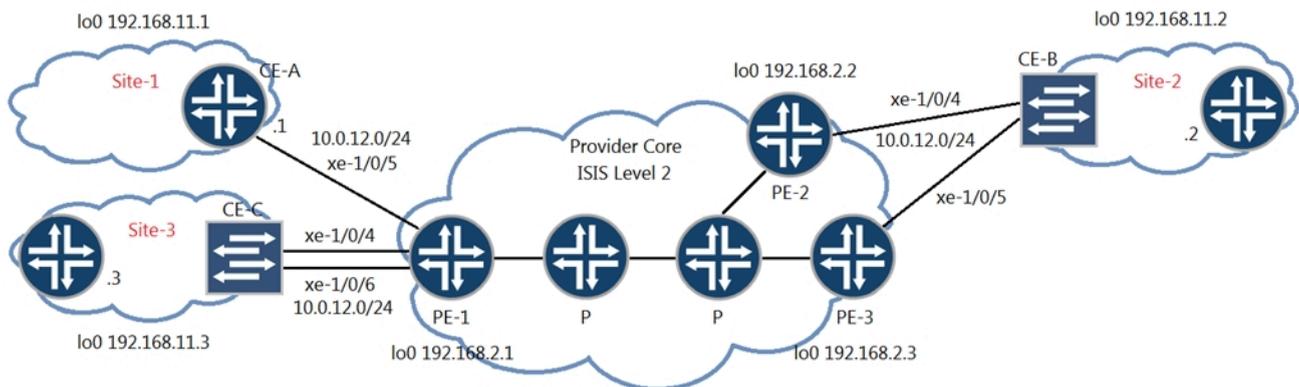
magmartin 2 years, 1 month ago

That's not right. Local preference is only sent to iBGP peers, but this parameter works for every route, either received through iBGP or eBGP, as is one of the first tiebreakers used in the BGP path selection process.
As routes received through an eBGP neighbor does not include this parameter, a default value of 100 is assigned, but this can be modified using an import policy.
upvoted 4 times

  **oceans1908** 2 years, 7 months ago

C is correct

upvoted 2 times



You have the LDP signaled VPLS topology as shown in the exhibit. CE-B at Site-2 is multihomed to both PE-2 and PE-3. In this scenario, where would you configure loop prevention?

- A. PE-1
- B. CE-B
- C. PE-3
- D. PE-2

Suggested Answer: A

Community vote distribution

A (80%)

B (20%)

manafo 1 year, 12 months ago

Selected Answer: A

https://supportportal.juniper.net/s/article/MX-Example-Configuring-LDP-based-VPLS-using-primary-backup-for-redundancy?language=en_US
upvoted 1 times

Panadol 2 years, 2 months ago

A is the correct answer
upvoted 1 times

Ahmed2021 2 years, 3 months ago

B is correct

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpn-vpls-multihoming.html>
upvoted 1 times

cartelluminoso 2 years, 5 months ago

Selected Answer: A

The correct answer is A. In LDP signaled VPLS, when you have a multihomed CE to two different PEs you can configure a neighbor and a backup neighbor in remote PEs. In this case PE-1. If the VPLS is signaled with BGP then you have to configure the preference in PE-2 and PE-3. The spanning tree option mentioned in other comments needs the configuration of spanning tree in the CE and you also need to configure a l2control instance in PEs attached to the CE. So B is not correct.
upvoted 3 times

mohdema 1 year, 9 months ago

What you are describing is redundancy

https://supportportal.juniper.net/s/article/MX-Example-Configuring-LDP-based-VPLS-using-primary-backup-for-redundancy?language=en_US
upvoted 1 times

mohdema 1 year, 9 months ago

In our case it seems automatic

Remote PE routers in the VPLS network need to determine which of the multihomed PE routers should forward traffic to reach the CE device. To make this determination, remote PE routers use the VPLS path-selection process to select one of the multihomed PE routers based on its NLRI advertisement. Because remote PE routers pick only one of the NLRI advertisements, it establishes a pseudowire to only

one of the multihomed PE routers,

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpn-vpls-multihoming.html>

upvoted 1 times

  **EBNPW9** 2 years, 6 months ago

Selected Answer: B

I agree with Rdpk's link.

