Question #1 *Topic 1*

Click the Exhibit button.

```
(65001)R1------R2------R3(65003)
[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] 00:00:50, localpref 100
                       AS path: 65001 I, validation-state: unverified
                     >  to 172.16.1.1 via ge-0/0/0.0
                     [BGP/170] 00:00:50, localpref 100
                       AS path: 65003 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 65003;
    }
}
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 configuration change will allow load balancing?

A. Configure the multipath parameter under the global BGP configuration.

B. Apply the prepend policy as an import policy under group R1.

C. Configure the multipath multiple-as parameter under the global BGP configuration.

D. Apply the prepend policy as an import policy under group R3.

Suggested Answer: *C*

Community vote distribution

C (100%)

emily_098 2 days, 21 hours ago

Selected Answer: C

C is the Correct Option

upvoted 1 times

piipo 11 months, 2 weeks ago

Selected Answer: C

C is Correct

upvoted 2 times

hbstyleboy 1 year, 2 months ago

C is correct: by default if you want to use multipath command for load balance perpose with EBGP neighbors it have to be in same AS. If you want to load balance with different EBGP neighbors (different AS) you have to include multiple-as also

upvoted 2 times

Question #2 *Topic 1*

Click the Exhibit button.

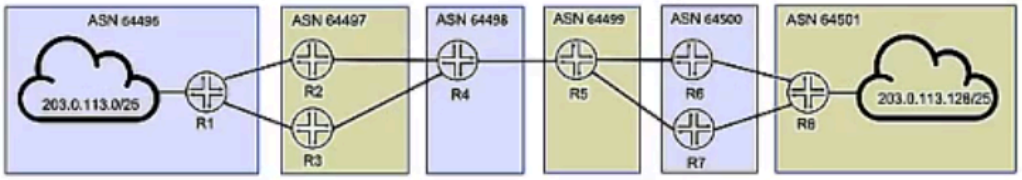


```
user@R1> show route forwarding-table matching 203.0.113.128/25
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
203.0.113.128/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.144/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.160/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.176/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.192/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.208/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.224/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
203.0.113.240/28   user     0 10.1.1.1           ucst      576    11 ge-0/0/4.0
```

You are troubleshooting the connection between AS 64496 and AS 64497 and notice that only one of the paths is being used for traffic forwarding. Referring to the exhibit, which three actions will ensure that R1 is configured properly for load balancing BGP routes? (Choose three.)

A. Verify that the routing table on R1 has BGP routes for 203.0.113.128/25 with multiple next hops.

B. Verify that the multipath option is configured under protocols bgp on both R2 and R3.

C. Verify that there is a load balancing export policy under routing-options for the received BGP routes on R1.

D. Verify that the multipath option is configured under protocols bgp on R1.

E. Verify that an import load balancing policy exists under protocols bgp for the received BGP routes on R1.

Suggested Answer: *ACD*

Community vote distribution

ACD (100%)

tsukasa123 Highly Voted 1 year, 7 months ago

Selected Answer: ACD

The requirements for R1 to load balance are A, C, D.

R2 and R3 settings are irrelevant.

upvoted 15 times

shumon1967 Most Recent 8 months, 1 week ago

ACD:

set protocols bgp group external multipath

set policy-options policy-statement loadbal then load-balance per-packet

set routing-options forwarding-table export loadbal

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

upvoted 1 times

Avee123 10 months, 1 week ago

ADE

Export Policy: In Juniper, an export policy under routing-options would typically be used to control which routes are advertised to other peers. This does not directly influence how R1 handles incoming BGP routes for load balancing.

Import Policy: For load balancing, you need to focus on import policies that affect how routes are accepted and processed. This is crucial for ensuring that multiple paths are used for forwarding traffic.

Multipath Configuration: The multipath option under protocols bgp on R1 is essential for enabling the router to use multiple paths for BGP routes.

To summarize, for Juniper devices, you should:

Verify that the routing table on R1 has BGP routes for 203.0.113.128/25 with multiple next hops (Option A).

Verify that the multipath option is configured under protocols bgp on R1 (Option D).

Verify that an import load balancing policy exists under protocols bgp for the received BGP routes on R1 (Option E).

upvoted 1 times

Guru29 1 year, 2 months ago

Correct Answer should be A,C,D

upvoted 1 times

hbstyleboy 1 year, 2 months ago

A,C, D are correct.

upvoted 1 times

M80T 1 year, 5 months ago

Selected Answer: ACD

B and E are clearly wrong - no config is needed on R2 and R3 to enable load-balancing of "received-routes on R1

upvoted 2 times

Koee 1 year, 6 months ago

So the answer should be ace?

upvoted 1 times

Question #3 *Topic 1*

After a recent power outage, your manager asks you to investigate ways to automatically reduce the impact caused by suboptimal routing in your OSPF and OSPFv3 network after devices reboot.
Which three configuration statements accomplish this task? (Choose three.)

A. set protocols ospf3 realm ipv4-unicast overload timeout 900

B. set protocols ospf overload

C. set protocols ospf overload timeout 900

D. set protocols ospf3 overload

E. set protocols ospf3 overload timeout 900

Suggested Answer: *ACE*

Community vote distribution

ACE (100%)

tsukasa123 Highly Voted 1 year, 1 month ago

Selected Answer: ACE

It says it will be reduced automatically, so need to return it after 900 seconds.

upvoted 10 times

zineeddine Most Recent 10 months, 3 weeks ago

CMD for answer E does not exist

upvoted 1 times

hbstyleboy 8 months, 1 week ago

Don't B.S pls. Tested out the cmd in CLI and it works.

upvoted 1 times

zineeddine 10 months, 3 weeks ago

ABC is the correct answer,

upvoted 1 times

Dimsop_Technology 11 months, 1 week ago

timeout seconds—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the overload statement is deleted or a timeout is set.

And the question says, reduce the impact automatically, so if we do not set time, the router will never come out of the overloaded state.

upvoted 1 times

Jicks 1 year ago

[edit logical-systems logical-system-name protocols (oospf | ospf3)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
Timeout value should be set otherwise network will not recover

upvoted 2 times

Question #4 *Topic 1*

You are configuring anycast RP for load balancing and redundancy in your PIM-SM domain. You want to share active sources between RPs.
In this scenario, what are two solutions that will accomplish this task? (Choose two.)

A. Configure MSDP on each RP router.

B. Configure anycast PIM with the rp-set statement on each RP router.

C. Configure anycast PIM with the rp-set statement on each source DR router.

D. Configure MSDP on each source DR router.

Suggested Answer: *AB*

Community vote distribution

AC (100%)

zineeddine Highly Voted 10 months, 3 weeks ago

Configure MSDP on each RP router: MSDP (Multicast Source Discovery Protocol) is used to share active multicast sources between RPs (Rendezvous Points) in different multicast domains. By configuring MSDP on each RP router, they can exchange information about active multicast sources, ensuring that all RPs are aware of the sources available. B. Configure anycast PIM with the rp-set statement on each RP router: Anycast PIM (Protocol Independent Multicast) allows multiple RPs to share the same IP address, providing load balancing and redundancy. By configuring anycast PIM with the rp-set statement on each RP router, they can function as anycast RPs, and the rp-set statement ensures that they share active sources between them.

Therefore, options A and B are the correct solutions to accomplish sharing active sources between RPs in this scenario.

upvoted 7 times

solen003 Most Recent 10 months, 3 weeks ago

I am going with A and B.

In the labs we were setting rp-set on all RPs with other RP information.

upvoted 2 times

raj1126 11 months, 3 weeks ago

Selected Answer: AC

A and C are the correct answer

upvoted 1 times

Question #5 *Topic 1*

Which two statements are correct about VPLS tunnels? (Choose two.)

A. BGP-signaled VPLS tunnels require manual provisioning of sites.

B. LDP-signaled VPLS tunnels only support control bit 0.

C. LDP-signaled VPLS tunnels use auto-discovery to provision sites.

D. BGP-signaled VPLS tunnels can use either RSVP or LDP between the PE routers.

Suggested Answer: *BD*

Community vote distribution

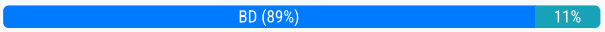

M80T Highly Voted 1 year ago

Selected Answer: BD

My answer is BD

It cannot be A, as Junos has an options to automatically assign site identifiers to VPLS sites:

https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/topics/concept/vpns-configuring-vpls-routing-instances.html#id-11510150__id-11568648

B is confirmed as correct on:

https://www.juniper.net/documentation/us/en/software/nce/feature-guide-virtual-private-lan-service/topics/task/vpls-ldp-signaling-solutions.html

D is correct, because as we are not relying on LDP for signalling, we can use RSVP and LDP is signal the LSPs.

upvoted 7 times

mrisler Highly Voted 1 year ago

B,D are correct. B: For LDP signaling within a VPLS routing instance, the Junos OS supports the following values only: Control Bit = 0; D: L2VPNs decouple the service from the transport by using

MPLS label stacking: the outer label takes the tunneled packets to the egress PE, and

the inner label identifies the VPN service.As for the PSN Tunnels, they are simply MPLS LSPs so all of the options discussed in

Chapter 2 apply here, too: LDP, RSVP-TE, BGP-LU, and SPRING.

upvoted 5 times

Guru29 Most Recent 8 months ago

https://www.juniper.net/documentation/us/en/software/nce/feature-guide-virtual-private-lan-service/topics/task/vpls-ldp-signaling-solutions.html

Juniper supports Control bit 0 for LDP based VPLS . So answer should be B,D

upvoted 1 times

zineeddine 9 months, 3 weeks ago

Selected Answer: BD

BD, not CB

upvoted 1 times

antigel8 9 months, 3 weeks ago

Not 'C' for sure, Layer2 VPN official course page 404 mentions that every remote PE needs to be manually configured. If you want auto-discovery, FEC129 or BGP0-signaled VPLS is required.

upvoted 2 times

zineeddine 10 months, 2 weeks ago

Selected Answer: CD

CB are the correct answers

upvoted 1 times

eknow 1 year ago

A,C are correct

upvoted 1 times

Question #6 *Topic 1*

Click the Exhibit button.

```
user@router> show route extensive
...
2:192.168.101.5:65101::22031::02:00:31:06:00:01/304 MAC/IP (2 entries, 1 announced)
TSI:
Page 0 idx 0, (group IBGP-EVPN-Core type Internal) Type 1 val 0xb225964 (adv_entry)
   Advertised metrics:
     Nexthop: 192.168.101.5
     Localpref: 100
     AS path: [65101] I (Originator)
     Cluster list:  2.2.2.2
     Originator ID: 192.168.101.5
     Communities: target:65101:268457487 encapsulation:vxlan(0x8)
     Cluster ID: 3.3.3.3
    Advertise: 00000001
Path 2:192.168.101.5:65101::22031::02:00:31:06:00:01 from 192.168.101.3 Vector len 4.  Val: 0
        *BGP    Preference: 170/-101
                Route Distinguisher: 192.168.101.5:65101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb2d3490
                Next-hop reference count: 10520
                Source: 192.168.101.3
                Protocol next hop: 192.168.101.5
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                State: <Active Int Ext>
                Local AS: 65101 Peer AS: 65101
                Age: 3d 19:56:57        Metric2: 0
                Validation State: unverified
                Task: BGP_65101.192.168.101.3
                Announcement bits (1): 1-BGP_RT_Background
                AS path: I (Originator)
                Cluster list:  2.2.2.2
                Originator ID: 192.168.101.5
                Communities: target:65101:268457487 encapsulation:vxlan(0x8)
                Import Accepted
                Route Label: 22031
                ESI: 05:00:00:fe:4d:00:00:56:0f:00
                Localpref: 100
                Router ID: 192.168.101.3
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 192.168.101.5
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 10.0.2.12 via et-0/0/0.0
                                Session Id: 0x0
                                Next hop: 10.0.2.22 via et-0/0/1.0
                                Session Id: 0x0
192.168.101.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 2
Nexthop: 10.0.2.12 via et-0/0/0.0
Session Id: 0
Nexthop: 10.0.2.22 via et-0/0/1.0
Session Id: 0
...
```

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

A. The devices advertising this route into EVPN are 10.0.2.12 and 10.0.2.22.

B. This route is learned through EBGP.

C. The device advertising this route into EVPN is 192.168.101.5.

D. This is an EVPN Type-2 route.

Suggested Answer: *CD*

Community vote distribution

CD (100%)

Guru29 8 months ago

10.0.2.12 and 10.0.2.22 are destination but 192.168.101.5 is the originator which advertise. Correct answer C,D

upvoted 2 times

snvy 10 months, 3 weeks ago

Selected Answer: CD

CD is correct

upvoted 2 times

Question #7 *Topic 1*

Click the Exhibit button.

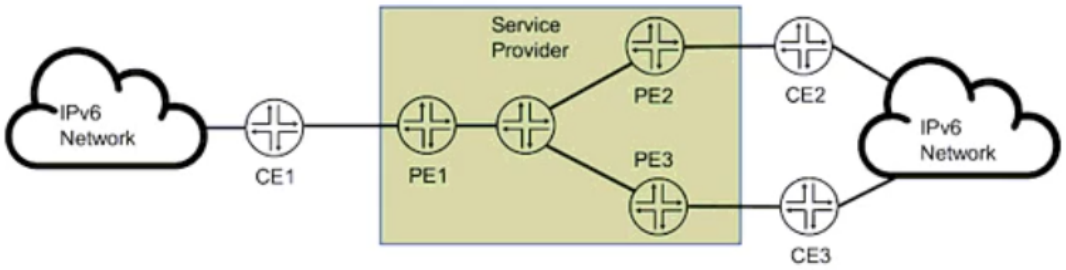

You are running a service provider network and must transport a customer's IPv6 traffic across your IPv4-based MPLS network using BGP. You have already configured mpls ipv6-tunneling on your PE routers.
Which two statements are correct about the BGP configuration in this scenario? (Choose two.)

A. You must configure family inet6 labeled-unicast between PE routers.

B. You must configure family inet6 add-path between PE and CE routers.

C. You must configure family inet6 unicast between PE and CE routers.

D. You must configure family inet6 unicast between PE routers.

Suggested Answer: *AC*

Community vote distribution

AC (82%) AD (18%)

zineeddine Highly Voted 10 months, 2 weeks ago

Selected Answer: AC

To transport IPv6 traffic over an IPv4-based MPLS network using BGP, you need to configure two address families: family inet6 labeled-unicast and family inet6 unicast. The former is used to exchange IPv6 routes with MPLS labels between PE routers, and the latter is used to exchange IPv6 routes without labels between PE and CE routers. The mpis ipv6-tunneling command enables the PE routers to encapsulate the IPv6 packets with an MPLS label stack and an IPv4 header before sending them over the MPLS network.

upvoted 6 times

Penieljacobpaul Most Recent 2 months ago

Selected Answer: AC

A. You must configure family inet6 labeled-unicast between PE routers.

Why It's Correct:

To transport IPv6 traffic over an IPv4 MPLS core, the PEs must exchange IPv6 routes with MPLS labels . This requires enabling family inet6 labeled-unicast in the BGP session between PEs.

Labeled-unicast allows BGP to carry both the IPv6 prefix and the MPLS label required for forwarding across the IPv4 core.

C. You must configure family inet6 unicast between PE and CE routers.

Why It's Correct:

The CE routers (customer edge) need to learn IPv6 routes from the SP's PEs. This requires the PE to advertise IPv6 routes (family inet6 unicast) to the CE.

The CE does not need labels (since it's at the edge), so standard unicast is sufficient.

upvoted 1 times

Harrrrrrrry 6 months, 3 weeks ago

Selected Answer: AD

A&D is correct. The question is about the provider network, need to enable both v6 labeled unicast and v6 unicast between PEs.

upvoted 1 times

nas1785 9 months, 1 week ago

Selected Answer: AC

A& C

upvoted 2 times

snvy 10 months, 3 weeks ago

AC, reason exhibit is given with PE-CE, we cant ignore its PE-PE

upvoted 1 times

M80T 12 months ago

Selected Answer: AD

Note that the question asks about BGP configuration which is PE to PE.

There is no mention of BGP running between PE and CE here, so any answer involving PE to CE is wrong.

upvoted 1 times

Question #8 *Topic 1*

When using OSPFv3 for an IPv4 environment, which statement is correct?

A. OSPFv3 only supports IPv4.

B. OSPFv3 is not backward compatible with IPv4.

C. OSPFv3 supports both IPv6 and IPv4, but not in the same routing instance.

D. OSPFv3 supports IPv4 only on interfaces with family inet6 defined.

Suggested Answer: *D*

Community vote distribution

D (100%)

Penieljacobpaul 2 months ago

Selected Answer: D

D: OSPFv3 supports IPv4 only on interfaces with family inet6 defined.
This is due to the protocol's reliance on IPv6 link-local addresses for neighbor discovery and transport.

upvoted 1 times

Sonu_860 9 months, 3 weeks ago

D is correct

upvoted 1 times

22cc424 9 months, 3 weeks ago

I believe the answer here is D.
You need to enable inet6 to be able to do OSPFv3. As OSPFv3 uses the link-local address which wont be there if you don't have inet6 defined.

upvoted 1 times

141f14b 1 year ago

Selected Answer: D

D is correct

i tested in lab and OSPFv3 realm-ipv4-unicast neighborship is not coming up if the interface does not have family inet6 configured

upvoted 1 times

zineeddine 1 year, 4 months ago

Selected Answer: D

if you want to use OSPFv3 for IPv4 routing, you would typically configure the family inet6 option along with the family inet option to enable OSPFv3 to handle both IPv6 and IPv4 routing.

upvoted 4 times

anonymonkey 1 year, 4 months ago

B
https://community.cisco.com/t5/ipv6/ospfv3-basic-question/td-p/2372187
or
https://www.ciscopress.com/articles/article.asp?p=3188198#:~:text=OSPF%20Version%203%20(OSPFv3)%2C,essentially%20the%20same%20for%20OSPFv3.

upvoted 1 times

BobbyAxelrod 2 months, 3 weeks ago

Why are you quoting Cisco?!!

upvoted 1 times

Jicks 1 year, 6 months ago

Selected Answer: D

OSPF must require IPv6 on interface

upvoted 3 times

Question #9 *Topic 1*

You are a network architect for a service provider and want to offer Layer 2 services to your customers. You want to use EVPN for Layer 2 services in your existing MPLS network.
Which two statements are correct in this scenario? (Choose two.)

A. Segment routing must be configured on all PE routers.

B. EVPN uses Type 2 routes to advertise MAC address and IP address pairs learned using ARP snooping.

C. EVPN uses Type 3 routes to join a multicast tree to flood traffic.

D. VXLAN must be configured on all PE routers.

Suggested Answer: *BC*

Community vote distribution

BC (100%)

zineeddine 10 months, 2 weeks ago

Selected Answer: BC

BC correct

upvoted 4 times

snvy 10 months, 3 weeks ago

BC is correct

upvoted 2 times

Question #10 *Topic 1*

You are configuring a Layer 3 VPN between two sites. You are configuring the vrf-target target: 65100:100 statement in your routing instance. In this scenario, which two statements describe the vrf-target configuration? (Choose two.)

A. This value is used to identify BGP routes learned from the remote PE device.

B. This value is used to add a target community to BGP routes advertised to the local CE device.

C. This value is used to add a target community to BGP routes advertised to the remote PE device.

D. This value is used to identify BGP routes learned from the local CE device.

Suggested Answer: *AC*

Community vote distribution

AC (100%)

zineeddine Highly Voted 10 months, 2 weeks ago

Selected Answer: AC

correct

upvoted 5 times

Penieljacobpaul Most Recent 2 months ago

Selected Answer: AC

Correct Answers: A and C

A: The VRF target filters routes received from remote PEs by matching the RT community.

C: The VRF target is added as a community to routes exported from the VRF to MP-BGP.

upvoted 1 times

Question #11 *Topic 1*

Click the Exhibit button.

```
user@R1> show route protocol bgp
inet.0: 8 destinations, 12 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
172.16.20.4/30     *[BGP/170] 00:49:55, localpref 100
                      AS path: 2 I, validation-state: unverified
                    > to 10.0.18.2 via ge-1/0/4.0
                      to 10.0.19.2 via ge-1/0/5.0
                    [BGP/170] 00:49:55, localpref 100
                      AS path: 2 I, validation-state: unverified
                    > to 10.0.19.2 via ge-1/0/5.0
```

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

A. The multihop configuration is used for load balancing.

B. This route is learned from the same AS number.

C. This route is learned from two different AS numbers.

D. The multipath configuration is used for load balancing.

Suggested Answer: *BD*

Community vote distribution

BD (100%)

Penieljacobpaul 2 months ago

Selected Answer: BD

Route is learned from the same AS & multipath along with load balancing policy is enabled

upvoted 1 times

crramos 9 months, 1 week ago

Selected Answer: BD

Correct B,D

upvoted 1 times

Guru29 1 year, 2 months ago

AS PATH for 2 I , so it's from same as and two interfaces pointing so multipath enabled. Correct B,D

upvoted 2 times

Question #12 *Topic 1*

A packet is received on an interface configured with transmission scheduling. One of the configured queues is exceeding the allocated bandwidth. In this scenario, which two actions will be taken by default on a Junos device? (Choose two.)

A. The exceeding queue will be considered to have negative bandwidth credit.

B. The excess traffic will be discarded.

C. The exceeding queue will be considered to have positive bandwidth credit.

D. The excess traffic will use bandwidth available from other queues.

Suggested Answer: *AD*

Community vote distribution

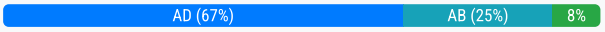

M80T Highly Voted 1 year, 7 months ago

Correct answer is AD
https://www.juniper.net/documentation/us/en/software/junos/cos-security-devices/topics/concept/cos-transmission-scheduling-security-overview.html
upvoted 23 times

Penieljacobpaul Most Recent 2 months ago

Selected Answer: AD
Correct Answers: A and D

A: Exceeding queues use negative credits to control transmission timing.
D: Excess traffic borrows unused bandwidth from other queues by default.
upvoted 1 times

quraitulain 8 months, 1 week ago

By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of offered load than the bandwidth allocated. A queue receiving traffic within its bandwidth configuration is considered to have positive bandwidth credit, and a queue receiving traffic in excess of its bandwidth allocation is considered to have negative bandwidth credit.
upvoted 1 times

Avee123 10 months, 1 week ago

The correct actions that will be taken by default on a Junos device when a queue exceeds its allocated bandwidth are:

A. The exceeding queue will be considered to have negative bandwidth credit.
When a queue exceeds its allocated bandwidth, Junos devices assign negative credit to that queue, which impacts its future ability to send traffic.

B. The excess traffic will be discarded.
Traffic exceeding the configured bandwidth limits is typically discarded to prevent congestion and maintain QoS (Quality of Service) standards.

The other options, C and D, are incorrect because:

C. Positive bandwidth credit is not given for exceeding queues.

D. Excess traffic does not automatically use bandwidth from other queues unless explicitly configured (like with a shared buffer or similar mechanism).
upvoted 2 times

Ant_OnioN 1 year, 2 months ago

Selected Answer: AD
A / D seem to be correct
upvoted 1 times

easynet13 1 year, 3 months ago

Selected Answer: AD

AD is correct.

upvoted 1 times

nas1785 1 year, 3 months ago

Selected Answer: AD

By default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of offered load than the bandwidth allocated. A queue receiving traffic within its bandwidth configuration is considered to have positive bandwidth credit, and a queue receiving traffic in excess of its bandwidth allocation is considered to have negative bandwidth credit.

https://www.juniper.net/documentation/us/en/software/junos/cos-security-devices/topics/concept/cos-transmission-scheduling-security-overview.html

upvoted 3 times

Dimsop_Technology 1 year, 3 months ago

Selected Answer: AD

De forma predeterminada, cada cola puede exceder el ancho de banda asignado si hay ancho de banda adicional disponible en otras colas. Cuando una clase de reenvío no utiliza completamente el ancho de banda de transmisión asignado, el ancho de banda restante puede ser utilizado por otras clases de reenvío si reciben una cantidad mayor de carga ofrecida que el ancho de banda asignado. Se considera que una cola que recibe tráfico dentro de su configuración de ancho de banda tiene un crédito de ancho de banda positivo, y una cola que recibe tráfico que excede su asignación de ancho de banda tiene un crédito de ancho de banda negativo.

upvoted 2 times

zineeddine 1 year, 4 months ago

Selected Answer: BD

In this scenario, the two default actions taken by a Junos device are:

B. The excess traffic will be discarded.
D. The excess traffic will use bandwidth available from other queues.

upvoted 1 times

tsukasa123 1 year, 7 months ago

Selected Answer: AB

If a queue exceeds its allocated bandwidth, it is considered to have negative bandwidth credit and its excess traffic will be discarded by default.

upvoted 3 times

Question #13 *Topic 1*

Which two statements are correct about a sham link? (Choose two.)

A. The PEs exchange Type 1 OSPF LSAs instead of Type 3 OSPF LSAs for the L3VPN routes.

B. It creates an OSPF multihop neighborship between two PE routers.

C. It creates a BGP multihop neighborship between two PE routers.

D. The PEs exchange Type 3 OSPF LSAs instead of Type 1 OSPF LSAs for the L3VPN routes.

Suggested Answer: *AB*

Community vote distribution

AB (100%)

crramos 9 months, 2 weeks ago

Selected Answer: AB

AB

https://www.juniper.net/documentation/us/en/software/junos/ospf/topics/topic-map/configuring-ospfv2-sham-links.html

upvoted 1 times

zineeddine 1 year, 4 months ago

Selected Answer: AB

correct

upvoted 4 times

Click the Exhibit button.