upvoted 2 times

  **oceans1908** 2 years, 6 months ago

Selected Answer: B

upvoted 1 times

  **zaibac** 2 years, 6 months ago

Selected Answer: A

Remote PE routers in the VPLS network need to determine which of the multihomed PE routers should forward traffic to reach the CE device. To make this determination, remote PE routers use the VPLS path-selection process to select one of the multihomed PE routers based on its NLRI advertisement. Because remote PE routers pick only one of the NLRI advertisements, it establishes a pseudowire to only one of the multihomed PE routers, the PE router that originated the winning advertisement. This prevents multiple paths from being created between sites in the network, preventing the formation of Layer 2 loops. If the selected PE router fails, all PE routers in the network automatically switch to the backup PE router and establish new pseudowires to it.

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/topic-map/vpls-bgp-multihoming.html>

upvoted 4 times

  **chris_2_a** 2 years, 5 months ago

Is LDP VPLS and not BGP signaled VPLS so the A is the correct one.

upvoted 1 times

  **oceans1908** 2 years, 7 months ago

B should it b

upvoted 1 times

  **Rdpk** 2 years, 8 months ago

BEST PRACTICE: To prevent the formation of Layer 2 loops between the CE devices and the multihomed PE routers, we recommend that you employ the Spanning Tree Protocol (STP) on your CE devices. Layer 2 loops can form due to incorrect configuration. Temporary Layer 2 loops can also form during convergence after a change in the network topology.

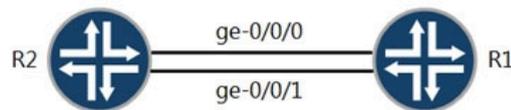
<https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/topic-map/vpls-bgp-multihoming.html>

upvoted 2 times

  **Mel7IT** 2 years, 8 months ago

B is correct

upvoted 2 times



```

user@R2# show interfaces
ge-0/0/0 {
  unit 0 {
    family iso;
    family inet6 {
      address 2001:db8::2/64;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 172.16.2.2/24;
    }
    family iso;
    family inet6 {
      address 2001:db8:1::2/64;
    }
  }
}
user@R2# show protocols isis
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface lo0.0 {
  passive;
}
topologies ipv6-unicast;

```

```

user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family iso;
    family inet6 {
      address 2001:db8::1/64;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family iso;
    family inet6 {
      address 2001:db8:1::1/64;
    }
  }
}
user@R1# show protocols isis
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface lo0.0 {
  passive;
}
topologies ipv6-unicast;

```

A network administrator is migrating from IPv4 to IPv6 and one of the IS-IS adjacencies is not coming up between R1 and R2. Which action will solve the problem?

- A. Remove topologies ipv6-unicast from protocols isis on R2.
- B. Configure topologies ipv4-unicast from protocols isis on R2.
- C. Remove topologies ipv6-unicast from protocols isis on R1.
- D. Configure an IPv4 address on interface ge-0/0/1 on R1.

Suggested Answer: D

Community vote distribution

D (100%)

ztw3587t 1 year, 3 months ago

Selected Answer: D

D is correct:

To configure IS-IS dual stacking:

Configure the interfaces, including both IPv4 and IPv6 addresses on each interface.

upvoted 1 times

ztw3587t 1 year, 3 months ago

Selected Answer: D

A is correct:

To configure IS-IS dual stacking:

Configure the interfaces, including both IPv4 and IPv6 addresses on each interface.

upvoted 1 times

NetworkGuy420 1 year, 4 months ago

Selected Answer: D

D

<https://www.juniper.net/documentation/us/en/software/junos/is-is/topics/example/isis-ipv6-dual-stacking.html>

upvoted 1 times

  **12_sdeEQ** 2 years, 10 months ago

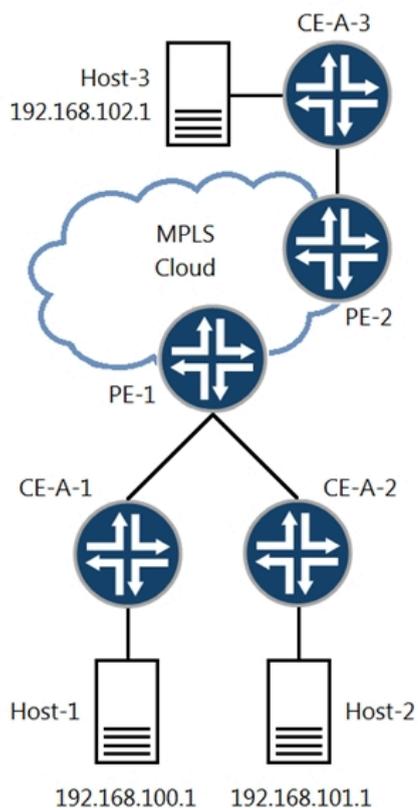
D is the correct answer.