```
user@R4> show pim rps
Instance: PIM.master
address-family INET
RP address       Type         Mode   Holdtime Timeout Groups Group prefixes
10.1.255.2     bootstrap    sparse       150     118      0 224.1.1.0/24
10.1.255.3     bootstrap    sparse       150     118      2 224.1.1.0/28
user@R4> show route 10.1.255.2
inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.2/32      *[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
user@R4> show route 10.1.255.3

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.3/32      *[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# show protocols pim
rp {
    bootstrap {
        family inet {
            priority 200;
        }
    }
    local {
        address 10.1.255.2;
        group-ranges {
            224.1.1.0/24;
        }
    }
}
interface all;
[edit]
user@R3# show protocols pim
rp {
    bootstrap {
        family inet {
            priority 210;
        }
    }
    local {
        address 10.1.255.3;
        group-ranges {
            224.1.1.0/28;
        }
    }
}
interface all;
```

R4 is directly connected to both RPs (R2 and R3). R4 is currently sending all joins upstream to R3 but you want all joins to go to R2 instead. Referring to the exhibit, which configuration change will solve this issue?

A. Change the bootstrap priority on R2 to be higher than R3.

B. Change the local address on R2 to be higher than R3.

C. Change the default route in inet.2 on R4 from R3 as the next hop to R2.

D. Change the group-range to be more specific on R2 than R3.

Suggested Answer: *D*

Community vote distribution

a65ecbc Highly Voted 1 year, 7 months ago

Selected Answer: D

More specific group range

upvoted 15 times

Jentti Most Recent 2 months, 2 weeks ago

Selected Answer: A

I'm just wondering that the question is stating "for ALL groups". Now if you change the group range to be more specific on R2, it will certainly send multicast to R2 for matching groups, but IMHO it won't fulfill the requiremen of ALL groups. Thus, I would match the group subnets to /24 on both and then increase the priority on R2. Your thoughts?

upvoted 2 times

udfiore 4 months ago

Selected Answer: D

Changing the bootstrap priority will only influence the bootstrap election. You need to chage the group range on R2 to be more specific.

upvoted 1 times

JP4CCNP 4 months ago

Selected Answer: D

1. Find all RPs with the most specific group range covering G.
2. From the subset in step 1, select all RPs with the highest priority (lowest priority value).
3. For the RPs that meet the requirements in steps 1 and 2, compute a hash value based on the group address G, the RP address, and the hask mask included in the Bootstrap messages. The RP with the highest hash value is the RP for the group.
4. In the case of a tie (that is, the same group address, priority and hash value), the RP with the highest IP address is the active RP.

upvoted 1 times

bradpope 6 months ago

Selected Answer: A

A - Priority

See the statement copied/pasted below from the information at this URL:
https://www.juniper.net/documentation/us/en/software/junos/multicast/topics/topic-map/mcast-pim-bootstrap-router.html

"By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device."

upvoted 2 times

Tony87 11 months, 3 weeks ago

D is correct, check this out
https://higherlogicdownload.s3.amazonaws.com/JUNIPER/MigratedInlineFiles/9531b1e0a25c46edaaa976789ad6305e_697bd1d7acea48deb9bd00536487d

upvoted 2 times

Guru29 1 year, 2 months ago

When two RPs advertise the exact same groups then we will prefer the RP with the highest priority. Unlike the BSR selection, highest priority for RP selection means the lowest priority value (0 is the highest priority). The priority is a value we can configure on the RP.
D is Correct

upvoted 1 times

solen003 1 year, 4 months ago

Absolutely priority.
Group can exclude only.

My R3 in lab is R3 in question. My R4 in lab is R2 in question
root@R4# set rp bootstrap family inet priority 220

Output on other router directly connected to both
address-family INET
RP address Type Mode Holdtime Timeout Groups Group prefixes
192.168.0.3 bootstrap sparse 150 131 0 224.0.0.0/4
192.168.0.4 bootstrap sparse 150 131 2 224.0.0.0/4 <<<<<<<<<<< 2 GROUPS

Group: 239.2.2.2
Source: *
RP: 192.168.0.4
Flags: sparse,rptree,wildcard
Upstream interface: ge-0/0/4.0 <<<<< To my R4 (R2 in question)

Group: 239.1.1.1
Source: *
RP: 192.168.0.4
Flags: sparse,rptree,wildcard
Upstream interface: ge-0/0/4.0 <<<<<< To my R4 (R2 in question)
upvoted 4 times

zineeddine 1 year, 4 months ago

Selected Answer: A

correct
upvoted 1 times

Question #15 *Topic 1*

Which two statements are correct regarding the PIM DR in a PIM-SM domain? (Choose two.)

A. By default. PIM DR election is performed on point-to-point links.

B. If the DR priorities match, the router with the lowest IP address is selected as the DR.

C. The source DR sends PIM register messages from the source network to the RP.

D. The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.

Suggested Answer: *CD*

Community vote distribution

CD (100%)

piipo 11 months, 1 week ago

Selected Answer: CD

Correct

upvoted 1 times

Guru29 1 year, 2 months ago

DR is elected by Highest Priority if Priority ties then Highest IP address .

upvoted 2 times

Question #16 *Topic 1*

An interface is configured with a behavior aggregate classifier and a multifield classifier.
How will the packet be processed when received on this interface?

A. The packet will be processed by the MF classifier first, then the BA classifier.

B. The packet will be processed by the BA classifier first, then the MF classifier.

C. The packet will be discarded.

D. The packet will be forwarded with no classification changes.

Suggested Answer: *B*

Community vote distribution

B (100%)

696501d 7 months, 3 weeks ago

Selected Answer: B

correct

upvoted 1 times

hbstyleboy 8 months, 1 week ago

Answer is B: You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first, and then the MF classification is performed. If the two classification results conflict, the MF classification result overrides the BA classification result.

upvoted 4 times

zineeddine 10 months, 2 weeks ago

Selected Answer: B

correct

upvoted 2 times

Question #17 *Topic 1*

You are configuring a BGP signaled Layer 2 VPN across your MPLS enabled core network. Your PE-2 device connects to two sites within the same VPN.
In this scenario, which statement is correct?

A. You must create a unique Layer 2 VPN routing instance for each site on the PE-2 device.

B. By default on PE-2, the site's local ID is automatically assigned a value of 0 and must be configured to match the total number of attached sites.

C. You must use separate physical interfaces to connect PE-2 to each site.

D. By default on PE-2, the remote site IDs are automatically assigned based on the order that you add the interfaces to the site configuration.

Suggested Answer: *D*

Community vote distribution

D (100%)

zineeddine Highly Voted 10 months, 2 weeks ago

Selected Answer: D

BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path.
In BGP Layer 2 VPNs, each site has a unique site ID that identifies it within a VFI. The site ID can be manually configured or automatically assigned by the PE device. By default, the site ID is automatically assigned based on the order that you add the interfaces to the site configuration. The first interface added to a site configuration has a site ID of 1, the second interface added has a site ID of 2, and so on.
upvoted 6 times

Penieljacobpaul Most Recent 2 months ago

Selected Answer: D

In BGP-signaled L2VPN configurations (e.g., Juniper Junos), when configuring multiple local sites under the same VPN instance on a PE router:

Local site IDs are automatically assigned starting from 1 (not 0) if not explicitly configured. The first site added gets ID 1, the second 2, and so on.

The statement in D inaccurately refers to "remote site IDs," but this behavior actually applies to local site IDs. However, given the options provided, D is the closest match to the correct behavior, assuming a possible wording error in the question (e.g., "remote" instead of "local").
upvoted 1 times

daff177 10 months, 3 weeks ago

D i assume
upvoted 1 times

Question #18 *Topic 1*

You are responding to an RFP for a new MPLS VPN implementation. The solution must use LDP for signaling and support Layer 2 connectivity without using BGP. The solution must be scalable and support multiple VPN connections over a single MPLS LSP. The customer wants to maintain all routing for their private network.
In this scenario, which solution do you propose?

A. circuit cross-connect

B. BGP Layer 2 VPN

C. LDP Layer 2 circuit

D. translational cross-connect

Suggested Answer: *C*

Community vote distribution

C (100%)

ThoNguyen 10 months, 4 weeks ago

C is correct

upvoted 1 times

marius2 1 year, 4 months ago

Selected Answer: C

C is correct.

upvoted 3 times

marius2 1 year, 4 months ago

C is correct.

upvoted 1 times

Question #19 *Topic 1*

Click the Exhibit button.

```
192.168.1.1:455:10.1.1.0/24
```

You are examining an L3VPN route that includes the information shown in the exhibit.
Which statement is correct in this scenario?

A. The information shows a route target.

B. The information shows a Type 1 route distinguisher.

C. The information shows a Type 0 route distinguisher.

D. The information shows a Type 2 route distinguisher.

Suggested Answer: *B*

Community vote distribution

B (100%)

M80T Highly Voted 1 year ago

Selected Answer: B

As per https://datatracker.ietf.org/doc/html/rfc4364, the correct answer is B (Type 1 RD)

upvoted 5 times

Penieljacobpaul Most Recent 2 months ago

Selected Answer: B

Type 0: The RD is composed of a 2-byte Autonomous System (AS) number (Administrator subfield) and a 4-byte assigned number (Assigned Number subfield). It is represented as AS_Number:Assigned_Number.
Type 1: The RD is composed of a 4-byte IP address (Administrator subfield) and a 2-byte assigned number (Assigned Number subfield). It is represented as IP_Address:Assigned_Number.
Type 2: The RD is composed of a 4-byte AS number (Administrator subfield, often from the AS4 range) and a 2-byte assigned number (Assigned Number subfield). It is represented as AS4_Number:Assigned_Number.
Looking at the RD 192.168.1.1:455:

192.168.1.1 is a 4-byte IP address.
455 is a number that fits within a 2-byte unsigned integer range (0-65535).

upvoted 1 times

marius2 10 months, 1 week ago

Selected Answer: B

Type1 is correct

upvoted 1 times

zineeddine 10 months, 1 week ago

Selected Answer: B

type 1

upvoted 2 times

StargateCommand 11 months, 3 weeks ago

Selected Answer: B

type 1, of course.

upvoted 3 times

a65ecbc 1 year, 1 month ago

Selected Answer: B

Type 1: When Type value is 1, the Administrator field is 4-bytes and Assigned Number field is 2-bytes.

The Administrator field should be set to the IP address (public IP addresses should be used). The Assigned Number field contains a number from a numbering space that is administered by the enterprise to which the IP address has been assigned by the appropriate authority.

upvoted 4 times

Arthur_Tumwesigye 1 year, 1 month ago

Correct answer is 1

upvoted 2 times

Arthur_Tumwesigye 1 year, 1 month ago

Answer is B - Type 1

upvoted 1 times

Question #20 *Topic 1*

Click the Exhibit button.

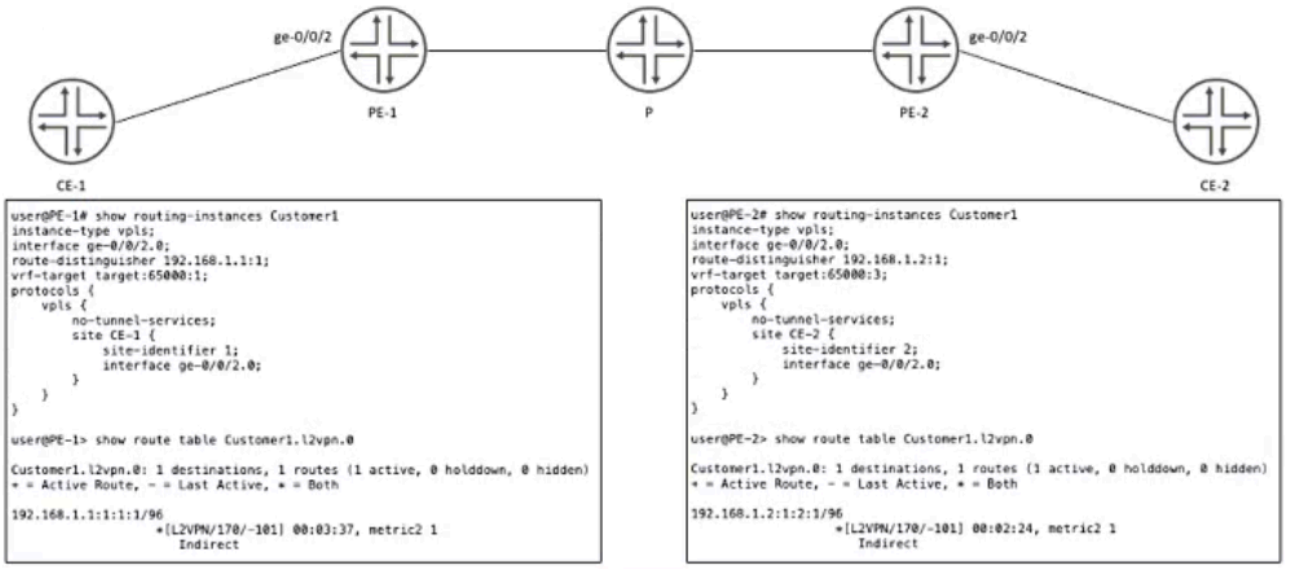

CE-1 and CE-2 are part of a VPLS called Customer1. No connectivity exists between CE-1 and CE-2. In the process of troubleshooting, you notice PE-1 is not learning any routes for this VPLS from PE-2, and PE-2 is not learning any routes for this VPLS from PE-1.
Referring to the exhibit, which statement is correct?

A. The route target must match on PE-1 and PE-2.

B. The instance type should be changed to l2vpn.

C. The no-tunnel-services statement should be deleted on both PEs.

D. The route distinguisher must match on PE-1 and PE-2.

Suggested Answer: *A*

Community vote distribution

A (88%) 13%

marius2 Highly Voted 1 year, 4 months ago

Selected Answer: A

A is correct.

upvoted 5 times

Penieljacobpaul Most Recent 2 months ago

Selected Answer: A

The lack of route exchange between PE-1 and PE-2 is due to mismatched route targets in their VPLS configurations:

PE-1 uses vrf-target target:65088::1.

PE-2 uses vrf-target target:65088::3.

In VPLS, route targets (RTs) determine how routes are imported/exported between PEs. For auto-discovery and route exchange in a VPLS domain:

All participating PEs must use the same route target (or compatible import/export policies).

If RTs do not match, PEs will not recognize each other as part of the same VPLS, leading to failed connectivity.

upvoted 2 times

piipo 11 months, 1 week ago

Selected Answer: D

route distinguisher

upvoted 1 times

piipo 11 months, 1 week ago

Sorry, A is Correct.

upvoted 1 times

Question #21 *Topic 1*

Click the Exhibit button.

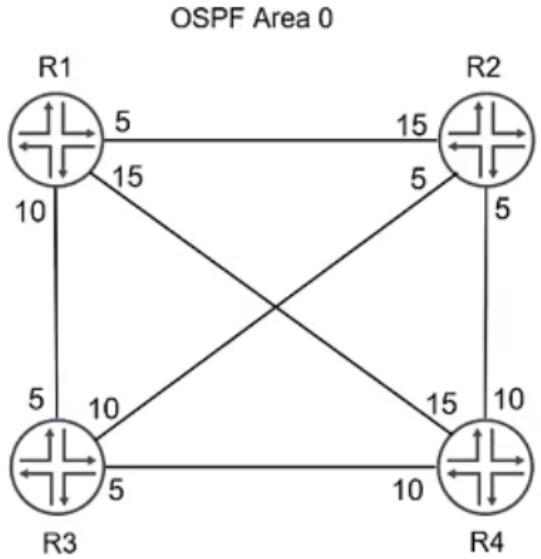

Referring to the exhibit, which path would traffic passing through R1 take to get to R4?

A. R1 -> R3 -> R4

B. R1 -> R2 -> R3 -> R4

C. R1 -> R2 -> R4

D. R1 -> R4

Suggested Answer: *C*

Community vote distribution

C (100%)

M80T Highly Voted 1 year, 6 months ago

Selected Answer: C

The OSPF cost is carried in the LSAs that are exchanged within an OSPF area. When a router calculates the cost to a destination it uses the cost of the exit interface of each router in the path to the destination.

upvoted 5 times

Penieljacobpaul Most Recent 2 months ago

Selected Answer: C

cost of the exit interface is consider while calculating best paths for ospf

upvoted 1 times

zineeddine 1 year, 4 months ago

Selected Answer: C

C is the shortest path

upvoted 2 times

augustgg21 1 year, 6 months ago

only see the outgoing interfaces

upvoted 1 times

Koee 1 year, 6 months ago

Anyone can explain this for me?

upvoted 1 times

MichelT 9 months, 2 weeks ago

Add up all outgoing interface costs. Lowest total cost is the big winner.

upvoted 1 times

Question #22 *Topic 1*

Click the Exhibit button.

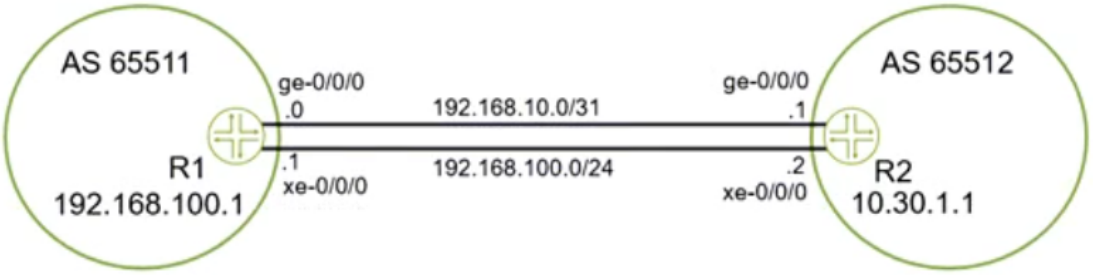

You want to use both links between R1 and R2. Because of the bandwidth difference between the two links, you must ensure that the links are used as much as possible.
Which action will accomplish this goal?

A. Ensure that the metric-out parameter on the Gigabit Ethernet interface is higher than the 10 Gigibit Ethernet interface.

B. Define a policy to tag routes with the appropriate bandwidth community.

C. Enable per-prefix load balancing.

D. Disable multipath.

Suggested Answer: *B*

Community vote distribution

B (100%)

Penieljacobpaul 2 months ago

Selected Answer: B

Answer: B. Define a policy to tag routes with the appropriate bandwidth community.

Explanation:
To maximize the use of both links (Gigabit Ethernet and 10 Gigabit Ethernet) proportionally to their bandwidths:

BGP bandwidth communities are critical. They encode the bandwidth of each next hop, enabling Junos to distribute traffic asymmetrically based on link capacity.

The text explicitly states that combining the bandwidth community with multipath and per-packet (per-flow) load balancing allows traffic to be split proportionally (e.g., 10:1 for a 10Gbps and 1Gbps link).

While enabling per-packet load balancing is required for this functionality, the question asks for the action to accomplish the goal. The bandwidth community must be configured via an import policy to reflect link speeds, which is the foundational step.

upvoted 1 times

easynet13 9 months, 2 weeks ago

Selected Answer: B

B is correct.

upvoted 1 times

StargateCommand 11 months, 3 weeks ago

Selected Answer: B

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

upvoted 4 times

M80T 1 year ago

Selected Answer: B

The answer is B as per https://www.juniper.net/documentation/us/en/software/junos/sampling-forwarding-monitoring/bgp/topics/concept/bgp-multipath-unequal-understanding.html

upvoted 4 times

Bob_Smith 1 year ago

B makes the most sense. Juniper can advertise a bandwidth community that will load balance flows proportionally based on bandwidth values

upvoted 2 times

eknow 1 year ago

Selected Answer: B

C is not correct. you should use per-packet instead of per-prefix.

upvoted 2 times

Question #23 *Topic 1*

Which two statements are correct about the customer interface in an LDP-signaled pseudowire? (Choose two.)

A. When the encapsulation is vlan-ccc or extended-vlan-ccc, the configured VLAN tag is included in the control plane LDP advertisement.

B. When the encapsulation is ethernet-ccc, tagged and untagged frames are both accepted in the data plane.

C. When the encapsulation is ethernet-ccc, only frames without a VLAN tag are accepted in the data plane.

D. When the encapsulation is vlan-ccc or extended-vlan-ccc, the configured VLAN tag is not included in the control plane LDP advertisement.

Suggested Answer: *AB*

Community vote distribution

AB (100%)

solen003 10 months, 1 week ago

[edit interfaces]
'ge-0/0/0'
ethernet-ccc encapsulation not allowed on tagged interfaces
error: configuration check-out failed

upvoted 1 times

zineeddine 10 months, 1 week ago

Selected Answer: AB

AB correct

upvoted 1 times

StargateCommand 11 months, 3 weeks ago

Selected Answer: AB

C is wrong!

A and B are correct

upvoted 3 times

eknow 1 year ago

Selected Answer: AB

A,B are correct.

upvoted 3 times

Question #24 *Topic 1*

Click the Exhibit button.

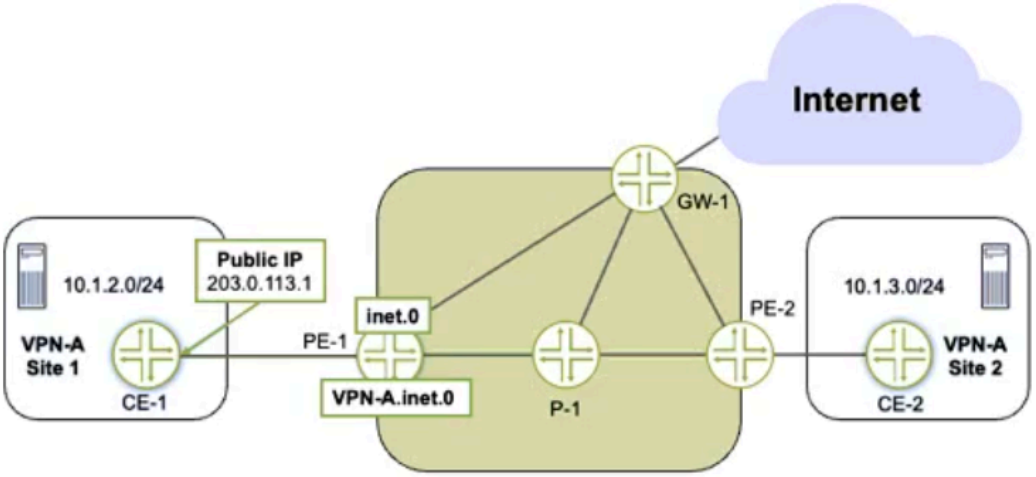

Referring to the exhibit, CE-1 is providing NAT services for the hosts at Site 1 and you must provide Internet access for those hosts.
Which two statements are correct in this scenario? (Choose two.)

A. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.

B. You must configure a static route in the main routing instance for the 10.1.2.0/24 prefix that uses the VPN-A.inet.0 table as the next hop.

C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-A.inet.0 table.

D. You must configure a RIB group on PE-1 to leak the 10.1.2.0/24 prefix from the VPN-A.inet.0 table to the inet.0 table.

Suggested Answer: *AC*

Community vote distribution

AC (100%)

a65ecbc Highly Voted 1 year, 7 months ago

Selected Answer: AC

It says the CE is doing the NAT.

upvoted 7 times

Bob_Smith Highly Voted 1 year, 5 months ago

Selected Answer: AC

We want a static route to the public IP. The private IP is hidden behind the NATed public IP, so a static route to the private range is useless to us.

upvoted 5 times

Avee123 Most Recent 10 months, 1 week ago

AC

A. You must configure a static route in the main routing instance for the 203.0.113.1/32 prefix that uses the VPN-A.inet.0 table as the next hop.

This static route is necessary because the traffic destined for the public IP (203.0.113.1) needs to be directed to the appropriate VPN routing instance (VPN-A) for proper forwarding.

C. You must configure a RIB group on PE-1 to leak a default route from the inet.0 table to the VPN-A.inet.0 table.

The RIB group is required to leak the default route into the VPN routing instance, so that traffic from Site 1 can reach the Internet through the default route.

These configurations help to ensure that the hosts behind CE-1 can access the Internet through proper routing.

upvoted 2 times

Tony87 11 months, 3 weeks ago

Selected Answer: AC

upvoted 1 times

zineeddine 1 year, 4 months ago

Selected Answer: AC

because of NAT

upvoted 2 times

StargateCommand 1 year, 5 months ago

Selected Answer: AC

Selected Answer: AC

upvoted 3 times

Question #25 *Topic 1*

Click the Exhibit button.