upvoted 1 times

  **minmon_6789** 2 years, 11 months ago

D is the correct answer.

upvoted 2 times



```
[edit routing-instances]
user@PE-1# show
CE-A-1 {
  instance-type vrf;
  interface ge-0/0/9.0;
  route-distinguisher 10.222.222.4:1;
  vrf-target target:65511:101;
  routing-options {
    static {
      route 192.168.100.0/24
    }
  }
  next-hop 192.168.0.2;
}
CE-A-2 {
  instance-type vrf;
  interface ge-0/0/8.0;
  route-distinguisher 10.222.222.4:3;
  vrf-target target:65511:101;
  routing-options {
    static {
      route 192.168.101.0/24
    }
  }
  next-hop 192.168.1.2;
}
}
```

Referring to the exhibit, there is a Layer 3 VPN setup that connects sites CE-A-1, CE-A-2, and CE-A-3 together. Host-1 can communicate with Host-3, but Host-1 cannot communicate with Host-2.

What must you do to solve the problem?

- A. Change the route distinguisher in both routing instances to the same value.
- B. Use the next-table configuration statement for static routes in the corresponding routing instances.
- C. Use BGP instead of static routing between the CE and PE devices.
- D. Use the auto-export command in both routing instances.

Suggested Answer: D

Community vote distribution

D (100%)

NetworkGuy420 1 year, 4 months ago

Selected Answer: D

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

auto-export:

This statement enables you to leak routes between VPN routing and forwarding (VRF) instances that are locally configured on a provider edge (PE) router. Auto export is always applied on the local PE router, because it applies to only local prefix leaking by evaluating the export policy of each VRF and determining which route targets can be leaked. The standard VRF import and export policies affect remote PE prefix leaking.

You can use this statement as an alternative to using the VRF import and export policies.

upvoted 1 times

kasqureshi 2 years, 4 months ago

Selected Answer: D

D is the answer

upvoted 1 times

somanyquestions 2 years, 10 months ago

https://kb.juniper.net/library/CUSTOMERSERVICE/GLOBAL_JTAC/auto-export-understanding.pdf

upvoted 1 times

```

user@host> show pim join 234.100.0.1 extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 234.100.0.1
  Source: 192.168.100.2
  Flags: sparse, spt
  Active upstream interface: ge-1/0/0.0
  Active upstream neighbor: 192.168.101.2
  MoFRR Backup upstream interface: ge-1/0/1.0
  MoFRR Backup upstream neighbor: 192.168.102.2
  Upstream state: Join to Source, No Prune to RP
  Keepalive timeout: 300
  Uptime: 00:00:15
  Downstream neighbors:
    Interface: ge-1/2/0.0
      192.168.103.2 State: Join Flags: S Timeout: Infinity
      Uptime: 00:00:15 Time since last Join: 00:00:15
  Number of downstream interfaces: 1

```

Which three statements are true about the show pim join output shown in the exhibit? (Choose three.)

- A. This is a source-specific multicast stream.
- B. The multicast receiver is still using the RP to receive the stream.
- C. The multicast stream has been configured with a backup path to allow for fast reroute.
- D. The multicast stream does not have an RP.
- E. The shortest path to the source is through the RP.

Suggested Answer: CDE

Community vote distribution

BCE (85%)

Other

🗨️ **ztw3587t** 1 year, 3 months ago

Selected Answer: CDE

A and B are a single answer and that one is wrong. Therefore the correct are C,D,E.
upvoted 1 times

🗨️ **mohdema** 1 year, 9 months ago

The Internet Assigned Numbers Authority (IANA) has reserved the address range 232.0. 0.0 through 232.255. 255.255 for SSM applications and protocols.
upvoted 1 times

🗨️ **mohdema** 1 year, 9 months ago

Selected Answer: BCE

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.
upvoted 1 times

🗨️ **Banand** 2 years, 2 months ago

Answer BCE agreed with ocean1908
upvoted 2 times

🗨️ **EBNPW9** 2 years, 6 months ago

Selected Answer: ACE

A <= Source: 192.168.100.2; Flags: spt; Upstream state: Join to Source
C <= MoFRR Backup upstream interface/neighbor

E <= Upstream state: No Prune to RP

Also see links in anonymonkey's comment.

upvoted 1 times

  **ptt** 2 years, 6 months ago

100% not ACE, because the multicast section i got zero percent correct after i chose ACE for this question.

upvoted 2 times

  **oceans1908** 2 years, 6 months ago

Selected Answer: BCE

Bce as the other two are wrong

upvoted 4 times

  **oceans1908** 2 years, 6 months ago

A and D are wrong so that leaves BCE

upvoted 1 times

  **junos_wizz** 2 years, 7 months ago

Selected Answer: BCE

agree with boyseven777

upvoted 3 times

  **boyseven777** 2 years, 7 months ago

Selected Answer: BCE

not A - SSM address space is 232.0.0.0/8; not D - sparse mode implicate presence of RP

upvoted 3 times

  **austin_powers** 2 years, 7 months ago

The answer is A, C and E.

PIM SSM can technically be used in the entire 224/4 multicast address range

upvoted 1 times

  **Anniesyed1234** 2 years, 10 months ago

A , C and E

upvoted 2 times

  **ak_sisko** 3 years ago

Agree, ACE

upvoted 1 times

  **juanangel128** 3 years ago

Exactly, agree with you, correct answers are ACE.

upvoted 1 times

  **anonymonkey** 3 years, 1 month ago

ACE

<https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/ref/command/show-pim-join.html>

<https://kb.juniper.net/InfoCenter/index?page=content&id=KB21903>

<https://kb.juniper.net/InfoCenter/index?page=content&id=KB21586&actp=METADATA>

upvoted 3 times

Why do interprovider option B VPNs scale better than interprovider option A VPNs?

- A. The ASBRs in interprovider option B VPNs do not need per-VPN VRF tables.
- B. The ASBRs in interprovider option A VPNs do not need per-VPN VRF tables.
- C. The ASBRs in interprovider option A VPNs only carry internal routes.
- D. The ASBRs in interprovider option B VPNs only carry internal routes.

Suggested Answer: A

Community vote distribution

A (100%)



 **oceans1908** 2 years, 6 months ago

Selected Answer: A

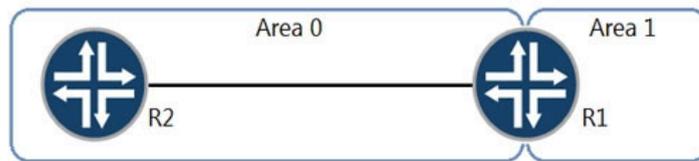
Answer A

upvoted 2 times

 **Anniesyed1234** 2 years, 10 months ago

A is correct

upvoted 2 times



```
users@R1> show ospf3 database inter-area-prefix detail
```

```

OSPF3 database, Area 0.0.0.0
Type      ID          Adv Rtr      Seq          Age   Cksum   Len
InterArPfx 0.0.0.11     172.16.1.1  0x80000001   4    0xaa9a  36
  Prefix 2001:db9:ffff:ff00::/64
  Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.12     172.16.1.1  0x80000001   4    0x8c6e  44
  Prefix 2001:db9:ffff:ff00::1/128
  Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.13     172.16.1.1  0x80000001   4    0xa899  36
  Prefix 2001:db9:ffff:ff01::/64
  Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.14     172.16.1.1  0x80000001   4    0x8a6d  44
  Prefix 2001:db9:ffff:ff01::1/128
  Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.15     172.16.1.1  0x80000001   4    0xa698  36
  Prefix 2001:db9:ffff:ff02::/64
  Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.16     172.16.1.1  0x80000001   4    0x886c  44
  Prefix 2001:db9:ffff:ff02::1/128
  Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.17     172.16.1.1  0x80000001   4    0xa497  36
  Prefix 2001:db9:ffff:ff03::/64
  Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.18     172.16.1.1  0x80000001   4    0x866b  44
  Prefix 2001:db9:ffff:ff03::1/128
  Prefix-options 0x0, Metric 0

```

Referring to the exhibit, which command would reduce the size of the OSPF database and corresponding routes?