```
user@R1> show isis database R3 level 2 extensive | find TLVs | match prefix
    IP prefix: 10.100.34.0/24, Internal, Metric: default 63, Up
    IP prefix: 10.100.13.0/24, Internal, Metric: default 63, Up
    IP extended prefix: 10.100.34.0/24 metric 63 up
    IP extended prefix: 10.100.13.0/24 metric 63 up
user@R1> show isis database R2 level 2 extensive | find TLVs | match prefix
    IP extended prefix: 10.100.12.0/24 metric 1000 up
    IP extended prefix: 10.100.23.0/24 metric 10 up
    IP extended prefix: 10.100.24.0/24 metric 1000 up
```

You enabled a new router (R3) in your network but all destinations using IS-IS routes are not properly load balancing over this new router. Referring to the exhibit, what is the problem?

A. R2 is missing internal routes for R1.

B. R1 is missing internal routes for R2.

C. R3 does not have wide-metrics enabled.

D. R1 does not have wide-metrics enabled.

Suggested Answer: *C*

Community vote distribution

C (100%)

Penieljacobpaul 2 months ago

Selected Answer: C

The IS-IS database outputs reveal a mismatch in metric types:

R3 advertises prefixes with narrow metrics (metric 63, the maximum value for 6-bit narrow metrics).

R2 uses wide metrics (metric 1000, which requires 24-bit wide metrics).

Why This Causes Load Balancing Failure:

IS-IS cannot compare or load balance between routes with narrow and wide metrics.

R1 likely has wide metrics enabled (as it processes R2's 1000 metric), but R3's narrow metrics (63) are incompatible. This prevents R1 from considering R3's paths as valid alternatives for load balancing.

upvoted 1 times

piipo 11 months, 1 week ago

Selected Answer: C

C is Correct

upvoted 3 times

Question #26 *Topic 1*

You are configuring schedulers to define the class-of-service properties of output queues. You want to control packet drops during periods of congestion.
In this scenario, which CoS configuration parameter would be used to accomplish this task?

A. buffer size

B. priority

C. shaping rate

D. drop profile

Suggested Answer: *D*

Community vote distribution

D (100%)

Penieljacobpaul 2 months ago

Selected Answer: D

A drop profile defines how packets are discarded during congestion based on queue fill levels. It uses thresholds and algorithms (e.g., Random Early Detection - RED) to proactively drop packets before a queue overflows, preventing uncontrolled tail drops. This directly addresses the requirement to "control packet drops during periods of congestion."

upvoted 1 times

696501d 7 months, 3 weeks ago

Selected Answer: D

A drop profile is a feature of the random early detection (RED) process that allows packets to be dropped before queues are full.

upvoted 3 times

Guru29 8 months ago

https://www.juniper.net/documentation/us/en/software/junos/cos-security-devices/topics/concept/cos-red-drop-profile-security-overview.html#:~:text=A%20drop%20profile%20is%20a,fullness%20and%20the%20drop%20probability.

Correct D

upvoted 1 times

Question #27 *Topic 1*

Click the Exhibit button.

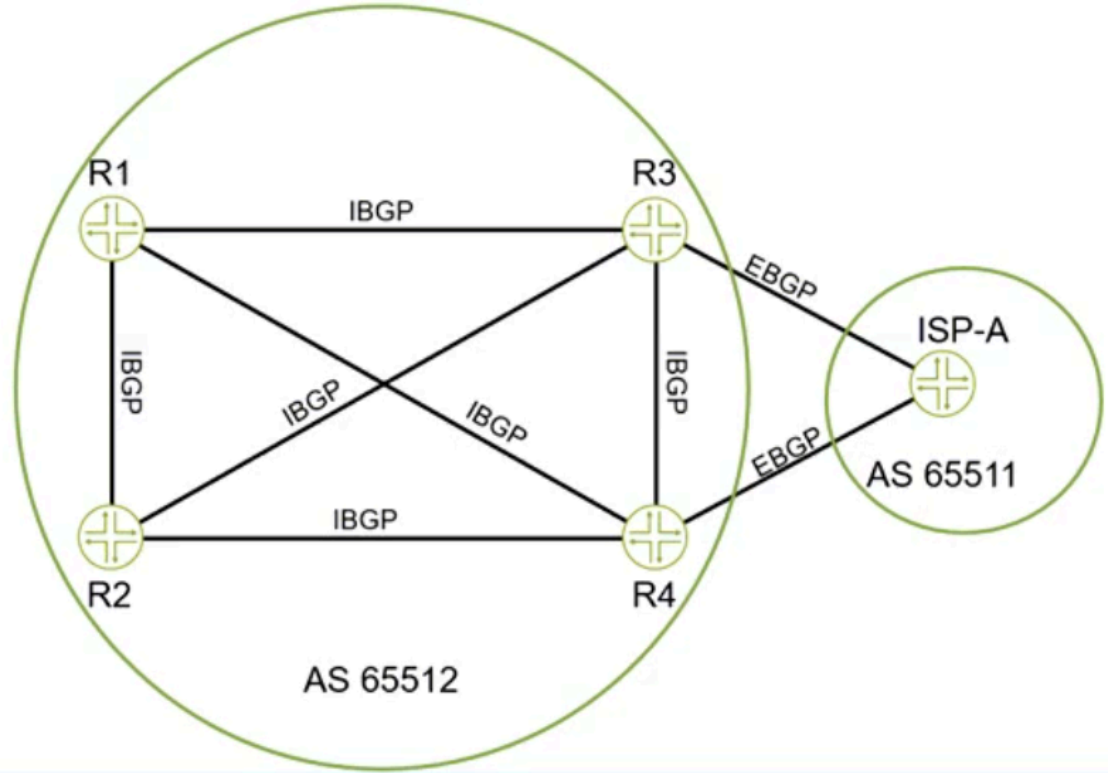

Referring to the exhibit, which two statements are correct about BGP routes on R3 that are advertised to R1? (Choose two.)

A. By default, the next-hop value for these routes is not changed by R3 before being sent to R1.

B. By default, the next-hop value for these routes is changed by R3 before being sent to R1.

C. By default, the BGP local-preference value that is assigned on R3 is advertised to R1.

D. By default, all BGP attributes values must be removed before advertising the routes to R1.

Suggested Answer: *AC*

Community vote distribution

AC (100%)

piipo 10 months, 3 weeks ago

Selected Answer: AC

Correct

upvoted 3 times

Guru29 1 year, 2 months ago

Correct A,C

upvoted 1 times

Question #28 *Topic 1*

You want to ensure that a single-area OSPF network will be loop free.
In this scenario, what are two requirements that satisfy this requirement? (Choose two.)

A. The DR/BDR ensures that each node within an area has the same information in their LSDBs.

B. The Shortest Path First algorithm must prune looped paths.

C. Nodes within an area must connect in a full mesh.

D. All nodes within an area must have the same information in their LSDBs.

Suggested Answer: *BD*

Community vote distribution

BD (100%)

Penieljacobpaul 2 months ago

Selected Answer: BD

B. OSPF uses the Shortest Path First (SPF) algorithm (Dijkstra's algorithm) to compute loop-free paths by pruning any redundant or looping routes. This ensures the topology is a tree with no cycles.

D. For OSPF to be loop-free, all routers in the area must have identical Link-State Databases (LSDBs). Consistent LSDBs guarantee that every router computes the same shortest paths, eliminating routing loops.

upvoted 1 times

4e29b77 1 year ago

Selected Answer: BD

BR/DR only deals with shared segments to reduce LSA flooding and is not about controlling loops.

upvoted 2 times

zineeddine 1 year, 4 months ago

Selected Answer: BD

The statement about the DR/BDR (Designated Router/Backup Designated Router) ensuring that each node within an area has the same information in their LSDBs is incorrect. While DR/BDR election helps reduce OSPF control traffic by limiting LSDB synchronization to only the DR and BDR, it does not ensure that each node within an area has the same LSDB information. OSPF routers exchange LSAs (Link State Advertisements) to build and maintain their LSDBs, and all routers within the same area should eventually have identical LSDBs regardless of the DR/BDR election.

upvoted 4 times

StargateCommand 1 year, 5 months ago

Selected Answer: BD

A is wrong!

upvoted 3 times

Jicks 1 year, 6 months ago

Selected Answer: BD

DR/BDR doesnt ensure information sharing within the domain

upvoted 4 times

Question #29 *Topic 1*

Click the Exhibit button.

You are asked to exchange routes between R1 and R4 as shown in the exhibit. These two routers use the same AS number.
Which two steps will accomplish this task? (Choose two.)

A. Configure the BGP group with the advertise-peer-as parameter on R2 and R3.

B. Configure the BGP group with the as-override parameter on R1 and R4.

C. Configure the BGP group with the advertise-peer-as parameter on R1 and R4.

D. Configure the BGP group with the as-override parameter on R2 and R3.

Suggested Answer: *AD*

Community vote distribution

AD (83%) CD (17%)

Penieljacobpaul 2 months ago

Selected Answer: AD

D. as-override on R2 and R3:

R1 and R4 are in the same AS (64522), and the intermediate AS (65411) includes R2 and R3.

When R2 receives a route from R1 (AS 64522), the AS_PATH becomes 64522. R2 adds its own AS (65411) when advertising to R3. R3 then sends the route to R4 (AS 64522), resulting in an AS_PATH of 65411 64522. R4 rejects the route because it sees its own AS in the path.

Configuring as-override on R2 and R3 replaces the originating AS (64522) with the local AS (65411) in the AS_PATH. This ensures R1 and R4 do not see their own AS in the path, allowing routes to be accepted.

A. advertise-peer-as on R2 and R3:

The advertise-peer-as parameter forces R2 and R3 to advertise the peer's AS (64522) instead of their own (65411). This avoids inserting AS 65411 into the AS_PATH entirely.

When R2 sends routes to R3, it advertises AS 64522 (R1's AS). R3 then advertises these routes to R4 with AS 64522 in the path, avoiding AS_PATH loop detection.

upvoted 1 times

Avee123 10 months, 1 week ago

BD
In the scenario where routers R1 and R4 are using the same AS number and you need to exchange routes between them, the correct answers are:

B. Configure the BGP group with the as-override parameter on R1 and R4.

The as-override option allows BGP to bypass the typical loop prevention mechanism that prevents a router from accepting routes that include its own AS number. Since R1 and R4 are in the same AS, this option is necessary to exchange routes successfully.
D. Configure the BGP group with the as-override parameter on R2 and R3.

Similarly, configuring as-override on R2 and R3 ensures that routes from R1 and R4 (which share the same AS number) can be exchanged, bypassing the default behavior that would prevent them from accepting routes originating from their own AS.
These configurations allow routes to be exchanged even when routers belong to the same AS.

upvoted 1 times

696501d 1 year, 1 month ago

as-override - if the as number of a router recving advertisement is included in as path, router drops this prefix. this can be changed by using as-override.

advertise-peer-as - junos do not advertise route learned from ebgp peer back to same ebgp peer. this can be changed by advertise-peer-as. Hence AD.

upvoted 1 times

Ant_OnioN 1 year, 2 months ago

Selected Answer: AD

Both options on R2 and R3 are meant to advertise the routes that have the same AS origin as the peer, only that with as-override the AS is replaced by the one on R2/R3 before advertising

upvoted 3 times

easynet13 1 year, 3 months ago

Selected Answer: AD

AD is correct.

upvoted 2 times

StargateCommand 1 year, 5 months ago

Selected Answer: AD

A and D

upvoted 2 times

M80T 1 year, 6 months ago

Selected Answer: AD

advertise-peer-as is explained in:
https://www.juniper.net/documentation/us/en/software/junos/routing-policy/bgp/topics/example/bgp-advertise-peer-as.html

Both advertise-peer-as and as-override are BGP settings applied on the PE, not the CE, therefore the correct answer is A and D

upvoted 4 times

eknow 1 year, 6 months ago

Selected Answer: AD

A,D are correct.

upvoted 3 times

tsukasa123 1 year, 7 months ago

Selected Answer: CD

The advertise-peer-as parameter allows a router to advertise its peer's AS number as part of the AS path attribute when sending BGP updates to other peers. This parameter is useful when two routers in the same AS need to exchange routes through another AS, such as in the case of R1 and R4. By configuring this parameter on R1 and R4, they can advertise each other's AS number to R2 and R3, respectively.
The as-override parameter allows a router to replace the AS number of its peer with its own AS number when receiving BGP updates from that peer. This parameter is useful when two routers in different AS need to exchange routes through another AS that has the same AS number as one of them, such as in the case of R2 and R3. By configuring this parameter on R2 and R3, they can override the AS number of R1 and R4 with their own AS number when sending BGP updates to each other.

upvoted 3 times

Question #30 *Topic 1*

A router running IS-IS is configured with an ISO address of 49.0001.00a0.c96b.c490.00.
Which part of this address is the system ID?

A. c96b.c490 is the system identifier.

B. c490 is the system identifier.

C. 00a0.c96b.c490 is the system identifier.

D. 0001.00a0.c96b.c490 is the system identifier.

Suggested Answer: *C*

Community vote distribution

C (100%)

piipo 11 months, 1 week ago

Selected Answer: C

C is Correct

upvoted 1 times

4e29b77 1 year ago

NET (Network Entity Title) consists of:
49—AFI
0001—Area ID
00a0.c96b.c490—System identifier
00—Selector

upvoted 3 times

Question #31 *Topic 1*

Click the Exhibit button.