A.

```

user@R1# show protocols ospf3
area 0.0.0.1 {
  area-range 2001:db9:ffff:ff00::/62;
}

```

B.

```

user@R1# show policy-options policy-statement summary-2001
term 10 {
  from {
    route-filter 2001:db9:ffff:ff00::/62 prefix-length-range /64-/128;
  }
  then accept;
}
user@R1# show protocols ospf3
area 0.0.0.0 {
  inter-area-prefix-import summary-2001;
}

```

C.

```
user@R1# show policy-options policy-statement summary-2001
term 10 {
  from {
    route-filter 2001:db9:ffff:ff00::/62 prefix-length-range /64-/128;
  }
  then accept;
}
user@R1# show protocols ospf3
area 0.0.0.1 {
  inter-area-prefix-export summary-2001;
}
```

D.

```
user@R1# show protocols ospf3
area 0.0.0.1 {
  stub no-summaries;
}
```

Suggested Answer: A

 **Dlbyam** 1 year, 10 months ago

D is correct, area-range works within same area
upvoted 1 times

 **EBNPW9** 2 years, 6 months ago

A is not technically correct (should be area 0)
B and C will not reduce the number of routes <https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/topic-map/configuring-ospf-routing-policy.html#id-import-and-export-policies-for-network-summaries-overview>
D is obvious and right solution
upvoted 1 times

 **nyanko** 2 years, 6 months ago

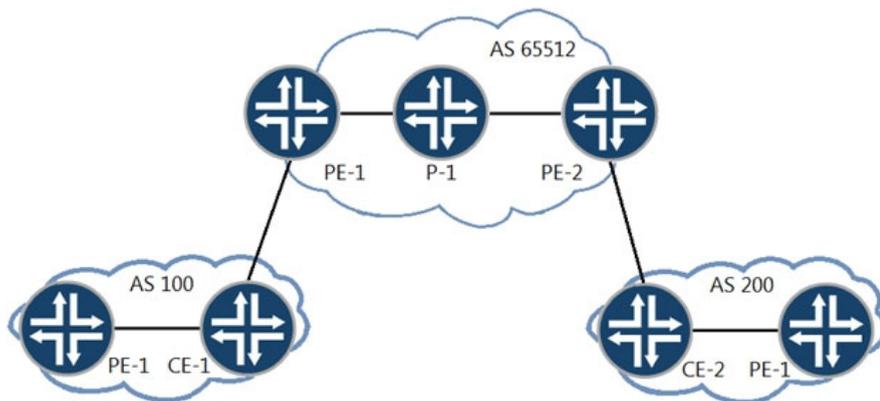
The show command displays Type3 LSA routes in area 0 database. It means those routes are originated in area 1. Hence, A can reduce those Type3 LSAs into only one Type3 LSA in area0 database.
upvoted 6 times

 **minmon_6789** 2 years, 11 months ago

A is correct.
upvoted 1 times

 **hcccc** 3 years, 3 months ago

<https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/topic-map/configuring-ospf-route-control.html>
upvoted 1 times



You are providing carrier-of-carrier VPN services for AS 100 and AS 200. You want to distribute MPLS labels between your PE routers and the AS 100 and AS 200 CE routers.

What must be enabled to accomplish this task?

- A. Use BGP with the inet-vpn address family enabled.
- B. Use BGP with the labeled-unicast address family enabled.
- C. Use RSVP with the lsp-set parameter enabled.
- D. Use RSVP with the tunnel-services parameter enabled.

Suggested Answer: A

Community vote distribution

B (75%)

A (25%)

vasiliy19 Highly Voted 3 years, 2 months ago

B

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-carrier-carrier.html>
upvoted 14 times

zaibac 2 years, 6 months ago

Question says "You are providing carrier-of-carrier VPN services for AS 100 and AS 200"

You are in charge of the ISP network only, no more
upvoted 1 times

ptt 2 years, 6 months ago

provider's PE need labeled-unicast to extend its LSPs to the customer CE router. Additionally, inet-vpn is for the multihop MB-EBGP session between customer's PE-1 and customer's PE-2
upvoted 2 times

ptt 2 years, 6 months ago

I mean the answer is B if the role(provider of the middle) is providing carrier-of-carrier VPN service.
upvoted 1 times

juanangel128 3 years ago

```

bgp {
  group group-name {
    export internal;
    peer-as as-number;
    neighbor address {
      family inet {
        labeled-unicast; <<<<<<----- here you are
      }
    }
  }
}

```

```
}
```

```
}
```

upvoted 4 times

  **muhammadradi** Most Recent 2 years ago

Selected Answer: A

A is correct as here asked about configuration on PE not CEs

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-carrier-carrier.html>

upvoted 1 times

  **Banand** 2 years, 2 months ago

As we are service provider we need to configure PE router -

Link : - https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-carrier-carrier.html#id-configuring-carrier-of-carriers-vpns-for-customers-that-provide-vpn-service_d157e485

To configure the routing instance on the provider's PE router (or switch) to send labeled routes to the carrier customer's CE router (or switch), include the following statements:

```
instance-type vrf;
interface interface-name;
route-distinguisher value;
vrf-import policy-name;
vrf-export policy-name;
protocols {
  bgp {
    group group-name {
      peer-as as-number;
      neighbor address {
        family inet {
          labeled-unicast;
        }
      }
    }
  }
}
```

upvoted 1 times

  **Banand** 2 years, 2 months ago

Ans - B (BGP with Labeled-unicast)

upvoted 1 times

  **Crece** 2 years, 2 months ago

A is correct.

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-carrier-carrier.html>

upvoted 1 times

  **EBNPW9** 2 years, 6 months ago

Selected Answer: B

inet-vpn will be on PE-1 in AS100 and on PE-1 in AS200. That's the whole point of Carrier-of-Carriers: we (AS 65512) don't need to know end customers' (customers of carrier AS100 and carrier AS200) routes. And with labeled-unicast we know addresses of PE-1 in AS 100 and PE-1 in AS 200, so these two PE-1s can form multihop EBGp session with each other. Also see vasilii19's answer.

upvoted 3 times

  **oceans1908** 2 years, 6 months ago

Selected Answer: B

Labeled unicast under bgp

upvoted 3 times

  **zaibac** 2 years, 6 months ago

Selected Answer: A

Question says "You are providing carrier-of-carrier VPN services for AS 100 and AS 200"

You are in charge of the ISP network only, no more

upvoted 1 times

  **Sct38** 2 years, 8 months ago

Miss-leading questions. Actual A+B. Need inet-vpn between instances and label-unicast between PE and CE.
upvoted 1 times

Which two statements regarding Ethernet segments (ES) are correct? (Choose two.)

- A. The Type-4 EVPN route will be used to elect the designated forwarder for the ES.
- B. The Type-3 EVPN route will be used for the aliasing function to load-balance to the ES.
- C. The Type-1 EVPN route will indicate if the ES is all-active or single-active.
- D. The Type-2 EVPN route will indicate if there is a designated forwarder on the ES.

Suggested Answer: AC

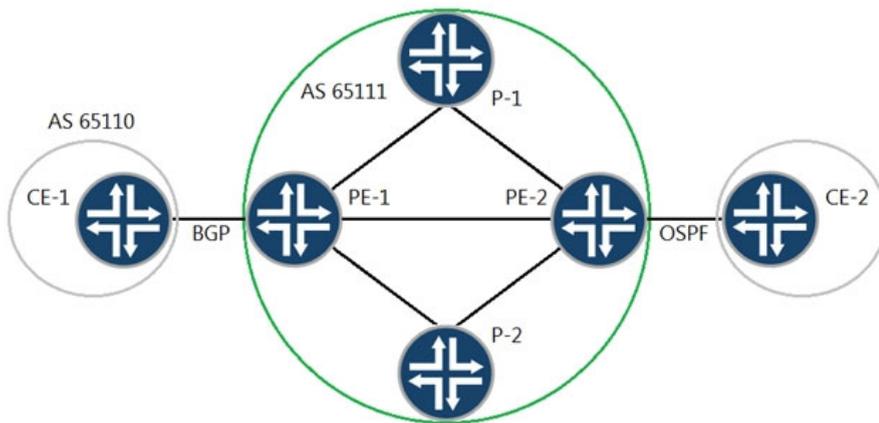
  **hcccc** 3 years, 3 months ago

Correct,

A designated forwarder (DF) is required when customer edge devices (CEs) are multihomed to more than one provider edge (PE) device. Without a designated forwarder, multihomed hosts would receive duplicate packets. Designated forwarders are chosen for an Ethernet segment identifier (ESI) based on type 4 route advertisements.