```
user@R1>show pim join extensive 232.1.1.1
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
Group: 232.1.1.1
    Source: *
    RP: 10.1.255.112
    Flags: sparse,rptree,wildcard
    Upstream interface: ge-0/0/0.0
    Upstream neighbor: 10.1.11.1
    Upstream state: Join to RP
    Uptime: 00:04:10
    Downstream neighbors:
        Interface: Local
        Interface: ge-0/0/2.0
            10.1.1.1 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 00:04:10 Time since last Join: 00:04:10
    Number of downstream interfaces: 2
    Number of downstream neighbors: 1
Group: 232.1.1.1
    Source: 172.16.1.2
    Flags: sparse,spt
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.1.21.1
    Upstream state: Join to Source, Prune to RP
    Keepalive timeout: 317
    Uptime: 00:01:39
    Downstream neighbors:
        Interface: Local
        Downstream neighbors:
            Interface: Local
            Interface: ge-0/0/2.0
                10.1.1.1 State: Join Flags: S   Timeout: Infinity
                Uptime: 00:01:39 Time since last Join: 00:01:39
        Number of downstream interfaces: 2
        Number of downstream neighbors: 1
```

Referring to the exhibit, which two statements are correct regarding the output shown in the exhibit? (Choose two.)

A. The multicast group is an ASM group.

B. The multicast group is an SSM group.

C. The multicast traffic is using the SPT.

D. The multicast traffic is using the RPT.

Suggested Answer: *AC*

Community vote distribution

AC (80%) BC (20%)

solen003 Highly Voted 1 year, 4 months ago

It is a trap.

The first entry is the key, STAR, G entry.

It is ASM despite using 232/8.

A, C

upvoted 8 times

Kondogbia3 Most Recent 2 weeks, 4 days ago

Selected Answer: AC

Any-source multicast (ASM)

• Supports one-to-many and many-to-many applications

Source-specific multicast (SSM)
• Supports only one-to-many applications
In this case we have both, means that A is correct.

If you don't know the source you start with RPT creation starting from the listener to the RP. When you know the source you can change RPT with SPT.
Uptime is an indicator.
So C is correct
upvoted 1 times

Jentti 2 months, 2 weeks ago

Selected Answer: AC

You can configure Junos OS to accept any-source multicast (ASM) join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

So it's really a trap, using SSM reserved range for ASM traffic.
upvoted 1 times

udfiore 4 months ago

Selected Answer: AC

The (*,G) entry indicates that is ASM. Also the Join sent to Source and Prune sent to RP means that traffic is flowing through the SPT
upvoted 1 times

crramos 9 months ago

Selected Answer: AC

A and C correct
upvoted 1 times

Tony87 11 months, 3 weeks ago

A and C are correct, if it's SSM then output should not have "Prune to RP".
upvoted 3 times

easynet13 1 year, 3 months ago

Selected Answer: AC

There is a (*,G) entry. That's only for ASM groups.
upvoted 4 times

StargateCommand 1 year, 5 months ago

Selected Answer: BC

B and C!
upvoted 2 times

StargateCommand 1 year, 5 months ago

Should be "C" and "B"
The group address is from SSM address block range (232.0.0.0 through 232.255.255.255)
upvoted 4 times

Question #32 *Topic 1*

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

A. CSNPs are flooded periodically.

B. PSNPs are flooded periodically.

C. PSNPs contain only descriptions of LSPs.

D. CSNPs contain only descriptions of LSPs.

Suggested Answer: *AD*

Community vote distribution

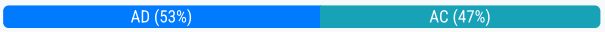

Jentti 2 months, 2 weeks ago

Selected Answer: AC

f the routing device is the designated router on a LAN, IS-IS sends CSNPs every 10 seconds. If the routing device is on a point-to-point interface, it sends CSNPs every 5 seconds. The general recommendation is to use the default values or to increase the CSNP interval if there are a large number of broadcast circuits that need to be supplied with fresh CSNPs. Increasing the interval can help protect against CSNP flooding.

upvoted 1 times

5f53711 2 months, 2 weeks ago

Selected Answer: AC

A:

CSNP=COMPLETE Sequence Number PDU

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

C:

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

upvoted 1 times

bradpope 6 months ago

Selected Answer: AD

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

"When a device receives a CSNP, it checks the database entries against its own local link-state database. If it detects missing information, the device requests specific link-state PDU details using a partial sequence number PDU (PSNP)."

CSNP has to contain only description of LSP, not complete LSP, or there would be no need to request missing/outdated LSP via PSNP.

upvoted 2 times

crramos 9 months, 2 weeks ago

Selected Answer: AD

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

upvoted 2 times

ak_sisko 11 months ago

Not C because PSNPs contain not only descriptions of LSPs.

upvoted 1 times

piipo 11 months, 1 week ago

Selected Answer: AD

PSNPs are used to request or acknowledge specific link-state information. When a router detects missing information via a CSNP, it uses PSNPs to request the necessary LSPs.

upvoted 2 times

4e29b77 1 year ago

Selected Answer: AD

It's actually AD because the CSNP only summarises LSPs in an index format, and routers may request a copy of the LSP using a PSNP if they are missing it.

upvoted 3 times

antigel8 1 year, 3 months ago

Selected Answer: AC

CSNPs contain a complete description of all LSPs in the IS-IS database. IS-IS sends CSNPs periodically on all links

upvoted 1 times

easynet13 1 year, 3 months ago

Selected Answer: AC

AC is correct. Contained within the CSNP is a link-state PDU identifier, a lifetime, a sequence number, and a checksum for each entry in the database. Periodically, a CSNP is sent on both broadcast and point-to-point links to maintain a correct database.

upvoted 2 times

solen003 1 year, 4 months ago

In my opinion this is A nd C.

CSNPs are flooded every 10 seconds on LAN segmnet.

CSNP contains descriptions of LSP.

If neighbor router notices that something is missing, it is requesting via PSNP missing prefixes.

upvoted 1 times

solen003 1 year, 4 months ago

A & D**********

upvoted 3 times

StargateCommand 1 year, 5 months ago

Selected Answer: AC

A and C

upvoted 1 times

tsukasa123 1 year, 7 months ago

Selected Answer: AC

LSPs contain information about the state and cost of links in the network, and are flooded periodically throughout the network. PSNPs are used to acknowledge receipt of LSPs and request retransmission of missing or corrupted LSPs. PSNPs contain only descriptions of LSPs, such as their sequence numbers and checksums.

CSNPs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases.

upvoted 2 times

Question #33 *Topic 1*

Which statement is correct about IS-IS when it performs the Dijkstra algorithm?

A. The local router moves its own local tuples into the candidate database.

B. When a new neighbor ID in the tree database matches a router ID in the LSDthe neighbor ID is moved to the candidate database.

C. Tuples with the lowest cost are moved from the tree database to the LSDB.

D. The algorithm will stop processing once the tree database is empty.

Suggested Answer: *B*

Community vote distribution

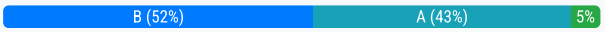

antigel8 Highly Voted 1 year, 3 months ago

Selected Answer: B

Based on AJSPR course;
Not A. because the local router moves its own local tuple into the tree database (and not into candidate database)
Not C. because SPF algorithm moves the tuple with the lowest cost from the candidate database into the tree database. (and not into the LSB)
Not D. because Step 1 (router determining cost from root to each neighbor ID) is repeated until the candidate database is empty. (and not until the tree database is empty)
Therefore the answer is B. because if a new neighbor ID appears in the tree database, any tuples in the LSDB with a router ID equal to the new tree entry's neighbor ID are moved into the candidate database.
upvoted 9 times

crisarro Most Recent 5 months ago

Selected Answer: A

As it is explained in "Advanced Junos Service Provider Routing (AJSPR) Student Guide"

"RTR-A begins by moving its own local database tupple into the candidate database. It calculates the total cost from the neighbor ID to the root, which results in 0 value. "
upvoted 3 times

boyseven777 6 months, 2 weeks ago

Selected Answer: A

Local tuple (LSDB) >> Candidate DB [SFP algo] >> Tree DB
reference: https://www.juniper.net/us/en/the-feed/summits/jumpstart-juniper/jumpstart-juniper-advanced.html
upvoted 3 times

Avee123 10 months, 1 week ago

A
The correct statement about IS-IS when it performs the Dijkstra algorithm is:

A. The local router moves its own local tuples into the candidate database.

Explanation: When IS-IS performs the Dijkstra algorithm (Shortest Path First algorithm), the local router begins by moving its own local tuples (information about its directly connected neighbors and links) into the candidate database. These tuples represent the starting point for calculating the shortest paths to other nodes in the network.

The other options are incorrect because:

B is inaccurate: the candidate database stores tuples based on cost, not just matching router IDs.
C is incorrect: tuples with the lowest cost are moved from the candidate database to the tree, not the LSDB.
D is incorrect: the algorithm stops when the candidate database is empty, not the tree database.
upvoted 3 times

ak_sisko 10 months, 2 weeks ago

Selected Answer: A

In the context of the Dijkstra algorithm, the local router includes its own local information (like its own metrics and links) into the candidate database, which represents possible next hops and their associated costs for routing decisions.

The other statements are not accurate in the context of how the Dijkstra algorithm operates in IS-IS:

B. This statement is incorrect because the algorithm does not move neighbor IDs based on a match with router IDs in the LSD (Link State Database) directly into the candidate database.
C. This is misleading; tuples with lowest cost are processed in the algorithm, but not necessarily "moved" to the LSDB (Link State Database) in this manner.
D. The algorithm does not stop processing simply when the tree database is empty; it continues until it has found the shortest path to all reachable nodes.

upvoted 3 times

easynet13 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 2 times

cname404 1 year, 5 months ago

candidate database is not the same as tree database

upvoted 2 times

StargateCommand 1 year, 5 months ago

Selected Answer: D

There is a part on "Advanced Junos Service Provider Routing On-Demand " that says: "The candidate database is not empty, so the algorithm continues".
Well, we can presume if the database is empty, algorithm stops.

upvoted 1 times

cname404 1 year, 5 months ago

candidate database is not the same as tree database

upvoted 1 times

Question #34 *Topic 1*

You have an L2VPN connecting two CEs across a provider network that runs OSPF. You have OSPF configured on both CEs.
Which two statements are correct in this scenario? (Choose two.)

A. OSPF neighborship is formed between the CEs and PEs.

B. OSPF neighborship is formed between the two CEs.

C. The CE and PE OSPF areas must match.

D. The CE and PE OSPF areas can be different.

Suggested Answer: *BD*

Community vote distribution

piipo 11 months, 1 week ago

Selected Answer: BD

BD is Correct

upvoted 1 times

Guru29 1 year, 2 months ago

No need to OSPF between PE<>CE , mean it can be different area

upvoted 3 times

Question #35 *Topic 1*

Click the Exhibit button.

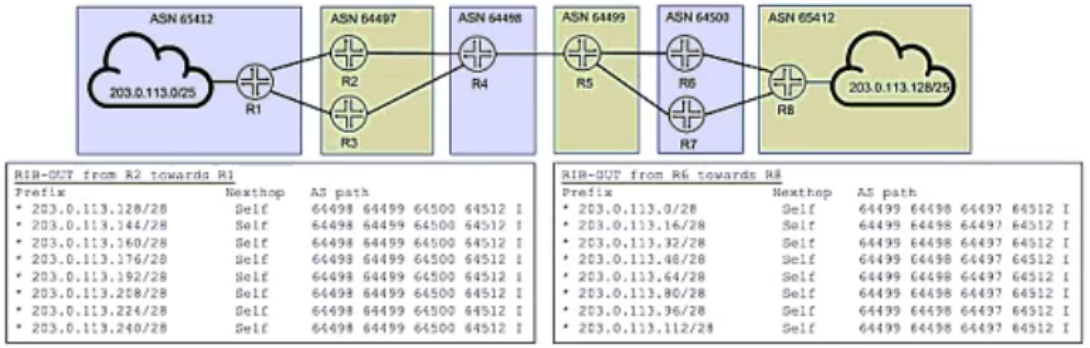

R1 and R8 are not receiving each other's routes.

Referring to the exhibit, what are three configuration commands that would solve this problem? (Choose three.)

A. Configure remove-private on advertisements from AS 64497 toward AS 64498.

B. Configure as-override on advertisement from AS 64500 toward AS 64512.

C. Configure remove-private on advertisements from AS 64500 toward AS 64499.

D. Configure loops and advertise-peer-as on routers in AS 64497 and AS 64450.

E. Configure loops on routers in AS 65412 and advertise-peer-as on routers in AS 64498.

Suggested Answer: *ABC*

Community vote distribution

ABC (44%) ACE (22%) AC (22%) 11%

None

Question #36 *Topic 1*

Click the Exhibit button.

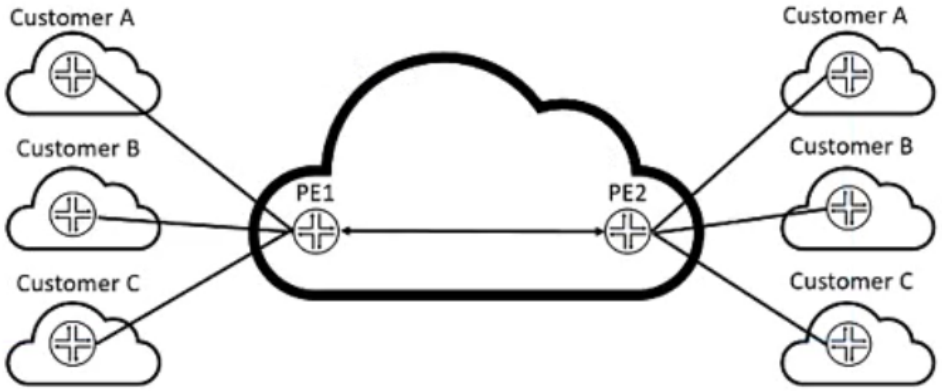

After adding Customer C to your Layer 3 VPN. you must ensure that PE2 is receiving VPN routes for all customers attached to PE1, as shown in the exhibit.

Which operational command displays this information?

A. show route table inet.0

B. show route table bgp.l3vpn.0

C. show route table customer-c.inet.0

D. show route summary

Suggested Answer: *B*

Community vote distribution

B (100%)

696501d 7 months, 2 weeks ago

Selected Answer: B

correct

upvoted 2 times

antigel8 9 months ago

Correct. 'bgp.l3vpn.0' command shows vpn routes table matching with route-target configured in the VRF.

upvoted 1 times

Question #37 *Topic 1*

Click the Exhibit button.

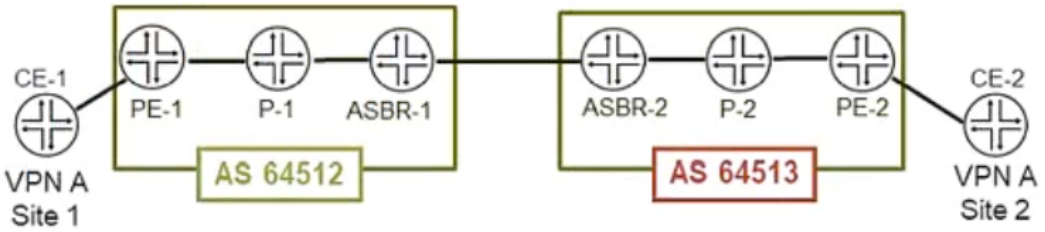

You are configuring an interprovider Option C Layer 3 VPN to connect two customer sites.
Referring to the exhibit, which three statements are correct? (Choose three.)

A. ASBR routers maintain the internal routes from its own AS, the loopback address from the other AS PEs, and the L3VPN routes.

B. ASBR routers maintain the internal routes from its own AS and the loopback addresses from the other AS PEs.

C. P routers maintain the internal routes from its own AS and the loopback address from the other AS PEs.

D. P routers only maintain the internal routes from their own AS.

E. PE routers maintain the internal routes from its own AS, the loopback address from the other AS Pes, and the L3VPN routes.

Suggested Answer: *BDE*

Community vote distribution

BDE (100%)

piipo 11 months, 1 week ago

Selected Answer: BDE

Correct

upvoted 2 times

Question #38 *Topic 1*

Click the Exhibit button.

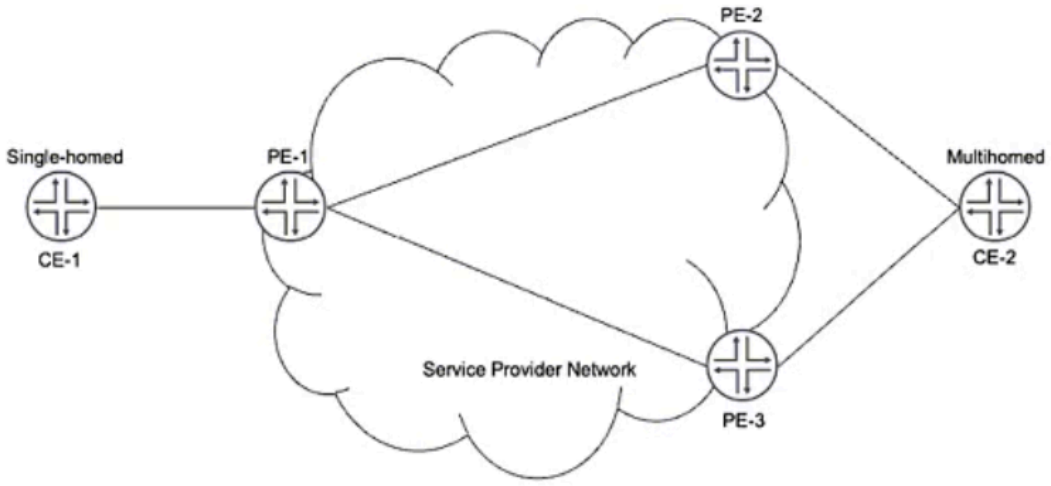

You have an EVI implemented between PE-1, PE-2, and PE-3 to allow communication between CE-1 and CE-2. CE-2 receives unicast traffic from CE-1 on both links to PE-2 and PE-3. When CE-1 sends broadcast traffic. CE-2 receives it on only one of the multihomed links.
Referring to the exhibit, which EVPN route type enables this behavior?

A. Type 2

B. Type 4

C. Type 1

D. Type 3

Suggested Answer: *B*

Community vote distribution

B (50%) C (50%)

albert1687 2 months, 2 weeks ago

Selected Answer: C

C, Type 1 is correct.

• Type 4 = Advertises multihoming (awareness)
• Type 1 = Controls who sends BUM traffic (actual behavior)

upvoted 1 times

piipo 11 months, 1 week ago

Selected Answer: B

Correct

upvoted 1 times

antigel8 1 year, 3 months ago

B is correct. EVPN route type-4 is used for electing designated forwarder

upvoted 1 times

Question #39 *Topic 1*

Click the Exhibit button.