<https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/ref/statement/designated-forwarder-election-hold-time-evpn.html>

upvoted 2 times



You have a Layer 3 VPN established between PE-1 and PE-2 to allow communication between CE-1 and CE-2.

Referring to the exhibit, which statement is correct?

- A. You will need an OSPF import policy on PE-1 to receive the BGP routes, learned from PE-2, through the Layer 3 VPN.
- B. You will need a BGP export policy on PE-1 to redistribute the routes, learned from CE-1, through the Layer 3 VPN.
- C. You will need a VRF import policy on PE-1 to advertise the BGP routes, learned from CE-1, through the Layer 3 VPN.
- D. You will need an OSPF export policy on PE-2 to redistribute the BGP routes, learned from PE-1, through the Layer 3 VPN.

Suggested Answer: D

ak_sisko Highly Voted 3 years ago

Correct answer - D

upvoted 8 times

anonymonkey Most Recent 3 years, 1 month ago

B

<https://www.juniper.net/documentation/us/en/software/junos/vpn-l3/topics/topic-map/l3-vpns-distributing-routes.html>

upvoted 3 times

Ahmed2021 2 years, 3 months ago

Wrong, CE-1 is using BGP, so it does not need export policy to exchange its routes

upvoted 2 times

```

[edit]
user@R4# run show route hidden extensive

inet.0: 7 destinations, 7 routes (5 active, 0 holddown, 1 hidden)
11.11.11.0/24 (1 entry, 0 announced)
    BGP      Preference: 170/-101
            Next hop type: Unusable, Next hop index: 0
            Address: 0xbc4dbb4
            Next-hop reference count: 2
            State: <Hidden Int Ext>
            Peer AS: 65002
            Age: 18
            Validation State: unverified
            Task: BGP_65002_65002.22.22.22.22
            AS path: 65001 I
            Communities: no-export no-advertise
            Accepted
            Localpref: 100
            Router ID: 22.22.22.22
            Indirect next hops: 1
                Protocol next hop: 172.16.1.1
                Indirect next hop: 0x0 - INH Session ID: 0x0

[edit protocols bgp]
user@R2# show
group 65001 {
    neighbor 172.16.1.1 {
        export no-advertise;
        peer-as 65001;
    }
}
group 65002 {
    type internal;
    local-address 22.22.22.22;
    neighbor 44.44.44.44 {
        export no-advertise;
    }
}
import no-export;
export nhs;
local-as 65002;

[edit]
user@R2# show policy-options
policy-statement no-advertise {
    term 1 {
        then {
            community add no-advertise;
        }
    }
}
policy-statement no-export {
    term 1 {
        then community add no-export;
    }
}
policy-statement nhs {
    term 1 {
        then {
            next-hop self;
        }
    }
}
community no-advertise members no-advertise;
community no-export members no-export;

```

R2 is receiving a route from an EBGp neighbor and is advertising the route to R4.

Referring to the exhibit, which configuration on R2 will solve the issue with the route on R4?

- Move the no-advertise export policy from group 65002 to a global BGP policy.
- Move the nhs policy from a global BGP export policy to an export policy under group 65002.
- Move the no-export policy from a global BGP import policy to an import policy under group 65001.
- Move the no-advertise export policy from group 65001 to a global BGP policy.

Suggested Answer: B

Community vote distribution

A (60%)

B (20%)

C (20%)

🗳️ 👤 **minmon_6789** Highly Voted 2 years, 11 months ago

A is the correct answer.

B is not correct. If move NHS policy from global level to group 65002 level, this policy still not be considered because no-advertise policy is under neighbor 44.44.44.44 level and it is the most specific policy and it is only policy to be considered. If NHS policy is not considered, R2 will not change Next-hop of route when advertising to R4 >>> R4 show next-hop unusable.

upvoted 5 times

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

Agree with A

upvoted 2 times

🗳️ 👤 **nyanko** 2 years, 6 months ago

As EBGP is established between R2 and R4, when R2 advertises routes to R4, R2 will change next-hop to its interface address without any policy. The NHS should be applied for IBGP and next hop address is loopback address. If NHS is applied to global, I wonder what is next hop address when R2 advertises routes to R4.

upvoted 1 times

🗳️ 👤 **nyanko** 2 years, 6 months ago

I missed both groups have export-policy. To enable NHS policy for group 65002, no-advertise must be moved from group 65002 neighbor 44.44.44.44 to the global in order not to prioritize no-advertise. A is the correct answer.

upvoted 1 times

🗳️ 👤 **much2furious** 1 year, 7 months ago

no-advertise means that R4 will not forward this prefix to another BGP neighbor, but it can use it and install it!

B is valid, because only the most explicit policy is applied (no advertise), so nhs never gets to be applied.

<https://www.juniper.net/documentation/us/en/software/junos/routing-policy/bgp/topics/example/policy-bgp-levels.html>

upvoted 1 times

🗳️ 👤 **ztw3587t** Most Recent 1 year, 3 months ago

Selected Answer: B

B is correct because there is a problem, next hop unusable. The problem of next hop resolution is being caused because the prefix is being advertised to iBGP peer without next-hop-self action on policy. It's necessary to move nhs policy to group to solve that problem.

upvoted 1 times

🗳️ 👤 **mohdema** 1 year, 9 months ago

Selected Answer: C

I think the correct answer is C

As the router will not send it to it's peers.

If you apply the no-advertise on export, the router will send it to it's peer while requesting they do not advertise it internally and externally...

upvoted 1 times

🗳️ 👤 **guesswho123** 1 year, 9 months ago

B seems more correct,

if we look at the output of show command, protocol next hop is 172.16.1.1 and this IP is of the neighbor in the AS 65001 (in different AS), i think the question is about solving nhs issue.

upvoted 1 times

🗳️ 👤 **manafo** 2 years ago

B is correct

upvoted 2 times

🗳️ 👤 **Panadol** 2 years, 2 months ago

A is the correct answer. Tested this in lab

upvoted 1 times

🗳️ 👤 **dexjov** 2 years, 4 months ago

agree with B, no-advertise tells R4 not to distribute the route to any other peer, but R4 itself can normally install route

upvoted 1 times

🗨️ 👤 **Tomerd** 2 years, 5 months ago

Selected Answer: A

A is the correct answer.

B is not correct, if "nhs" policy move to group policy the is still more specific policy under the neighbor 44.44.44.44

test on lab.

upvoted 3 times

🗨️ 👤 **Sct38** 2 years, 8 months ago

This route is not usable due to no NHS performed when advertising across the iBGP peer. This can be corrected by either A or B. Considering NHS should only be done for iBGP peers, would choose B as best practice.

upvoted 4 times

🗨️ 👤 **manafo** 2 years ago

agree B is the most used config

upvoted 1 times

```
[edit routing-instances]
user@PE-1# show
vpn=a {
  instance-type vrf;
  interface ge-1/1/4.0;
  route-distinguisher 192.168.1.1:1;
  vrf-target target:65111:101;
  protocols {
    bgp {
      group my-ext-group {
        type external;
        peer-as 65601;
        neighbor 10.0.10.2;
      }
    }
  }
}
```

You have an established Layer 3 VPN between two PE devices. You are asked to only send certain routes from PE-1 over the VPN to the remote site while maintaining all the routes on the PE-1 device. You created a policy that matches the specific routes and then tags these routes with the appropriate target community values.

In this scenario, which configuration changes must be made to satisfy the requirement?

- A. Configure the export parameter and apply the policy to the my-ext-group BGP group configuration.
- B. Configure the vrf-export parameter and apply the policy under the edit routing-instances vpn-a hierarchy.
- C. Configure a RIB group and apply the policy as an import policy to routes distributed into the bgp.13vpn.0 routing table.
- D. Configure the import parameter and apply the policy to the my-ext-group BGP group configuration.

Suggested Answer: B

Community vote distribution

B (100%)

ztw3587t 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 1 times

Anniesyed1234 2 years, 10 months ago

B is the correct answer

upvoted 2 times