```
Communities: target:64512:5678 mac-mobility:0x0 (sequence 4)
```

You have MAC addresses moving in your EVPN environment.

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

A. It identifies MAC addresses that should be discarded.

B. It resolves conflicting MAC address ownership claims.

C. It is advertised using a Type 2 message.

D. It helps the local PE to identify the latest advertisement.

Suggested Answer: *CD*

Community vote distribution

CD (100%)

696501d 7 months, 1 week ago

Selected Answer: CD

refer rfc7432 and section 15

upvoted 3 times

Guru29 8 months ago

C&D Correct

upvoted 2 times

Guru29 8 months ago

To determine the correct MAC address location, PE devices use the MAC mobility extended community field, as defined in RFC 7432, in the MAC advertisement route message. The MAC mobility extended community includes a static flag and a sequence number. The static flag identifies pinned MAC addresses that should not be relocated. The sequence number identifies newer MAC advertisement messages. Starting at 0, the sequence number is incremented for every MAC address mobility event. PE devices running Junos OS apply the following precedence order in determining the MAC advertisement route to use: Type 2 is Kind of ARP(MAC to IP resolution) Correct answer . A,D

upvoted 1 times

ARSE_TOP 8 months, 3 weeks ago

Correct Answer: C & D

EVPN uses a Type 2 route, tagged with a MAC mobility BGP community.

upvoted 3 times

Question #40 *Topic 1*

Click the Exhibit button.

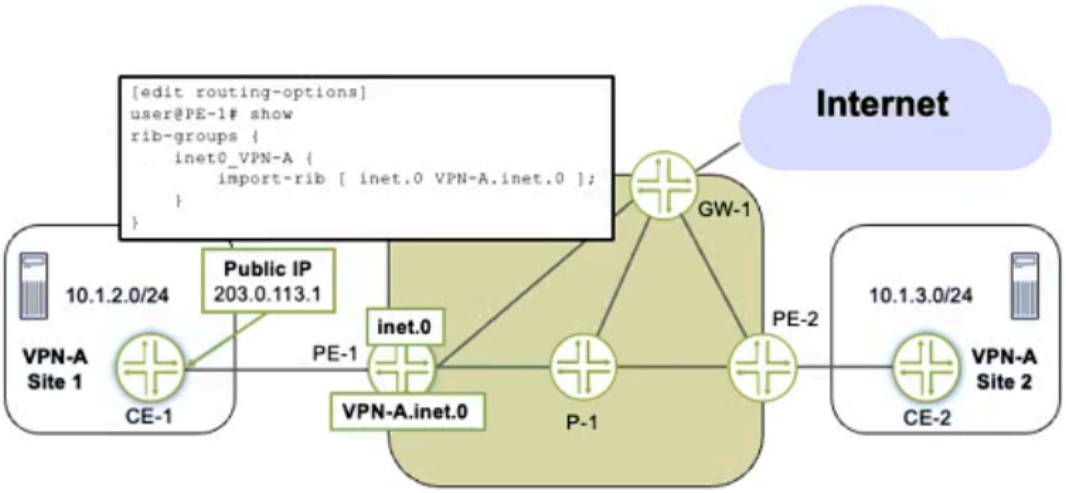

Referring to the exhibit, you must provide VRF Internet access over a single connection for VPN-A Site 1, which connects to PE-1.
Which two statements are correct in this scenario? (Choose two.)

A. You must use the RIB group to move interface routes from the inet.0 table to the VPN-A.inet.0 table.

B. You must use the RIB group to move a default route, which is learned through BGP, from the inet.0 table to the VPN-A.inet.0 table.

C. You do not need to use the RIB group to move interface routes from the inet.0 table to the VPN-A.inet.0 table.

D. You do not need to use the RIB group default route, which is learned through BGP, from the inet.0 table to the VPN-A.inet.0 table.

Suggested Answer: *AB*

Community vote distribution

BC (50%) AB (50%)

shan_online 1 week, 6 days ago

Selected Answer: AB

A. You must use the RIB group to move interface routes from the inet.0 table to the VPN-A.inet.0 table.
✓ Correct –
When a physical interface is in the global routing-instance (inet.0), and the VRF (VPN-A.inet.0) needs to access that next-hop (such as for Internet-bound traffic), you must leak the connected route using a RIB group, otherwise the VRF won't know how to reach the next-hop.

B. You must use the RIB group to move a default route, which is learned through BGP, from the inet.0 table to the VPN-A.inet.0 table.
✓ Correct –
If the default route (e.g., 0.0.0.0/0) is learned in inet.0 via BGP (e.g., from the Internet), and VPN-A needs to access the Internet, then you must use a RIB group to leak that default route from inet.0 into VPN-A.inet.0.
upvoted 1 times

1a3c9cd 4 months, 3 weeks ago

Selected Answer: BC

Answer is correct as it simply states what to do with the RIB group.

What is also required, but not asked is the static route for the NAT IP in the main table pointing to the VPN-A table.
upvoted 1 times

crramos 9 months, 1 week ago

Selected Answer: BC

The rib-goups is required and the default route with internet access should be imported from inet.0 into VPN-A and the interface routes related to VPN-A must be exported to inet.0 to complete the 2-way communication. For that reason it is B&C.
upvoted 2 times

1a3c9cd 5 months, 4 weeks ago

I follow the explanation fully.., but then the answer is A & B right?

upvoted 1 times

piipo 11 months, 1 week ago

Selected Answer: AB

AB is Correct

upvoted 2 times

Guru29 1 year, 1 month ago

Answer A & B

upvoted 3 times

Question #41 *Topic 1*

Click the Exhibit button.

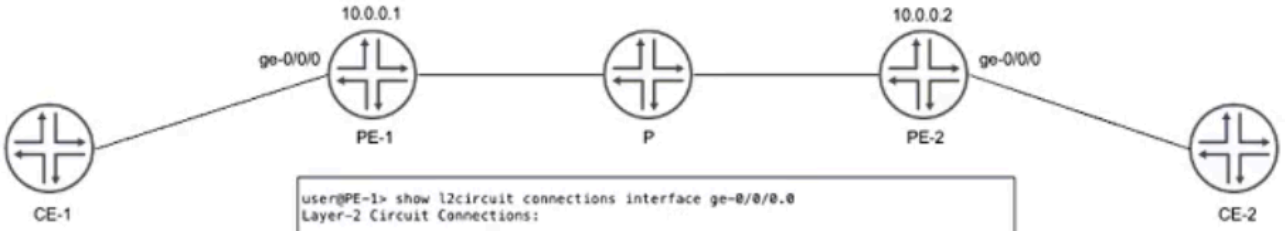


```
user@PE-1> show l2circuit connections interface ge-0/0/0.0
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.0.0.2
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/0.0(vc 5748)       rmt   OL
```

PE-1 and PE-2 are configured with LDP-signaled pseudowires to provide connectivity between CE-1 and CE-2. You notice no connectivity exists between CE-1 and CE-2.

Referring to the exhibit, which two statements describe potential causes for this fault? (Choose two.)

A. There is no LSP configured from PE-1 to PE-2.

B. Interface ge-0/0/0 on PE-1 is down.

C. The VC IDs are mismatched.

D. There is no LSP configured from PE-2 to PE-1.

Suggested Answer: *CD*

Community vote distribution

CD (80%) AD (20%)

crramos 9 months, 1 week ago

Selected Answer: AD

OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit.
Source: https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/command/show-l2circuit-connections.html

If there is VC missmatch the session will be directly down

upvoted 1 times

ak_sisko 10 months, 2 weeks ago

Selected Answer: CD

Correct answed: CD, no LSP from PE-2 to PE-1 and VC ID mismatch.

upvoted 4 times

Question #42 *Topic 1*

Which two statements are correct about reflecting inet-vpn unicast prefixes in BGP route reflection? (Choose two.)

A. Clients add their originator ID when advertising routes to their route reflector.

B. Route reflectors add their cluster ID to the AS path when readvertising client routes.

C. Route reflectors do not change any existing BGP attributes by default when advertising routes.

D. A BGP peer does not require any configuration changes to become a route reflector client.

Suggested Answer: *CD*

Community vote distribution

CD (100%)

Jenia1 11 months, 3 weeks ago

Selected Answer: CD

The given answer is correct. C and D

B is not correct because: route reflectors add the cluster ID to the route's cluster list attribute!!! , not to the AS PATH

upvoted 4 times

fukJUNIPER 1 year ago

CORRECT: B AND C

B. Route reflectors add their cluster ID to the AS path when readvertising client routes:
This statement is correct because a BGP route reflector uses a cluster ID to identify its cluster. When a route reflector readvertises routes from its clients, it appends its cluster ID to the Cluster List attribute. This helps in preventing routing loops within the cluster.

C. Route reflectors do not change any existing BGP attributes by default when advertising routes:
This statement is correct. By default, BGP route reflectors do not modify existing BGP attributes when they advertise routes to other peers. They simply reflect the routes while preserving the attributes as received from the client.

upvoted 1 times

696501d 1 year, 1 month ago

RR add cluster ID while advertising route to other RR but not by adding it to the as path. C and D are correct. B is wrong because cluster ID is an addition to the existing attribute. This is not a modification to the existing attribute.

upvoted 2 times

Guru29 1 year, 2 months ago

A- is wrong RR add originator ID
B- is correct Cluster-ID added for loop prevent
C-Wrong , Cluster-ID is a BGP attribute
D- is Correct
so correct answer B,D

upvoted 3 times

Question #43 *Topic 1*

Which three statements about IS-IS in a multi-area network are correct? (Choose three.)

A. Internal L1 PDUs are flooded to all L1 routers in other areas.

B. External L2 PDUs are only flooded to the local area's L2 routers.

C. Internal L1 PDUs are flooded to the local area's L2 routers.

D. External L2 PDUs are flooded to all L2 routers in other areas.

E. Internal L1 PDUs are only flooded to the local area's L1 routers.

Suggested Answer: *CDE*

Community vote distribution

CDE (100%)

piipo 11 months ago

Selected Answer: CDE

Correct

upvoted 3 times

Question #44 *Topic 1*

Which two statements describe PIM-SM? (Choose two.)

A. Routers without receivers must periodically prune themselves from the SPT.

B. Traffic is only forwarded to routers that request to join the distribution tree.

C. Traffic is initially flooded to all routers and an S,G is maintained for each group.

D. Routers with receivers send join messages to their upstream neighbors.

Suggested Answer: *BD*

Community vote distribution

BD (100%)

piipo 11 months ago

Selected Answer: BD

Correct

upvoted 2 times

Guru29 1 year, 2 months ago

A & C -PIM DM

so correct B&D

upvoted 1 times

Question #45 *Topic 1*

Click the Exhibit button.

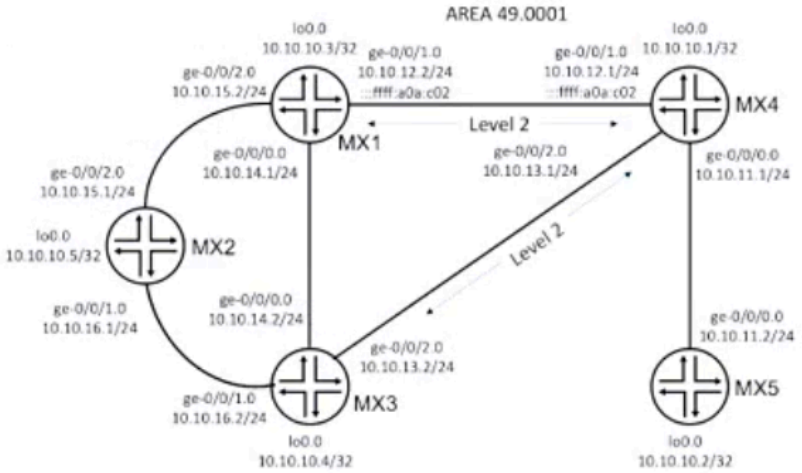


```
user@MX1> show isis interface ge-0/0/1
IS-IS interface database:
Interface L      CirID      Level 1 DR      Level 2 DR              L1/L2 Metric
ge-0/0/0/0       2          0x1         Disabled Point to Point     10/100
```

```
user@MX1> show isis database MX1 extensive | find TLV | match 10.100.12.0/24
  IP prefix: 10.100.12.0/24, Internal, Metric: defaut 63, Up
  IP extended prefix: 10.100.12.0/24 metric 63 up
```

A network is using IS-IS for routing.

In this scenario, why are there two TLVs shown in the exhibit?

A. There are both narrow and wide metric devices in the topology.

B. Both IPv4 and IPv6 are being used in the topology.

C. Wide metrics have specifically been requested.

D. The interface specified a metric of 100 for L2.

Suggested Answer: *A*

Community vote distribution

A (67%) D (17%) B (17%)

5f53711 2 months, 2 weeks ago

Selected Answer: A

Type 135 TLV is for wide metrics.

https://www.rfc-editor.org/rfc/rfc5305

(Go to number 4 or look up type 135)

upvoted 1 times

FNT_27 6 months, 2 weeks ago

Selected Answer: B

Extended prefix is a type 135 tlv that carries both ipv4 and ipv6 prefixes

upvoted 1 times

ak_sisko 10 months, 2 weeks ago

Selected Answer: A

Interface L2 metric - 100, prefix metric - 63, so both narrow and wide

upvoted 3 times

Jenia1 11 months, 3 weeks ago

Selected Answer: D

D seems to be correct.

upvoted 1 times

Question #46 *Topic 1*

You are asked to protect your company's customers from amplification attacks.
In this scenario, what is Juniper's recommended protection method?

A. ASN prepending

B. BGP FlowSpec

C. destination-based Remote Triggered Black Hole

D. unicast Reverse Path Forwarding

Suggested Answer: *B*

Community vote distribution

B (100%)

piipo 11 months ago

Selected Answer: B

B is Correct

upvoted 2 times

Question #47 *Topic 1*

Click the Exhibit button.

```
root@R1> show ospf interface extensive
Interface            State   Area            DR ID           BDR ID          Nbrs
et-0/0/33.0          DR      0.0.0.0         192.168.252.0   0.0.0.0            0
  Type: LAN, Address: 192.168.254.0, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  DR addr: 192.168.254.0, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
root@R4> show ospf interface extensive
Interface            State   Area            DR ID           BDR ID          Nbrs
et-0/0/48.0          Waiting 0.0.0.0         0.0.0.0         0.0.0.0            0
  Type: LAN, Address: 192.168.254.1, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  Priority: 128
    Adj count: 0
  Hello: 5, Dead: 20, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
et-0/0/49.0          DR      0.0.0.0         192.168.253.0   192.168.252.1      1
  Type: LAN, Address: 192.168.254.9, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  DR addr: 192.168.254.9, BDR addr: 192.168.254.8, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
root@R2> show ospf interface et-0/0/33.0 extensive
Interface            State   Area            DR ID           BDR ID          Nbrs
et-0/0/33.0          BDR     0.0.0.0         192.168.253.0   192.168.252.1      1
  Type: LAN, Address: 192.168.254.8, Mask: 255.255.255.254, MTU: 9202, Cost: 1
  DR addr: 192.168.254.9, BDR addr: 192.168.254.8, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

You have an OSPF environment. You have recently added a router called R4 that is directly connected to R1 and R2. You discover that R4 is only peering with R2.

Referring to the exhibit, how would you correct the peering?

A. Adjust the Hello Interval on R1 and R2 to match the Hello Interval on R4.

B. Change the MTU size on R1 and R2 to be 22 bytes higher than R4's MTU size.

C. Adjust the Dead Interval on R4 to match the Dead Interval on R1 and R2.

D. Adjust the Priority on R1 to be lower than the Priority on R4.

Suggested Answer: *A*

Community vote distribution

A (100%)

3574e4e Highly Voted 1 year, 2 months ago

Both hello and dead timers have to be adjusted on R4?

upvoted 5 times

dre8051 Most Recent 5 months, 3 weeks ago

Selected Answer: A

To adjust hello-interval will automatically convert the default dead-interval value into 4 times of the hello-interval that you just adjusted. However, you cannot do just by adjusting dead-interval to change the default hello-interval automatically. So basically I will select answer A.

upvoted 2 times

dre8051 5 months, 3 weeks ago

To adjust hello-interval will automatically convert the default dead-interval value into 4 times of the hello-interval that you just adjusted. However, you cannot do just by adjusting dead-interval to change the default hello-interval automatically. So basically I will select answer A.

upvoted 1 times

quraitulain 9 months ago

it should be A and C

upvoted 1 times

quraitulain 9 months ago

shouldn't it be C

upvoted 1 times

piipo 11 months ago

Selected Answer: A

I would like to combine both hello interval and dead interval, but if I had to choose between them, I would choose A.

upvoted 1 times

Guru29 1 year, 2 months ago

None of them are correct answer.

upvoted 3 times

hbstyleboy 1 year, 2 months ago

The hello interval and the dead interval optimize convergence times by efficiently tracking neighbor status. By lowering the values of the hello interval and the dead interval, you can increase the convergence of OSPF routes if a path fails. These intervals must be the same on all routing devices on a shared network.

upvoted 2 times